Sun StorageTek™ Enterprise Backup Software 7.6

Administration Guide

Sun Microsystems, Inc.
4150 Network Circle
Santa Clara, CA 95054 U.S.A.
650-960-1300

Part No. 821-1026-10
November 2009, Revision A

Send comments about this document to: docfeedback@sun.com

Sun StorageTek Enterprise Backup Softare 7.6 Administration Guide

Copyright © 2009 by EMC Corporation. All rights reserved.

Sun Microsystems, Inc. has intellectual property rights relating to technology that is described in this document. In particular, and without limitation, these intellectual property rights may include one or more of the U.S. patents listed at http://www.sun.com/patents and one or more additional patents or pending patent applications in the U.S. and in other countries.

This document and the product to which it pertains are distributed under licenses restricting their use, copying, distribution, and decompilation. No part of the product or of this document may be reproduced in any form by any means without prior written authorization of Sun and its licensors, if any.

Third-party software, including font technology, is copyrighted and licensed from Sun suppliers.

Parts of the product may be derived from Berkeley BSD systems, licensed from the University of California. UNIX is a registered trademark in the U.S. and in other countries, exclusively licensed through X/Open Company, Ltd.

Sun, Sun Microsystems, the Sun logo, Java, docs.sun.com, and Solaris are trademarks or registered trademarks of Sun Microsystems, Inc., or its subsidiaries, in the U.S. and in other countries.

All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the U.S. and in other countries. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc. NetWorker® is a trademark of EMC Corporation.

The OPEN LOOK and Sun™ Graphical User Interface was developed by Sun Microsystems, Inc. for its users and licensees. Sun acknowledges the pioneering efforts of Xerox in researching and developing the concept of visual or graphical user interfaces for the computer industry. Sun holds a non-exclusive license from Xerox to the Xerox Graphical User Interface, which license also covers Sun's licensees who implement OPEN LOOK GUIs and otherwise comply with Sun's written license agreements.

U.S. Government Rights—Commercial use. Government users are subject to the Sun Microsystems, Inc. standard license agreement and applicable provisions of the FAR and its supplements.

DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID.

Copyright © 2008 by EMC Corporation. Tous droits réservés.

Sun Microsystems, Inc. a les droits de propriété intellectuels relatants à la technologie qui est décrit dans ce document. En particulier, et sans la limitation, ces droits de propriété intellectuels peuvent inclure un ou plus des brevets américains énumérés à http://www.sun.com/patents et un ou les brevets plus supplémentaires ou les applications de brevet en attente dans les Etats-Unis et dans les autres pays.

Ce produit ou document est protégé par un copyright et distribué avec des licences qui en restreignent l'utilisation, la copie, la distribution, et la décompilation. Aucune partie de ce produit ou document ne peut être reproduite sous aucune forme, par quelque moyen que ce soit, sans l'autorisation préalable et écrite de Sun et de ses bailleurs de licence, s'il y ena.

Le logiciel détenu par des tiers, et qui comprend la technologie relative aux polices de caractères, est protégé par un copyright et licencié par des fournisseurs de Sun.

Des parties de ce produit pourront être dérivées des systèmes Berkeley BSD licenciés par l'Université de Californie. UNIX est une marque déposée aux Etats-Unis et dans d'autres pays et licenciée exclusivement par X/Open Company, Ltd.

Sun, Sun Microsystems, le logo Sun, Java, docs.sun.com, et Solaris sont des marques, ou des marques enregistrées de Sun Microsystems, Inc., ou ses filiales, aux Etats-Unis et autres pays.

Toutes les marques SPARC sont utilisées sous licence et sont des marques de fabrique ou des marques déposées de SPARC International, Inc. aux Etats-Unis et dans d'autres pays. Les produits portant les marques SPARC sont basés sur une architecture développée par Sun Microsystems, Inc. NetWorker® est un marque de fabrique de EMC Corporation.

L'interface d'utilisation graphique OPEN LOOK et Sun™ a été développée par Sun Microsystems, Inc. pour ses utilisateurs et licenciés. Sun reconnaît les efforts de pionniers de Xerox pour la recherche et le développement du concept des interfaces d'utilisation visuelle ou graphique pour l'industrie de l'informatique. Sun détient une license non exclusive de Xerox sur l'interface d'utilisation graphique Xerox, cette licence couvrant également les licenciées de Sun qui mettent en place l'interface d 'utilisation graphique OPEN LOOK et qui en outre se conforment aux licences écrites de Sun.

LA DOCUMENTATION EST FOURNIE "EN L'ÉTAT" ET TOUTES AUTRES CONDITIONS, DECLARATIONS ET GARANTIES EXPRESSES OU TACITES SONT FORMELLEMENT EXCLUES, DANS LA MESURE AUTORISEE PAR LA LOI APPLICABLE, Y COMPRIS NOTAMMENT TOUTE GARANTIE IMPLICITE RELATIVE A LA QUALITE MARCHANDE, A L'APTITUDE A UNE UTILISATION PARTICULIERE OU A L'ABSENCE DE CONTREFAÇON.



Please
Recycle

Adobe PostScript

Contents

Preface

Chapter 1 Overview

The Sun StorageTek EBS environment 40
Sun StorageTek EBS components 40
Sun StorageTek EBS features 43
Performance features 43
Ease of use 44
Scalability 45
Optional software additions 45
IPv4 to IPv6 transition 46
Common Sun StorageTek EBS tasks 46
How to add a new host 46
How to schedule a backup 47
How to configure devices 47
How to label media 47
How to view failed backups 48
Sun StorageTek EBS user interfaces 48
Sun StorageTek EBS Management Console interface 49
Sun StorageTek EBS client interface 59
Sun StorageTek EBS character-based interface 65
Sun StorageTek EBS command-line interface 65
Sun StorageTek EBS services 65
Services and programs on the Sun StorageTek EBS server . 66
Services and programs on the Sun StorageTek EBS client .. 67
Manually stopping and starting Sun StorageTek EBS services 68

Chapter 2 Backing Up Data

Scheduled backups 74
Planning a scheduled backup 74
Setting up a scheduled backup 75
Save sets 84
Scheduling predefined save sets for backup 85
The All save set 86
Manual backups 88
Performing a manual backup on Windows 88
Performing a manual backup from the command prompt . 91
Verifying backed-up data 94
Probe-based backups 95
Deduplication backups 95
BMR backups 96
Encrypting backup data 96
Compressing backup data 98
Applying compression to a scheduled backup 98
Applying compression to a manual backup 98
Special data handling for Sun StorageTek EBS clients on Windows 98
Backing up Sun StorageTek EBS Console Management data ... 99
Scheduling backups for the Sun StorageTek EBS Console database 99
Performing a manual backup of the Console database 100
Managing the size of the Console database log file 101
Backing up Windows mount points 102
Including mount points in scheduled backups 102
Performing a manual backup of a mount point and its data ... 102
Performing a manual backup of nested mount points and their data 103
Backing up the Windows Content Index Server 104
Backing up CIS on Windows 2000 or later 104
Directing Sun StorageTek EBS software to skip CIS catalog backups 105
Backing up Windows DHCP and WINS databases 105
Backing up a DHCP database 105
Backing up a WINS database 106
Windows backup and recovery notes 106
Enabling short filename support 106
Enabling hard link support 107
Recovery and case-sensitivity 107
Failed backup and recovery attempts 108
Granting full permissions for backup of Disk Quota database

108
Security settings for logging operations performed by backup operator 109
Customizing the backup command 110
Using the save command with a customized backup script 110
Using the savepnpc command with a customized backup program 116
Considerations for backing up raw partitions 120
Backing up raw partitions on UNIX 120
Backing up raw partitions on Windows 120
Backing up a mapped drive 120
Backing up access control lists 121
Support for backing up renamed directories 121
Backing up only client file indexes and the bootstrap 122

Chapter 3 Configuring Libraries and Devices

Devices and libraries 124
Support for LTO-4 hardware-based encryption 125
Storage nodes 126
Prerequisites 126
Licensing 127
Storage node configuration 127
Configuring a storage node 127
Modifying timeout attribute for storage node operations . 128
Timeouts for storage node remote devices 129
Storage node affinity 130
Dedicated storage nodes 131
Tips for troubleshooting storage nodes 131
Deduplication nodes and replication nodes 132
Dynamic drive sharing 132
Introduction to DDS 132
Block-size compatibility 134
SCSI Reserve/Release 135
Enabling DDS with NDMP 138
DDS attributes in the device properties 139
High availability and DDS 140
Licensing 141
Autodetection of libraries and devices 141
Scanning for libraries and devices 141
Configuring the Sun StorEdge Network Foundation HBA/Driver 143

Add (configure) libraries 145
How to add a library device 146
Queuing device resources for AlphaStor 149
Virtual Tape Library (VTL) configuration 150
Library reconfiguration 150
Reconfiguring a library 150
Using the jbedit command 151
Add (configure) stand-alone devices 152
Autodetecting and configuring a stand-alone tape drive . 152
Adding a stand-alone device manually 153
Configuring Sun StorageTek EBS DiskBackup 154
Support for cloud backup devices 155
Cloud backup devices compared to other device types 155
Cloud backup prerequisites and considerations 156
Cloud best practices 156
Creating and labeling a cloud storage device 157
Gathering report information on cloud backup 161
Staging with a cloud storage device 161
Cloning to a cloud storage device 161
Deleting libraries and devices 162
How to delete a library 162
How to delete a device 162
Library slots 163
How to configure library slots 163
Troubleshooting autoconfiguration failure 164

Chapter 4 Backup Groups and Schedules

Overview of Sun StorageTek EBS scheduling 166
Backup groups 166
The Sun StorageTek EBS server and time-based backup groups 167
Creating and scheduling a probe group 169
Preconfigured groups 170
Key Group attributes 170
Aborted backup groups 172
How to create a group 172
How to edit a group 172
How to delete a group 173
How to copy a group 173
Setting the backup group time interval 174
Limiting full backups when the time interval is less than 24 hours 175

Managing backup groups 176
Previewing a backup group 176
Moving clients between groups 177
Estimating save set sizes of a backup group 177
Backing up status reports 177
Generating and printing bootstrap reports 178
Backing up open files 179
Opening files owned by the operating system 179
Opening files owned by a specific application 179
Files that change during backup 180
Backing up open files with VSS 180
Schedules 180
Schedules for clients 181
Preconfigured Sun StorageTek EBS schedules 181
Backup cycles 182
Scheduling and planning considerations 184
Scheduling large client file systems 185
Key components of a schedule 186
Working with schedules 186
Overriding a client's regular backup schedule 188
Client timezone adjustment for Windows clients 188
Backup levels 190
How Sun StorageTek EBS backup levels work 191
The Sun StorageTek EBS server and backup levels 193
Backup levels for Windows SYSTEM and VSS SYSTEM save sets 196
Save set consolidation 198
Requirements 199
Save set consolidation usage 199
Limitations of save set consolidation 200
Performing save set consolidation 200
Setting up save set consolidation on a schedule 201
Directing data from a consolidation backup to a specific pool 202

Chapter 5 Browse and Retention Polices

About browse and retention policies 206
Browse policies 206
Retention policies 209
Managing the data lifecycle 213
Assigning multiple policies to a single client 215
Preconfigured policies 215

Editing a policy 216
Delete a policy 216
Snapshot policies 217
Working with snapshot policies 217
Browse and retention policies for manual backups 219
Modifying the browse and retention policy on a save set 220
Reports on browse and retention policies for save sets 221

Chapter 6 Directives

Directives overview 224
Types of local and global directives 224
Global directives 224
Sun StorageTek EBS User local directives (Windows only) 224
Local directive files 224
Creating a global directive resource 225
Editing a global directive resource 225
Deleting a global directive resource 226
Copying a global directive resource 226
Applying directives 226
Order of precedence of global and local directives 228
Local directives within the NetWorker User program 228
Set up a NetWorker User program local directive 229
Preconfigured global directive resources 230
Format of directive statements 232
Directory specifications 233
Precautions when using rawasm to back up UNIX raw partitions 238
File matching with multiple ASMs in a directive 239

Chapter 7 Sorting Backup Data

Media pools 242
Using media pools 242
Sun StorageTek EBS media pool types 243
Sorting data with media pools 243
Directing client file indexes and bootstrap to a separate media pool 245
Directing consolidated backup data to a specific media pool 246
Meeting the criteria for more than one media pool configuration 246

When no customized media pool criteria is met 248
Configuring media pools 249
Using storage devices and media pool configuration to sort data 251
Creating a media pool 251
Supporting WORM and DLTWORM tape drives 254
Working with media pools 258
Label templates 260
Using label templates 260
Using preconfigured label templates 261
Completing Label Template attributes 262
Naming label templates 264
Working with label templates 266

Chapter 8 Archiving

Overview of archiving 270
Archive requirements 270
How the Sun StorageTek EBS server archives data 271
Indexed and nonindexed archiving 272
Permissions for archiving 273
Enabling archive services for the client 273
Enabling or restricting archive access 273
Enabling public archive access 274
About archive pools 274
Preconfigured Indexed Archive pool and PC Archive pool 275
Preconfigured archive pool 275
Creating custom Archive pools 275
Archiving data procedures 275
Enabling archive services for a Sun StorageTek EBS client 276
Manually archiving data 276
Scheduling data archives 278
Retrieving archived data 281
Retrieval permissions 281
Retrieving nonindexed archives from a client on Windows 281
Retrieving nonindexed archives from a client on UNIX 282
Recovering indexed archive data from a client on Windows 283
Recovering indexed archive data from a client on UNIX 284
Retrieving archive data from the command prompt 284

Archive request management 285
Starting a scheduled archive at any time 285
Stopping a scheduled archive while in progress 285
Disabling a scheduled archive 285
Viewing details of a scheduled archive 286

Chapter 9 Cloning

Overview of cloning 288
Cloning requirements 288
Cloning performance 289
Cloning compared to a basic copy 289
Specifying retention policies for clone data 290
Volume cloning 291
Creating a clone volume 291
Viewing clone volume details 291
Save set cloning 292
Viewing the clone status of a save set 292
Cloning a save set manually 292
Cloning a save set automatically 295
Examples of manual and automatic clone operations 295
Manual clone operations 301
Viewing manual clone history 301
Stopping a manual clone operation 301
Recovering cloned data 301
Recovering a cloned save set from the command prompt 302
Recovering a save set when all cloned instances have expired 303
Cloning archived data 303
Directing clones to a special storage node 304
Storage node selection criteria for reading the clone data 304
Storage node selection criteria for writing the clone 306
Storage node selection criteria for recovering cloned data 307
Using file type devices for clone operations 308
Differences in the cloning process 308
Manual cloning with advanced file type device 308
Backup-to-tape for deduplication clients 309
Using the nsrclone command 309

Chapter 10 Staging Backups

Save set staging 314
Working with staging policies 314

Creating a staging policy 315
Editing a staging policy..... 317
Copying a staging resource 317
Deleting a staging policy 318
Staging and cloning from the command prompt..... 318

Chapter 11 Device Operations

Overview of device operations 322
File type and advanced file type devices..... 322
File type device..... 323
Advanced File Type Device (AFTD) 324
Device and volume display 336
Propagating changes to the Enable attribute from read-write to read-only..... 337
Differences between file type devices and AFTDs..... 337
Support for Sun StorageTek EBS DiskBackup configurations . 340
Sun StorageTek EBS DiskBackup configuration 341
Sun StorageTek EBS DiskBackup licensing..... 348
Compatible block size for UNIX and Microsoft Windows 349
Determine the allowable block size..... 349
Solutions to block-size compatibility problems 349
Setting the block size for a device type..... 350
Library device maintenance 351
Periodic maintenance 351
Automatic tape device cleaning..... 351
Tape alert..... 352
Service mode..... 354
Device parameter settings..... 355
Device settings in the Sun StorageTek EBS Administration interface 355
Device settings as environment variables 356
Common device interface 359
Device ordering 361
Persistent binding and naming..... 361
Detecting device ordering issues..... 362
Correcting driver ordering changes..... 363
Nonrewinding tape device usage (UNIX/Linux only) 364
Block-size mode (UNIX/Linux only)..... 365
Silos 365
Sun StorageTek EBS software interactions with a silo 365
Installing a silo..... 366

Naming conventions for silo devices 366
Releasing a silo device 369
Cleaning silo devices..... 369
Environment variables for DAS and StorageTek silos..... 371
Miscellaneous device operations..... 372
Refreshing enterprise library views on request 373
Changing the polling interval for enterprise library views 373
Stopping a device operation 373
Display device operations messages 373
Sharing libraries among Sun StorageTek EBS hosts..... 374
Sleeping periods for library tasks 375
Server Network Interface attribute 376
Tips for using libraries 376
Library notifications..... 376
Recycling compared to adding more volumes 377
Resetting a library..... 378
Using pools with libraries 379
Device calibration..... 379
Reordering tape drive numbers (Microsoft Windows only).... 379
Adding and removing media by using the library front panel 380

Chapter 12 Media Management

Storage management operations..... 384
How the Sun StorageTek EBS server uses volume labels.. 384
How the Sun StorageTek EBS server selects a volume..... 384
Auto Media Management..... 386
Using Auto Media Management..... 387
Volume operations..... 389
Using the read-only mode..... 392
Changing a volume's mode 392
Recycling volumes..... 393
Labeling volumes 393
Using barcode labels 396
Mounting and unmounting volumes 399
Using libraries with a volume import and export capability... 404
Inventorying library volumes..... 406
Working with volumes..... 407
Media handling errors 410

Reenable a device 411
Media management in a silo 411
Numbering a silo slot 411
Mounting and unmounting silo volumes 412
Labeling a silo volume 412
Using a silo with volume import/export capability 413
Barcode IDs 414
Inventorying a silo 418
Volume save sets 418
Viewing save set details in the Volume Save Sets window 418
Viewing save set details from the Save Set detail table 420

Chapter 13 Recovering Data

Choosing a recovery method 424
When to recover by file selection 424
When to recover by save set selection 424
Recovering by file selection 425
Recovering data on UNIX systems 425
Recovering data on Windows systems 426
Recovering data from the command prompt 428
Recovering an earlier version of a file 429
Relocating recovered data 433
Viewing versions of a directory or file 434
Searching for a recoverable file 434
Viewing the volumes required for data recovery 435
Recovering by save set selection 436
Recover a save set on Windows 436
Recover a save set on UNIX 437
Recover a save set from the command prompt 437
Recover individual files or directories from a save set on Windows 437
Recover individual files or directories from a save set on UNIX 438
View volumes required for save set recovery 438
Relocate recovered save set data 439
Resolve conflicts with recovered save set data 440
Recovering deduplication data 440
Recovering with BMR 441
Recovering encrypted data 441
Disaster recovery 441
Directing recoveries to another client 442

Uses for directed recovery 443
Requirements for directed recovery 444
Enabling directed recoveries..... 445
Perform a directed recovery between Windows clients from the GUI..... 448
Perform a directed recovery between UNIX clients from the GUI 449
Perform a directed recovery from the command line..... 450
Recovering online indexes for a Sun StorageTek EBS server .. 452
Recover a client file index..... 452
Restore a save set entry in the online indexes..... 454
Restoring a save set entry to the client file index only..... 455
Recovering the Console database..... 456
Recovering backed-up client files from an old Sun StorageTek EBS server 457
Recovering the Windows system configuration 458
Temporary disk space..... 460
Recovering the Windows SYSTEM from the command prompt..... 461
Point-in-time recovery of the SYSTEM and VSS SYSTEM save sets 463
Point-in-time recovery of Microsoft SQL Server or Exchange Server from earlier versions of NetWorker 464
Preparing to recover the Windows SYSTEM STATE save set . 465
Preparing to recover the SYSTEM DB save set 466
Recovering Windows volume mount points..... 467
Recovering mount points 467
Recovering a mount point and its data 467
Recovering Windows DHCP and WINS databases..... 470
Recover a DHCP database 470
Recover a WINS database 470
Restoring Windows Content Index Server on Windows 2000 or later 470

Chapter 14 Reporting

Gathering report data..... 474
Enable the gathering of report data 475
Disable the gathering of report data..... 475
Data retention and expiration policies 475
Setting expiration policies for data retention 477
Report categories 477

Report types 478
Basic reports 479
Drill-down reports 480
Customized reports 480
Configuring reports 481
How to configure a report 482
Date and time formats 482
Input formats 483
Viewing reports 484
Interactive mode 485
Document mode 487
Interactive and document mode chart types 489
View reports 494
Restricting report views 494
Background processing of reports 495
Preconfigured reports 495
Sun StorageTek EBS backup statistics reports 495
Sun StorageTek EBS backup status reports 502
Recovery reports 508
Deduplication Statistics reports 514
Event reports 517
Host reports 521
User reports 522
Device reports 524
Cloud backup and recover reports 526
Inactive files report 529
Customizing and saving reports 531
Naming reports 532
Saved file ownership and deleted users 532
Sharing reports 533
Sharing a report 533
Exporting reports 534
Exporting a report 534
Exporting non-ASCII characters 535
Command line reporting 535
The command line reporting program 536
Java runtime environment 537
Printing reports 537

Chapter 15 Monitoring

Overview of monitoring 540
Events 540

Polling for System Events 541
Enabling or disabling the Capture Events option 541
Viewing events..... 542
Event priorities..... 542
Working with notes..... 543
Working with annotations 544
Monitoring window 545
Viewing the Monitoring window 546
Viewing backup group information 547
Viewing session information 550
Other panel information..... 551
Monitoring the status of devices 552
Monitoring device operations 553
Stop an operation..... 555
Viewing log files from the console..... 555
Viewing archive information..... 556

Chapter 16 Console Server Management

Managing Sun StorageTek EBS Console server users 562
Login authentication for Console users 562
Adding Console users and assigning to Console roles..... 565
Restricting a Console user's view of managed servers..... 566
Granting Console users privileges on a NetWorker server 568
Resetting user preferences 569
Deleting a user 570
Editing user properties 570
Resetting the administrator password (native authentication only)..... 571
Moving the Console server 572
Setting system options 575
Setting a system option..... 575
Individual User Authentication 578
Setting environment variables 580
Setting environment variables..... 581
Accessing the Console Configuration Wizard..... 582
Sun StorageTek EBS Console server maintenance tasks 582
Changing the service port used by the Sun StorageTek EBS Console database 583
Changing database connection credentials 584
NMC server IP address/hostname updates..... 585
Displaying international fonts in non-US locale environments.....

585

Chapter 17 Sun StorageTek EBS Server Management

Enterprise 588
- Enterprise components 588
- Organizing Sun StorageTek EBS servers 588
- Viewing the enterprise 590
- Managing various servers in the enterprise 590
- Configuring remote access for a Retrospect host 593
- Managing folders in the enterprise 593
- Adding or deleting multiple servers by using a hostname file 595

Configuring a Sun StorageTek EBS server 599
- Set up the server 599

Configuring ConnectEMC 600
- Configure ConnectEMC in the ConnectEMC Console 601
- Configure ConnectEMC in NMC 602

Enabling BMR server support 604

Parallelism and multiplexing 604
- Parallelism 604
- Multiplexing 608

Managing server access 609
- Granting access to manage Sun StorageTek EBS servers from Console server 610
- Administrator privileges 610

Sun StorageTek EBS User Groups 613
- Preconfigured User Groups 613
- Customizing privileges 614
- Working with User Group resources 616

Server communication issues within Microsoft Windows 618
- Name resolution 618
- Backup Operators group 618
- Dynamic Host Configuration Protocol 618
- Backup and Recover Server service 619

Notifications 619
- Preconfigured notifications 619
- Customizing notifications 622
- Logging event notifications 626
- Creating a custom notification 627
- Editing a notification 628
- Copying a notification 628
- Deleting a custom notification 629

Indexes 629
Characteristics of the online indexes 629
Automated index activities 630
Checking online indexes 631
Viewing information about the indexes 631
Index save sets 632
Querying the media database 633
Cross-checking client file indexes 633
Refreshing index information 634
Client file index locations 634
Managing the size of the online indexes 636
Querying deduplication save sets using mminfo 640
Updating Sun StorageTek EBS clients by using the software distribution feature 641
Message log management 641
Setting the UNIX startup script to trim log files 642
Savegroup job information 643
Internationalization 644
Interoperability with previous releases of Sun StorageTek EBS 645
Display issues 645
Maximum path and save set length 646
Group completion messages 646
Locale-specific configuration issues on UNIX/Linux 647

Chapter 18 Sun StorageTek EBS Client Management

Sun StorageTek EBS client overview 650
Client configuration 650
Creating a client 650
Editing a client 650
Copying a client 651
Changing a client name 651
Deleting a client 652
Recovering a deleted client 652
Editing a client NSRLA database 653
Creating a client probe 655
Associating a probe with a Client resource 656
Deduplication client 656
Configuring a client for BMR support 657
Lockbox pass phrase management 657
Set the Datazone pass phrase for a Sun StorageTek EBS server 658

Error messages and error handling for the Datazone pass phrase 658
Sun StorageTek EBS authentication 658
Strong authentication (nsrauth) 658
Authentication for backwards compatibility (oldauth) 659
Access privileges for authentication configuration 659
Specifying the minimum authentication strength between hosts 660
Maintaining Sun StorageTek EBS local host authentication credentials 663
Maintaining local host Peer resources 668
Creating a custom certificate and private key for a host.... 670
Multiple clients from the same computer 671
Redefining a file system into multiple client and save set instances 671
Defining a client and save set combination 672
Scheduled backups of non-ASCII files or directories 672
Adding or changing the Sun StorageTek EBS servers for a client.. 673
Editing the servers file 674
Client priority 674

Chapter 19 Cluster Support

Cluster terminology 678
Cluster-aware and cluster-unaware Sun StorageTek EBS servers . 678
Cluster licensing 680
Backing up data on UNIX cluster environments 680
Configuring a scheduled save from a physical cluster client... 680
Configuring a scheduled save from a virtual cluster client..... 680
Performing a manual backup from a virtual cluster client 682
Configuring a virtual client to back up to a local storage node 684
Directing a virtual client backup to a local storage node... 685
Configuring a virtual client to be recovered from a local storage node 685
Directing a virtual client recovery from a local storage node .. 686
Recovering a virtual client's data by using the currechost keyword 686

Configuring an external client to a virtual server 686
Recovering data in a UNIX cluster environment 687
Recovering data from a private disk that belongs to a physical client 687
Recovering data from a shared disk that belongs to a virtual client 687
Recovering the Console server database 688
Tracking scheduled saves 688
Forcing save sets to use a specific index 689
Path-ownership rules 689
Using a physical storage node and virtual server on UNIX platforms 691
Using autochangers and stand-alone tape devices 692
Configurations for autochangers and stand-alone tape devices 693
DDS within a cluster 698
Considerations when using jbconfig for a virtual server .. 703
Client Backup Configuration Wizard support requirements .. 704
NetWorker cluster performance issues 706
Editing the cluster cache timeout attribute 706
Setting NetWorker environment variables in a Sun cluster 708

Chapter 20 Support for VMware

Introduction to VMware support 710
Licensing Sun StorageTek EBS support for VMware 710
Configuring Sun StorageTek EBS clients for virtual machines 710
Considerations for VCB backups 710
Task 1: Configure the Hypervisor resource 712
Task 2: Configure the VM proxy host 713
Task 3: Configure a virtual client for backups 717
Automatic discovery of VMware environments 722
Performing on-demand auto-discovery of VMware environments 723
Notifications of changes to VMware environments 723
Visual Representation of VMware environments 724
Virtualization node hierarchical display of the VMware environment 724
Graphical display of the VMware environment 725
Recovering VMware Consolidated Backups 728
Performing a file-based recovery on the local host 728
Performing a file-based recovery using CIFS share 728
Performing a file-based recovery using directed recovery 729

Performing a single step recover of the full virtual machine ... 730

Chapter 21 Sun StorageTek EBS Server Support for NDMP

Overview of NDMP ... 736
Licensing ... 736
NDMP connection support ... 736
Supported NDMP configurations ... 737
Local backup by using an NDMP tape server ... 737
Backup by using Sun StorageTek EBS storage node support ... 739
Three-party backup with NDMP tape server ... 742
Requirements for and limitations of NDMP support ... 743
Requirements for Sun StorageTek EBS software with NDMP support ... 744
Limitations of the Sun StorageTek EBS software with NDMP support ... 744
Sun StorageTek EBS storage node support for NDMP clients ... 747
Vendor-specific information ... 750
Configuring the Sun StorageTek EBS server for NDMP operations ... 750
Configuring library support for NDMP ... 751
Configuring an NDMP device resource ... 757
Configuring an NDMP Client resource ... 758
Additional considerations for NDMP operations ... 762
Performing NDMP backups ... 763
Tape server support for NDMP client backup ... 763
Performing a local backup by using the NetWorker User program ... 764
Performing an NDMP backup from the command prompt ... 765
Performing a three-party backup from the command prompt 766
Performing DAR backups ... 767
Storage node support for NDMP client backup ... 768
Performing NDMP recoveries ... 770
Storage node support for NDMP client recovery ... 771
Performing a save set recovery ... 771
Destructive save set recoveries ... 773
Performing a directory-level recovery ... 774
Performing a file-by-file recovery ... 775

Performing a directed recovery 777
NDMP cloning 780
Cloning between storage nodes 780
Cloning functionality 781
NDMP index and bootstrap cloning 781
NDMP Path-to-Tape cloning for regular save sets 782
Ancillary product information 785
Filers that act as data servers 785
Network Appliance 786
SnapImage software 801
Filers that act as tape servers only 802

Chapter 22 SNMP Module

SNMP traps 808
Configuring NetWorker SNMP notifications 808
Command line options for nsrtrap 809
Modifying preconfigured NetWorker SNMP notification 809
Creating NetWorker SNMP notifications 810
Configuring SNMP management software 811
Sun StorageTek EBS SMI Network Management Private Enterprise Code 811
Receiving traps in the SNMP network management software 811

Chapter 23 Support for Microsoft Automated System Recovery

Microsoft Automated System Recovery 814
Microsoft ASR documentation 814
Sun StorageTek EBS support for ASR 814
Sun StorageTek EBS ASR save set 815
Usage overview 815
Network connection names 816
ASR Limitations and special considerations 816
FAT16 partitions are not supported 816
OEM recovery CDs are not supported 817
Vendor-specific drivers must be installed after Windows installation 817
ASR backup 817
How to include the ASR or VSS ASR DISK save set in a scheduled backup 818
How to include ASR or VSS ASR DISK save set in a manual backup 818

Creating an ASR disk 819
Prerequisites 819
Create an ASR disk locally 820
Creating an ASR disk by using directed recovery 821
Posting ASR disk creation task 822
Using the ASR disk to recover a Sun StorageTek EBS client 823
Requirements for an ASR recovery 823
Performing an ASR recovery 824
Components that require special handling after an ASR recovery 825
Verifying the Sun StorageTek EBS client recovery 826

Chapter 24 Volume Shadow Copy Service

Overview of VSS 830
VSS and the backup process 831
The importance of writers 833
Controlling VSS from Sun StorageTek EBS software 833
Controlling VSS from the Administration window 834
Controlling VSS from the Sun StorageTek EBS client 836
Control VSS from the command-prompt 838
Globally disabling VSS 838
VSS commands 839
Handling MSDE, SQL Server, and Microsoft Exchange databases 840
Authoritative restores of the Active Directory Application Mode (ADAM) and DFS Replication writers 841

Chapter 25 DiskXtender Data Manager File System Support

Supported configurations 844
Path information 844
Permissions 844
DiskXtender Data Manager file system overview 844
File data in a DXDM file system 845
Backup of DXDM file systems 846
Aborted backups 848
Recovery of DXDM file systems 849
Initiating a recovery 850
Restoring deleted files and previous file versions 851
File system synchronization 851
Automatic synchronization 852
Manually synchronizing a file 852

Chapter 26 Support for Avamar Deduplication

Sun StorageTek EBS deduplication node 856
Deduplication backups 856
Replication of deduplication backups 857
Schedules for deduplication clients 857
Backup levels for deduplication clients 857
Retention policies 859
Directives 859
Creating a deduplication node and a replication node 860
Avamar setup 860
NetWorker setup 862
Cloning and deduplication 867
Backup-to-tape for deduplication clients 868
Recovering deduplicated data 868
Viewing deduplication events in NMC 869
Disaster recovery 870
Deduplication save sets 870
Deleting deduplication save sets 871
Querying deduplication save sets using mminfo 871
Refreshing deduplication node information 872
Tips for troubleshooting deduplication nodes 873

Chapter 27 Support for HomeBase

Introduction to HomeBase Support (Bare Metal Recovery) 876
BMR Recovery 876
NetWorker BMR HomeBase Server 876
Installation Notes 876
BMR backups 877
Configuring a client for BMR support 877
Recovering with BMR 878
Excluding files for BMR recovery 880

Chapter 28 Licensing

Product licenses 882
Base enabler 882
Evaluating the NetWorker software 882
How to apply an evaluation enabler code 883
Deleting an enabler code 883
Evaluation enabler alert 884
Permanently licensing the NetWorker software 884
NetWorker License Manager 884

Entering an enabler code 885
Deleting an enabler code..... 885
Entering an authorization code..... 885
Changing the License Manager server 886
Licensing Sun StorageTek EBS software components 886
Archiving licensing..... 886
Bare Metal Recovery (BMR) licensing 887
Client connection licenses..... 887
Cluster licensing..... 887
DDS licensing 888
HomeBase server/agent 889
NDMP licensing..... 889
NetWorker Modules..... 889
Sun StorageTek EBS cloud backup option 890
Sun StorageTek EBS storage node licensing 890
Virtual environments simplified licensing..... 890
Virtual Edition Client Connection license..... 891
License Conformance Summary 893
How to obtain NetWorker license information 895
Querying the local server..... 896

Chapter 29 Troubleshooting

Before contacting technical support 902
Determining the version of Sun StorageTek EBS software running on a client..... 903
Displaying diagnostic mode attributes..... 903
Viewing log files..... 904
Rendering log files in the current locale at runtime 904
How to view log files with the nsr_render_log program .. 906
Viewing log files from remote host machines 906
Log files from previous releases of Sun StorageTek EBS... 907
Filtering log file information displayed by nsr_render_log 907
Locating savegroup job logs..... 907
Sun StorageTek EBS functionality issues..... 909
Backup and recovery..... 909
Unsynchronized client clock errors..... 911
Backups fail to start when daylight savings time change occurs..... 912
Clone ID timestamp does not reflect the time the cone was created..... 912
Backups fail to stop..... 912

Memory usage when browsing large save sets 913
Memory usage and nsrjobd 913
Media position errors encountered when auto media verify is enabled 913
PACKET RECEIVE BUFFER and NO ECB counters increase . 914
The scanner program marks a volume read-only 914
The scanner program requests an entry for record size 914
Limitations for groups containing a bootstrap 914
Index recovery to a different location fails 915
Illegal characters in configurations 915
Hostname aliases 915
Directory pathname restrictions 916
Failed recover operation directly after new installation 916
Recovering files from an interrupted backup 917
Backup of a new client defaults to level full 917
Non-full backup of Solaris files with modified extended attributes 917
Renamed clients cannot recover old backups 918
Client file index errors 918
Cannot use the Console interface to stop the savegrp command 919
Aborting a recovery 919
RPC error 920
Error message when relocating data 920
Desktop heap size limitation 920
The All save set and duplicate drive serial numbers 920
Disk label errors 921
Cannot print bootstrap information 921
Server index not forced 921
Copy violation 922
Xview errors 922
Converting sparse files to fully allocated files 923
Backing up large sparse files 923
The mminfo -N command is case-sensitive regarding save set names 924
Devices and autochangers 925
Additional attributes in the Autochanger resource 926
Maintenance commands 926
Autodetected SCSI jukebox option causes server to stop responding 926
Autochanger inventory problems 926
Destination component full messages 927

Tapes do not fill to capacity 927
Tapes get stuck in drive when labelling on Linux Red Hat platform 929
Increasing the value of Save Mount Time-out for label operations 929
Server cannot access autochanger control port 930
Nonrewinding device requirement 931
Scanner command behaves differently with adv_file type device 931
Sleep times required for TZ89 drive types 932
Managing optical drives with Solaris 9 932
Message displayed when CDI enabled on NDMP or disk FTD 933
Verifying firmware for switches and routers 933
Commands issued with nsrjb on a multi-NIC host fail 934
SCSI reserve/release with dynamic drive sharing 934
Device ordering issues 936
Sun StorageTek EBS locale and code set support 936
Issues with fonts when using X servers on Linux 936
Resource database notes 936
Viewing resources 937
Repairing resource database corruption 937
Network and server communication errors 938
General issues 938
Firewall issues 940
UNIX communication issues 941
Troubleshoot IP errors 941
Binding to server errors 945
These messages indicate that the Sun StorageTek EBS services **nsrd**, **nsrexecd**, **nsrindexd**, **nsrmmd**, and **nsrmmdbd** might not be running. Saving remote file systems 945
Microsoft Windows communication issues 946
Sun StorageTek EBS archiving and retrieval 950
Remote archive request from server fails 950
Multiple save sets appear as a single archive save set 951
Wrong archive pool is selected 951
Second archive request does not execute 951
The nsrarchive program does not start immediately 951
Archive request succeeds but generates error when nsrexecd is not running 951
Filenames with carriage return not allowed 952
Storage nodes 952

Storage node affinity errors 952
Storage node timeout errors 952
Console error messages and corrective actions 954
Console log files 957
The install log 957
The gstd log 959
Console troubleshooting notes and tips 959
Making sure the Console server is running 959
Enabling Java script 961

Appendix A SYSTEM and VSS SYSTEM Save Sets

SYSTEM save sets 964
Components of the SYSTEM STATE save set 964
Components of the SYSTEM FILES save set 965
Components of the SYSTEM DB save set 966
Components of the SHAREPOINT save set 968
Components of the ASR save set 969
VSS SYSTEM save sets 970
Components of the VSS SYSTEM BOOT save set 970
Components of the VSS SYSTEM FILESET save set 970
Components of the VSS SYSTEM SERVICES save set 971
Components of the VSS USER DATA save set 972
Components of the VSS OTHER save set 972
Components of the VSS ASR DISK save set 973

Appendix B Firewall Support

Overview 976
Configuring the port ranges in Sun StorageTek EBS software 976
Connection ports range 976
Service ports range 976
Using the nsrports command to update port ranges 977
Using the Console to update the port ranges 978
Using nsradmin to update the port ranges 978
Ports needed in the service ports range for Sun StorageTek EBS 7.3.x and later 979
The Sun StorageTek EBS server 980
Sun StorageTek EBS storage node 981
Sun StorageTek EBS client 982
Ports needed in the service ports range for Sun StorageTek EBS 7.2.x. 982
Sun StorageTek EBS server 982
Sun StorageTek EBS storage node 983

Sun StorageTek EBS client 983
Ports used by Sun StorageTek EBS add-on products 983
Sun StorageTek EBS Console 983
AlphaStor 985
Client Backup Configuration Wizard 986
NDMP 986
Sun StorageTek EBS modules 986
Example of a firewall configuration 987
Additional ports that may be required 988
Access control to local host resources 988
Firewall port enhancement 989

Appendix C Backing Up and Restoring a Microsoft DFS

Overview of a Microsoft Distributed File System (DFS) 992
Domain-based DFS versus registry-based DFS 992
DFS junctions 992
DFS backups and restores 993
DFS topology information 993
DFS backup consideration 993
Configuring a scheduled DFS backup 994
Restoring a DFS 995

Appendix D Additional Features of the Microsoft Windows Server

NetWorker Module for Microsoft Applications 998
Active Directory 998
Backing up Active Directory 998
Recovering Active Directory 999
Preparing to recover the system state on a domain controller 999
Performing a nonauthoritative restore 1000
Performing an authoritative restore 1001
Encrypting file system 1004
Event logs 1006
Internet information server 1006
Windows registry 1007
Backward compatibility with earlier Sun StorageTek EBS releases 1007
Files and folders not backed up in Windows 2000 1008
Sparse files 1008
Terminal server licensing 1008
Backing up TSL data 1008
Recovering TSL data 1009

Windows Change Journal 1010
Sun StorageTek EBS support for Change Journal..... 1010
Configuring Sun StorageTek EBS software to use the Change Journal 1012
Advanced Configuration and Power Interface 1014
Sun StorageTek EBS support for ACPI 1014

Appendix E UNIX and Linux Platform-Specific Notes

Solaris 1018
Support for Solaris zones 1018
Using the All keyword for the Save Set attribute on Solaris 10 1018
Sun StorageTek EBS executables not found for Solaris client. 1019
Extended file attribute data included in Save Set File Size attribute..... 1020
The inquire command and Solaris 10..... 1020
Linux 1021
Backup considerations for Linux raw disk partitions..... 1021
Configure Linux operating system to detect SCSI devices 1021
Configuration requirements for the inquire command .. 1021
Linux Journaled file system support 1022
HP-UX 1023
Autochanger installation on an HP-UX system..... 1023
How to test the device driver and device file installation..... 1025
The “Hardware address not found” message..... 1026
Errors from unsupported media in HP tape drives 1027
Unloading tape drives on an HP-UX storage node..... 1027
Symbolic link entries in the fstab file..... 1027
Font warnings when exporting display to a Solaris computer 1028
AIX 1029
STK-9840 drives attached to AIX 1029
IRIX 1030
IRIX displays a “no space left in tape” message..... 1030

Appendix F MAC OS X Support

Support for Mac OS X 1032
Mac OS X metadata support..... 1032
Supported file systems..... 1032

Known limitations 1032
Mac OS X backup considerations 1032
Scheduling a Sun StorageTek EBS client backup on Mac OS X 1033
Performing a manual backup on Mac OS X..... 1035
Recovering files and directories on Mac OS X 1035
Task 1: Browse backed-up Mac OS X data 1035
Task 2: Recover individual files or directories 1036

Appendix G Direct SCSI Backup and Recover

Introduction to direct SCSI backup and recover 1038
System requirements 1038
Unsupported features..... 1039
Performing direct SCSI backup 1040
Backing up data on a Symmetrix BCV device 1040
Backing up data on a raw device..... 1042
Backing up data from the command line 1042
Performing direct SCSI recover 1044
Recovering data to a Symmetrix BCV device 1044
Recovering data to a raw device..... 1046
Licensing 1047

Glossary

Index

Preface

The *Sun StorageTek Enterprise Backup Software 7.5 Administrator's Guide* contains information on how to configure and manage the Sun StorageTek Enterprise Backup™ software.

This guide also provides information about how to use and administer the Sun StorEdge EBS software when it is installed in a typical networked environment. For more detailed technical information about Sun StorEdge EBS commands, refer to the online Sun StorEdge EBS manual (man) pages after you install the software.

The information in this guide is intended for system administrators who are responsible for installing software and maintaining the servers and clients on a network. Operators who monitor the daily backups may also find this manual useful.

This guide is part of the Sun StorageTek EBS® documentation set.

This guide provides information on how to configure and manage the Sun StorageTek EBS™ software.

Install the Sun StorageTek EBS software on the appropriate server and clients before using the information presented in this guide. Refer to the *Sun StorageTek EBS Installation Guide* for installation instructions.

Before You Read This Book

In order to fully use the information in this document, you must have thorough knowledge of the topics discussed in these books:

- *Sun StorageTek Enterprise Backup 7.5 Installation Guide*
- *Sun StorageTek Enterprise Backup 7.5 Release Notes*

Using UNIX Commands

This document might not contain information on basic UNIX® commands and procedures such as shutting down the system, booting the system, and configuring devices.

See one or more of the following for this information:

- *Solaris Handbook for Sun Peripherals*
- AnswerBook2™ online documentation for the Solaris™ operating environment
- Other software documentation that you received with your system

Organization

To find information about the user interface and basic tasks, see:

- Chapter 1, "Overview,"

To set up backups and configure devices for backups, see:

- Chapter 2, "Backing Up Data,"
- Chapter 3, "Configuring Libraries and Devices,"

To find detailed information about managing backup data, see:

- Chapter 4, "Backup Groups and Schedules,"
- Chapter 5, "Browse and Retention Polices,"
- Chapter 6, "Directives,"
- Chapter 7, "Sorting Backup Data,"
- Chapter 8, "Archiving,"
- Chapter 9, "Cloning,"
- Chapter 10, "Staging Backups,"

To find detailed information about devices and media, see:

- Chapter 11, "Device Operations,"
- Chapter 12, "Media Management,"

To recover data, see:

- Chapter 13, "Recovering Data,"

To create reports and monitor Sun StorageTek EBS activity, see:

- Chapter 14, "Reporting,"
- Chapter 15, "Monitoring,"

To manage hosts including the Console server, Sun StorageTek EBS server, and Sun StorageTek EBS clients, see:

- Chapter 16, “Console Server Management,”
- Chapter 17, “Sun StorageTek EBS Server Management,”
- Chapter 18, “Sun StorageTek EBS Client Management,”

To manage Sun StorageTek EBS software in a cluster environment.

- Chapter 19, “Cluster Support,”

To find licensing information, see:

- Chapter 28, “Licensing,”

To manage a Sun StorageTek EBS server in a Network Data Management Protocol (NDMP) or Simple Network Management Protocol (SNMP) environment, see:

- Chapter 21, “Sun StorageTek EBS Server Support for NDMP,”
- Chapter 22, “SNMP Module,”

To find detailed information about how Sun StorageTek EBS software works with various supported platforms and platform features, see:

- Chapter 20, “Support for VMware,”
- Chapter 23, “Support for Microsoft Automated System Recovery,”
- Chapter 25, “DiskXtender Data Manager File System Support,”
- Chapter 24, “Volume Shadow Copy Service,”
- Chapter 26, “Support for Avamar Deduplication,”
- Chapter 27, “Support for HomeBase,”
- Appendix A, “SYSTEM and VSS SYSTEM Save Sets,”
- Appendix B, “Firewall Support,”
- Appendix C, “Backing Up and Restoring a Microsoft DFS,”
- Appendix D, “Additional Features of the Microsoft Windows Server,”
- Appendix E, “UNIX and Linux Platform-Specific Notes,”
- Appendix F, “MAC OS X Support,”

To find detailed troubleshooting information, see:

- Chapter 29, “Troubleshooting,”

To find terms used in this guide, see "Glossary".

Conventions used in this guide

EMC uses the following conventions for notes, cautions, warnings, and danger notices.

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



CAUTION

A caution contains information essential to avoid data loss or damage to the system or equipment. The caution may apply to hardware or software.

Typographical conventions

The following type style conventions are used in this guide:

bold	• User actions (what the user clicks, presses, or selects) • Interface elements (button names, dialog box names) • Names of keys, commands, programs, scripts, applications, utilities, processes, notifications, system calls, services, applications, and utilities in text
italic	• Book titles • New terms in text • Emphasis in text
`Courier`	• Prompts • System output • Filenames • Pathnames • URLs • Syntax when shown in command line or other examples
`Courier, bold`	• User entry • Options in command-line syntax
`Courier italic`	• Arguments in examples of command-line syntax • Variables in examples of screen or file output • Variables in pathnames
<>	Angle brackets for parameter values (variables) supplied by user.
[]	Square brackets for optional values.
|	Vertical bar symbol for alternate selections. The bar means or.
...	Ellipsis for nonessential information omitted from the example.

Accessing Sun Documentation Online

A broad selection of Sun system documentation is located at:

`http://www.sun.com/products-n-solutions/hardware/docs`

A complete set of Solaris documentation and many other titles are located at:

`http://docs.sun.com`

Sun Welcomes Your comments

Sun is interested in improving its documentation and welcomes your comments and suggestions. You can email your comments to Sun at:

`docfeedback@sun.com`

Please include the part number of your document in the subject line of your email.

1 Overview

This chapter covers these topics:

- The Sun StorageTek EBS environment 40
- Sun StorageTek EBS features 43
- IPv4 to IPv6 transition 46
- Common Sun StorageTek EBS tasks 46
- Sun StorageTek EBS user interfaces 48
- Sun StorageTek EBS services 65

The Sun StorageTek EBS environment

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

The Sun StorageTek EBS software is a cross-platform, client-server application that provides the ability to remotely manage all Sun StorageTek EBS clients and servers from a web-enabled, graphical interface.

Sun StorageTek EBS components

Figure 1 on page 43 illustrates the main components in a Sun StorageTek EBS environment.

Console server

All Sun StorageTek EBS servers and clients are managed from the Sun StorageTek EBS Console server. The Console server also provides reporting and monitoring capabilities for all Sun StorageTek EBS servers and clients.

Console user interface

The Console server is accessed through a graphical interface that can be run from any computer that has a supported web browser and Java Runtime Environment (JRE). The Sun StorageTek EBS Installation Guide provides information on supported web browsers and 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 Sun StorageTek EBS client.

Sun StorageTek EBS server

Sun StorageTek EBS servers provide services to back up and recover data for the Sun StorageTek EBS client computers in a datazone.

Datazone

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

Sun StorageTek EBS storage node

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

Sun StorageTek EBS client

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

Sun StorageTek EBS deduplication node

A Sun StorageTek EBS deduplication node is an EMC Avamar® server that stores deduplicated backup data. The initial backup to a deduplication node should be a full backup. During subsequent backups, the Avamar infrastructure identifies redundant data segments at the source and backs up only unique segments, not entire files that contain changes. This reduces the time required to perform backups, as well as both the network bandwidth and storage space used for backups.

Avamar server installation is separate from Sun StorageTek EBS installation, and is performed by EMC Customer Support.The Avamar server must be configured as a Sun StorageTek EBS deduplication node, and the Avamar server must be available when:

- A deduplication client resource is created.
- The Avamar server receives backup data from Sun StorageTek EBS deduplication clients.

Chapter 26, “Support for Avamar Deduplication” provides detailed information on deduplication.

Virtual environments

Sun StorageTek EBS clients can be created for virtual machines for either traditional backup or VMware Consolidated Backup (VCB). Additionally, the Sun StorageTek EBS 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.

Sun StorageTek EBS bare metal recovery (BMR) HomeBase Server

A Sun StorageTek EBS BMR server is an EMC HomeBase Server that manages and restores operating system and server state data,

collected as profiles. The HomeBase Server is the collection point for profiles produced by HomeBase Agents. HomeBase Agents are automatically installed with the Sun StorageTek EBS client software. The profile data includes:

- hardware configurations
- operating system levels
- system tuning
- network connections
- security
- storage layouts.

The HomeBase Agent captures this information on a scheduled basis with a NetWorker save set, and sends the profile results to a HomeBase Server for management and analysis.

A BMR recovery enables the restoration or migration of a server from one hardware type to another. This type of recovery begins with the HomeBase Server and the profile data that is gathered and includes the file system and application data restore from the Sun StorageTek EBS server through the Console server.

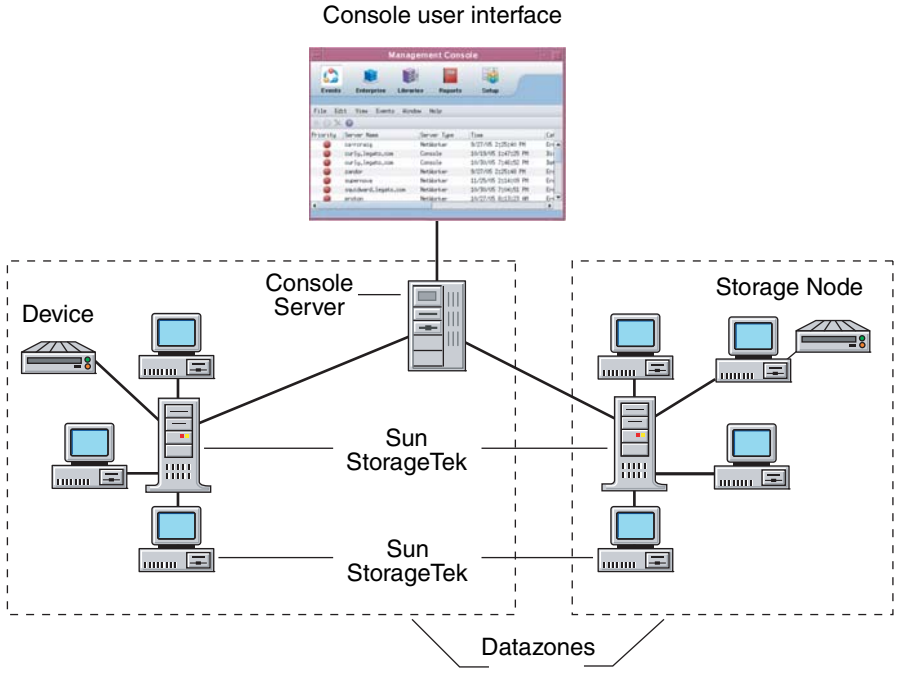
HomeBase Server installation is separate from Sun StorageTek EBS server installation. The HomeBase Server must be configured and available when:

- A BMR client resource is created.
- BMR backup data (profiles) are generated from Sun StorageTek EBS clients.

Chapter 27, "Support for HomeBase" provides detailed information on BMR. The EMC HomeBase product documentation provides

details about HomeBase Server installation, profile management, and recovery specifics.

Figure 1 Sun StorageTek EBS components



Sun StorageTek EBS features

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

Performance features

Standard Sun StorageTek EBS 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.
- Session management, which enables one to control the number of save streams per device to maximize the performance of each device.
- Backup to file devices and optional subsequent staging to near-line or offline volumes.
- Backup to a private or public cloud configuration.

Ease of use

Sun StorageTek EBS 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 Sun StorageTek EBS tasks and functions.
- Use wizards to set up the following Sun StorageTek EBS items:
  - Client resources
  - Common Console configuration tasks
  - LDAP user authentication
- Administer and configure Sun StorageTek EBS 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 Sun StorageTek EBS command reference guide, which provides information similar to the UNIX man pages.
- Take advantage of the automatic media management feature to enable the Sun StorageTek EBS 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 Powerlink® website to find answers to common questions.
- Automatically discover and view a graphical map of virtual environments.

- Set up Sun StorageTek EBS Console server authentication to an external LDAP v3.1 compliant server.

Scalability

Sun StorageTek EBS 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 Sun StorageTek EBS 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 Sun StorageTek EBS 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

IPv4 to IPv6 transition

Internet Protocol version 6 or IPv6 is designed to replace the current Internet Protocol which is IPv4. IPv6 fixes a number of problems in IPv4, such as the limited number of available IPv4 addresses. It also adds many improvements to IPv4 in areas such as routing and network autoconfiguration. IPv6 is expected to gradually replace IPv4, with the two coexisting for a number of years during a transition period. The *Sun StorageTek EBS Installation Guide* provides information about using Sun StorageTek EBS in an IPv6 environment.

Common Sun StorageTek EBS tasks

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

- "How to add a new host" on page 46
- "How to schedule a backup" on page 47
- "How to configure devices" on page 47
- "How to label media" on page 47
- "How to view failed backups" on page 48

"Sun StorageTek EBS Management Console interface" on page 49 has information about starting the Sun StorageTek EBS 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 **Create Host** dialog box, type in the new host name and click **Next**. The **Select Managed Application** dialog box appears.
6. Click **Next**. The **Manage NetWorker** dialog box appears.
7. Make sure the following checkboxes are enabled:
   - For NetWorker: **Capture Events** and **Gather Reporting Data**

- For Retrospect®: **Capture Events**
- For Avamar: **Capture Events**

8. Click **Finish**.

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 **Sun StorageTek EBS** 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 74 provides information about scheduling backups.

How to configure devices

To configure devices 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 **Sun StorageTek EBS** 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**.

Chapter 3, "Configuring Libraries and Devices" information about configuring media.

How to label media

To label tapes from the Sun StorageTek EBS 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 **Sun StorageTek EBS** 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 12, "Media Management" provides information about labeling tapes.

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 **Sun StorageTek EBS** 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 15, "Monitoring" provides information about viewing backup status information

Sun StorageTek EBS user interfaces

The Sun StorageTek EBS application consists of these user interfaces:

- "Sun StorageTek EBS Management Console interface" on page 49
- "Sun StorageTek EBS client interface" on page 59
- "Sun StorageTek EBS character-based interface" on page 65

- "Sun StorageTek EBS command-line interface" on page 65

Sun StorageTek EBS Management Console interface

The interface for Sun StorageTek EBS Management Console, also called the Sun StorageTek EBS Console, consists of two main windows:

- Console window
- Administration window

Console window

When Sun StorageTek EBS software is started, the Console window appears as shown in Figure 2.

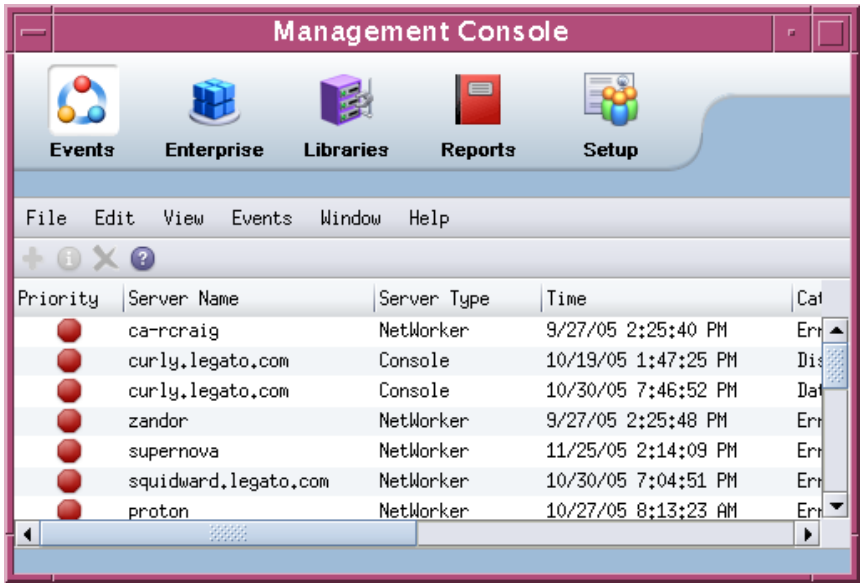


Figure 2 **Console window**

The Console window is the first point of access for Sun StorageTek EBS tasks. Table 1 on page 50 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
	Events	View important messages about all Sun StorageTek EBS servers that have been added as Enterprise applications, as well as the Console server, and Avamar server. "Events" on page 540 provides information about managed events.
	Enterprise	Select a Sun StorageTek EBS server to manage and monitor the server and its backup clients. The Enterprise window allows you to open the Administration window for a Sun StorageTek EBS server.
	Libraries	View summary information and manage libraries for all Sun StorageTek EBS servers. The Administration window can also be opened from this location.
	Reports	Configure and view Console reports. Chapter 14, "Reporting." provides information about reports.
	Setup	Control administrative functions: • User management — Add, edit, and delete Console user accounts, restrict user views of servers. "Managing Sun StorageTek EBS Console server users" on page 562 provides information about user management. • License management — Manage Sun StorageTek EBS licenses. "Licensing" on page 881 provides information about license management.

Administration window

Sun StorageTek EBS servers are managed through the Administration window as shown in Figure 3.

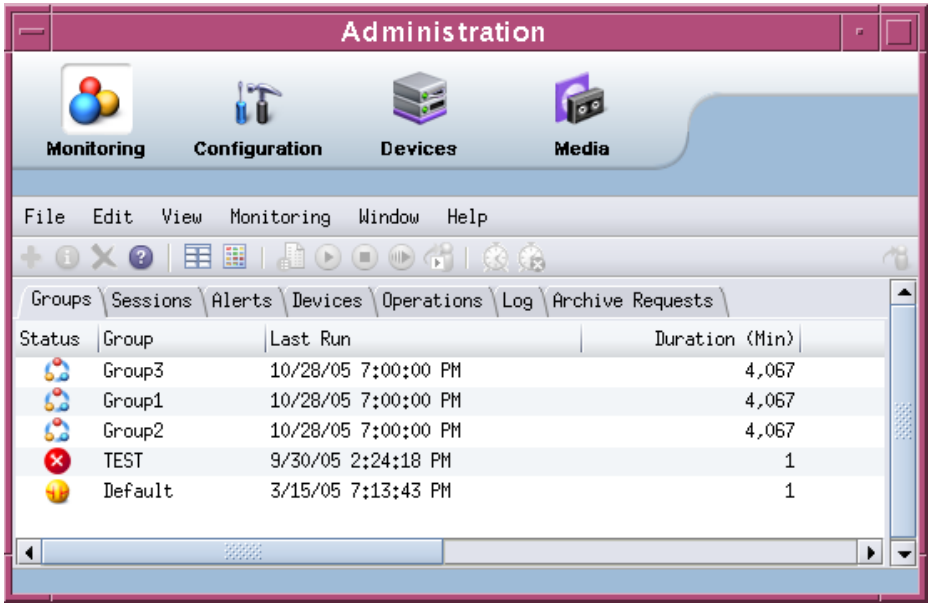


Figure 3 **Administration window**

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

Table 2 on page 52 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 Sun StorageTek EBS 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 Sun StorageTek EBS server and its resources such as clients, backup schedules, and policies. For example, you can create a backup schedule, add Sun StorageTek EBS clients, and apply the backup schedule to several Sun StorageTek EBS clients.
	Devices	Add, configure, and operate single or multiple devices, libraries, and silos for the Sun StorageTek EBS 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.

Starting Console for the first time

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

To open Console for the first time:

1. Start a web browser session.
2. Type the following command to specify the Console server URL:

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

   where:

   - *server_name* is the name of the computer where the Console server component is installed.
   - *http_service_port* is the IP port for the embedded HTTP server. The HTTP port is specified during installation. The default HTTP port is 9000.

3. From the **Welcome** page, click **Start Sun StorageTek EBS Management Console**.
4. From the **Security Warning** screen, click **Start** to run Sun StorageTek EBS Console.
   - If the JRE was *not* already installed on the system, a prompt to install it appears. Follow the on-screen instructions to install JRE.
   - Once the JRE has been installed, the Java Web Start Desktop Integration dialog box appears.
5. Complete the **Java Web Start Desktop Integration** banner dialog box by selecting one of these options:
   - To place a shortcut on the desktop, click **Yes**.
   - To decline having a shortcut placed on the desktop, click **No**.
   - To have the option to decide later, click **Ask Later**.
6. In the **Sun StorageTek EBS Management Console Login** banner dialog box, type the username and password.

   Note: The default password for the administrator account is *administrator*. For added security, change the password during the first login session. "Resetting the administrator password (native authentication only)" on page 571 provides information if the administrator password needs to be changed.

7. Click **OK** and the **Console** window appears.

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 52 provides information.
- Double-click the **Sun StorageTek EBS 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 Sun StorageTek EBS server and open the Administration window:

1. From the Console window, click **Enterprise**.
2. Add one or more Sun StorageTek EBS 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 Sun StorageTek EBS server is running and click **Next**.
   d. Select **Sun StorageTek EBS** for the type of application to be managed.
   e. Click **Finish**.
   f. Repeat for all Sun StorageTek EBS 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 14, "Reporting" 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 **Ctl + 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**.

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 56
- "Client and group management in the Administration window" on page 57
- "Library operations in the Devices window" on page 57
- "Copy and paste tabular information to operating system clipboard" on page 58

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.

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

Client and group management in the Administration window

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

- Copy selected clients to a new Sun StorageTek EBS 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 Sun StorageTek EBS group to a new group. Select a group in the directory tree and move clients from the Client Summary table to another Sun StorageTek EBS group.
- Move selected clients to a new Sun StorageTek EBS 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 Sun StorageTek EBS 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 Sun StorageTek EBS 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.

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

Multiple library devices and slots

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

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

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

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

- Mount
- Load without mount
- Withdraw
- Label
- Inventory
- Remove (STL and EMC AlphaStor®only)

Setting user interaction preferences

Depending on the window button that was selected (see Table 2), 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.

Sun StorageTek EBS 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 Sun StorageTek EBS client immediately. Scheduled backup and archive operations are set up through the Console interface. For information about the Console interface, see "Sun StorageTek EBS Management Console interface" on page 49.

Windows client interface

The Sun StorageTek EBS User program shown in Figure 4 is the Windows client interface.

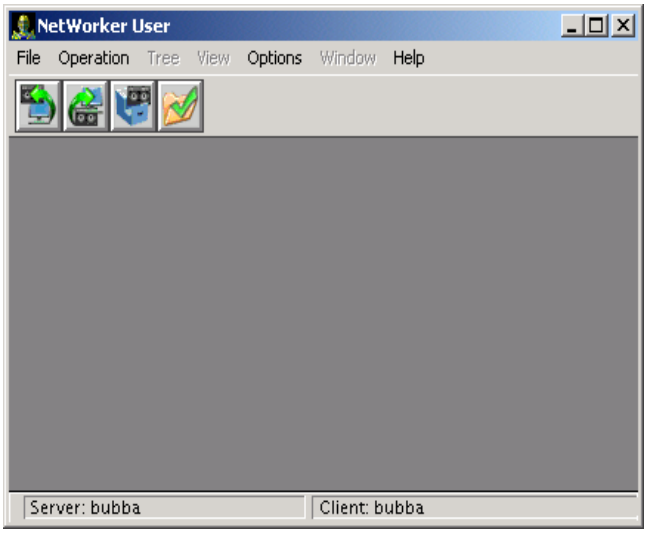


Figure 4 Sun StorageTek EBS User program

Starting the User program

To start the Sun StorageTek EBS User program, perform one of the following:

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

Note: The Sun StorageTek EBS client package must be installed on the host where you start the Sun StorageTek EBS User program.

To start the Sun StorageTek EBS User program, you must belong to the appropriate Windows groups. Table 3 lists the groups that you must belong to in order to run the Sun StorageTek EBS User program.

The Backup Operators and Administrators groups are the local and remote Microsoft security groups.

Table 3 Sun StorageTek EBS 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 Sun StorageTek EBS User program has a toolbar with buttons for common User program tasks. Table 4 describes the function of each button.

Table 4 Sun StorageTek EBS User toolbar functions

Button	Name	Function
	Backup	Starts an manual (unscheduled) backup of the client's data to a Sun StorageTek EBS 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 Sun StorageTek EBS User program appears when you select one of the following:

- A toolbar button
- A Backup, Recover, Archive, Verify, or Local Directive command from the Sun StorageTek EBS User File menu

The browse window, shown in Figure 5 displays the directory tree of the file system that is being browsed.

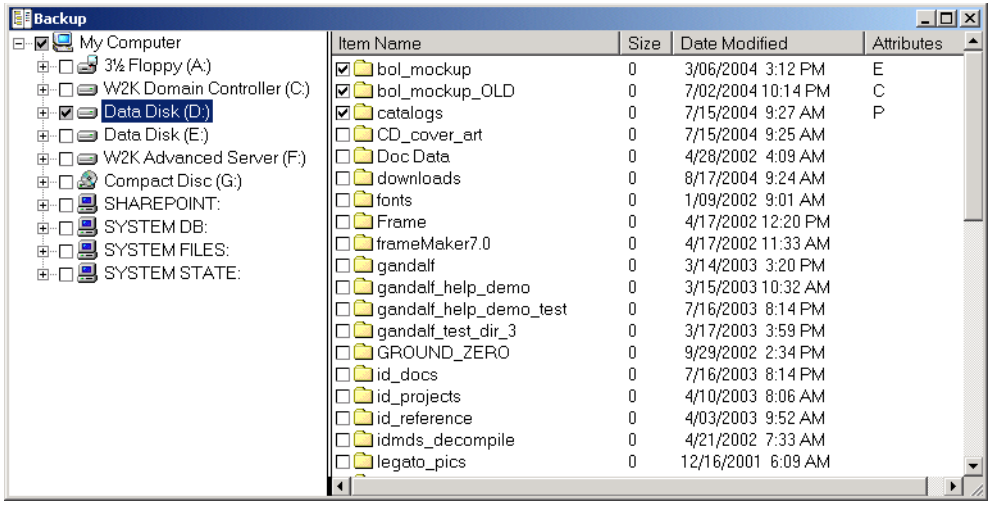


Figure 5 **Example of the browse window**

Note: 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 96 provides more information.
- **E** — The item is marked for password-protection and encryption, using the PW2 ASM. "Encrypting backup data" on page 96 provides more information.
- **C** — The item is marked for compression. "Compressing backup data" on page 98 provides more information.

Connecting to a Sun StorageTek EBS server

A typical user that runs the Sun StorageTek EBS User program needs to connect to the Sun StorageTek EBS 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 Sun StorageTek EBS server.

Before the Sun StorageTek EBS User program can connect to a Sun StorageTek EBS server, the client computer must be set up as a Client resource on that Sun StorageTek EBS server. "Task 6: Create a backup Client resource" on page 81 provides information about creating a Client resource.

To connect to a Sun StorageTek EBS server:

1. From the **Operation** menu, select **Change Sun StorageTek EBS Server**.
2. In the **Change Server** dialog box, select a server from the list of available Sun StorageTek EBS servers. If the server is not listed, do one of the following:
   - Click **Update List** to search the network for available Sun StorageTek EBS servers.
   - Type the server's hostname.
3. Click **OK**.

UNIX client interface

The UNIX version of the Sun StorageTek EBS client interface includes a graphical interface program, **nwrecover**, for performing recoveries of backed up data and retrieving archived data.

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

The Recover program

Use the **nwrecover** program to recover backed-up files.

- To specify the Sun StorageTek EBS server to connect to, type the **-s** *servername* option from the command line:

  nwrecover -s *servername*

  The /servers file, located in /nsr/res/servers, contains an entry for each available server.

  - If the **-s** option is not entered and there is only one server detected, that server will be connected automatically.

- If there are no servers detected, or if there is more than one server available, the **Change Server** dialog box appears, which allows you to choose the server.

◆ To specify the Sun StorageTek EBS client, use the **-c** *client* option:

```
nwrecover -s servername -c client
```

If the **-c** option is not entered, the local computer is used by default.

Figure 6 displays the **nwrecover** program.

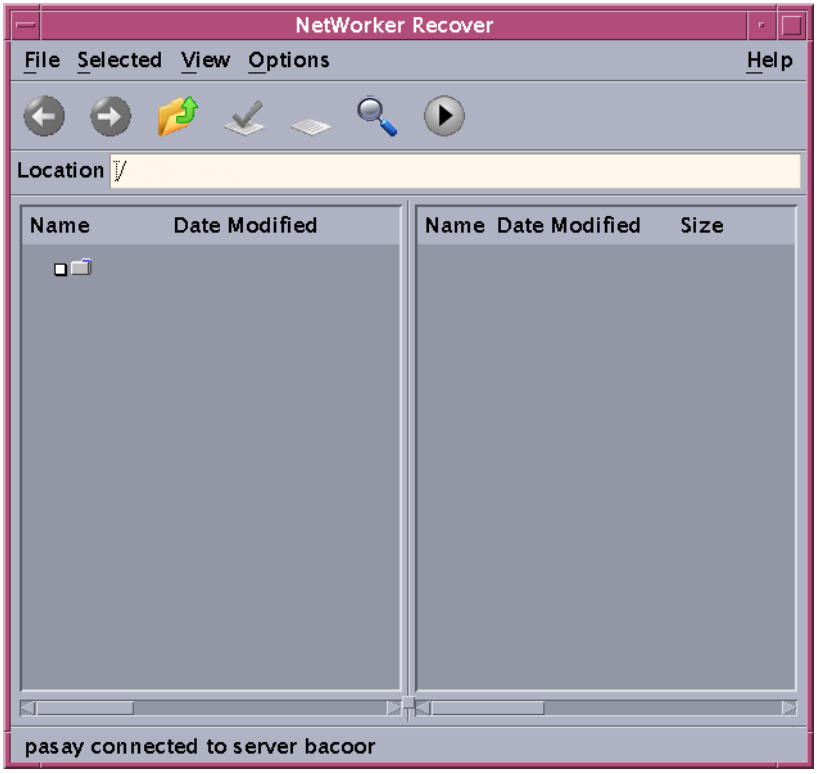


Figure 6 The nwrecover Program

Sun StorageTek EBS character-based interface

Use the Sun StorageTek EBS character-based interface (**nsradmin**) to perform configuration and management tasks for a Sun StorageTek EBS 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.

Sun StorageTek EBS 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.

Sun StorageTek EBS services

This section provides information and procedures that use services and programs for the NetWorker server and NetWorker Client. It also describes how to start and stop these services. For more information about:

- Main Sun StorageTek EBS 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 Sun StorageTek EBS server

Table 5 describes the services and programs that provide Sun StorageTek EBS software functionality.

Table 5 Services or programs on the server (1 of 2)

Service or program	Function
ansrd	Monitors an active **recover** session. It is called by **nsrd** in response to a **recover** session.
gstd	Known as the Generic Services Toolkit (GST), controls other services provided by the Console server.
hagentd	Manages the NetWorker host agent modules and maintains the NetWorker host agent resource database.
jobsd	Manages monitoring and other communication between the NetWorker server and the NetWorker console.
nsrd	Monitors an active **save** or **recover** program session. This is an agent process spawned by **nsrd** in response to a **save** or **recover** session. It (**nsrd**) is also the master service that controls other services on the NetWorker server, clients, and storage nodes.
nsrck	Checks the consistency of the client file index. It is invoked by **nsrd** whenever the consistency of the client file index needs to be confirmed.
nsrim	Automatically manages the server media database. It is invoked by **nsrmmdbd** when it starts up at the end of the **savegrp** program. It removes both aborted and expired save sets once every 24 hours after a savegroup is completed (if Volume Recycle is set to Auto). It (**nsrim**) is also called by **nsrd** when a user removes the oldest backup cycle.
nsrindexd	Provides a method for inserting entries into the client file index that is based on information passed by the **save** program.
nsrlogd	Allows logging of security audit messages from any Sun StorageTek EBS machine.
nsrmmd	Provides device support, generates mount requests, and multiplexes save set data during a multiclient backup. It writes the data sent by **save** to storage media. An additional **nsrmmd** service is started while mount requests are pending.The **nsrd** service starts one **nsrmmd** service for each enabled device. Forwards storage information to **nsrmmdbd** for recording in the Sun StorageTek EBS server media database.
nsrmmdbd	Provides media database management services to the local **nsrd** and **nsrmmd** services and records entries in the media database. This is the media management database service.

Table 5 Services or programs on the server (2 of 2)

Service or program	Function
nsrmmgd	Manages media library operations. It provides an RPC-based service that manages all jukebox operations on behalf of the **nsrd** service. The **nsrd** service starts one and only one instance of **nsrmmgd** on the Sun StorageTek EBS server as needed.
nsrlcpd	Provides **nsrmmgd** with a uniform interface to supported libraries. One **nsrlcpd** service is started by the **nsrmmgd** service for each virtual library instance. The **nsrlcpd** services run on the storage node that controls the library that the service manages.
nsrpm	Works with the energy-saving features of Windows to keep NetWorker components running properly during power-saving modes. If it is not required, it can be disabled. **Note:** This service is for the Windows OS only.
savegrp	Runs a group of Sun StorageTek EBS clients through the **save** process.

Services and programs on the Sun StorageTek EBS client

The **nsrd** service calls on the Sun StorageTek EBS client service, **nsrexecd**, and several programs on the client when a scheduled or on-demand backup request is received. A temporary server agent, the **ansrd** service, starts on the Sun StorageTek EBS server to monitor the progress of the backup session.

Table 6 describes the services or programs on the Sun StorageTek EBS client.

Table 6 Services or programs on the client(1 of 2)

Service or program	Function
nsrlogd	Allows logging of security audit messages from any Sun StorageTek EBS machine.
nsrmmd	Provides device support, generates mount requests, and multiplexes save set data during a multiclient backup. It writes the data sent by **save** to storage media. An additional **nsrmmd** service is started while mount requests are pending.The **nsrd** service starts one **nsrmmd** service for each enabled device. Forwards storage information to **nsrmmdbd** for recording in the Sun StorageTek EBS server media database.
nsrexecd	Authenticates and processes the Sun StorageTek EBS server remote execution requests and executes the **save** and **savefs** programs on the client.

Table 6 Services or programs on the client(2 of 2)

Service or program	Function
recover	Browses the Sun StorageTek EBS server client file index and restores the specified file to primary disk storage.
save	Sends specified files in a multiplexed data stream to the Sun StorageTek EBS server for backup to media by **nsrmmd** and entry in the client file indexes by **nsrindexd** and the media database by **nsrmmdbd**. A level-based system is used to save only the files modified since some previous save (a partial save).
savefs	Gathers all the necessary information, such as file system and level, and sends it to **savegrp**. The **savegrp** program then requests that the **save** program sends the data to the Sun StorageTek EBS server.

Manually stopping and starting Sun StorageTek EBS services

Sun StorageTek EBS hosts might need to be manually stopped and restarted in conjunction with configuration and troubleshooting activities. Table 7 lists the main services that are required for each type of Sun StorageTek EBS host.

Table 7 Main services for Sun StorageTek EBS hosts

Sun StorageTek EBS host	Sun StorageTek EBS service
Console server	**gstd, nsrexecd, nsrlogd**
Sun StorageTek EBS server	**nsrd**, **nsrexecd**, **nsrindexd**, **nsrjobd**, **nsrlogd**, **nsrmmd**, **nsrmmdbd**, **nsrmmgd**, **nsrlcpd** (**nsrlcpd** is running only if a library is enabled, configured, and communication has been established)
Sun StorageTek EBS client	**nsrexecd**
Sun StorageTek EBS storage node	**nsrexecd**, **nsrmmd**, **nsrlcpd** (**nsrlcpd** is running only if a library is enabled, configured, and communication has been established)
License Manager (available only from EMC)	**lgtolmd**

Note: The service **nsrmmd** is present only if one or more devices are enabled. For releases prior to 7.3, more than one **nsrexecd** may be running. For release 7.3 and later, only one **nsrexecd** will be running.

Stopping and starting the Console server

This section describes how to stop and start the Console server service.

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 Sun StorageTek EBS client is running.

   The Sun StorageTek EBS Remote Exec Service (**nsrexecd**) should have a status of **Started**. If the service has not started:

   a. Right-click **Sun StorageTek EBS 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:
   - **/etc/init.d/gst stop**

Start the Console server on UNIX

To start the Console server:

1. Log in as root.
2. Verify that the Sun StorageTek EBS 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 Sun StorageTek EBS host on UNIX" on page 71 provides information about starting the client.

3. Start the Console server by typing one of the following commands:

   - **/etc/init.d/gst start**

Stopping and starting a Sun StorageTek EBS server, client, or storage node

This section describes how to manually stop and start the services for a Sun StorageTek EBS server, client, or storage node.

Stop a Sun StorageTek EBS host on Windows

To stop a host:

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 **Sun StorageTek EBS Remote Exec Service** (**nsrexecd**) and select **Stop**.

Start a Sun StorageTek EBS host on Windows

To start a host:

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.

   - Sun StorageTek EBS client or storage node: Right-click the **Sun StorageTek EBS Remote Exec Service** (**nsrexecd**) and select **Start**.

Stop a Sun StorageTek EBS host on UNIX

To stop the Sun StorageTek EBS services, log in as root and type the following command:

```
nsr_shutdown
```

Start a Sun StorageTek EBS host on UNIX

To start Sun StorageTek EBS services, log in as root and type the appropriate startup command listed in Table 8.

Table 8 Sun StorageTek EBS startup commands

Operating system	Startup command
Solaris, Linux, IRIX	**/etc/init.d/networker start**
HP-UX	**/sbin/init.d/networker start**
HP Tru64 UNIX	**/sbin/init.d/NSRstartstop start**
AIX	**/etc/rc.nsr**

Stop a Sun StorageTek EBS host on Mac OS X

To stop the Sun StorageTek EBS 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 Sun StorageTek EBS
```

Start a Sun StorageTek EBS host on Mac OS X

To start the Sun StorageTek EBS 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 Sun StorageTek EBS
```

2

Backing Up Data

This chapter covers these topics:

- Scheduled backups 74
- Save sets 84
- Manual backups 88
- Verifying backed-up data 94
- Probe-based backups 95
- Deduplication backups 95
- BMR backups 96
- Encrypting backup data 96
- Compressing backup data 98
- Special data handling for Sun StorageTek EBS clients on Windows 98
- Backing up Sun StorageTek EBS Console Management data 99
- Backing up Windows mount points 102
- Backing up the Windows Content Index Server 104
- Backing up Windows DHCP and WINS databases 105
- Windows backup and recovery notes 106
- Customizing the backup command 110
- Considerations for backing up raw partitions 120
- Backing up a mapped drive 120
- Backing up access control lists 121
- Support for backing up renamed directories 121

Scheduled backups

The Sun StorageTek EBS server backs up client data regularly by using scheduled backups. They are preferred over the "Manual backups" on page 88 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 74
- "Setting up a scheduled backup" on page 75

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 9 on page 74. This table also maps each requirement to specific Sun StorageTek EBS features.

Table 9 Accounting department backup requirements(1 of 2)

Requirement	Sun StorageTek EBS feature	More information
Backups occur at the same time.	Backup Schedule Backup Group	"Task 1: Set up a schedule for backups" on page 77 "Task 2: Set up a group for backup clients" on page 78
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 79
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 79
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 80 (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 84

Table 9 Accounting department backup requirements(2 of 2)

Requirement	Sun StorageTek EBS feature	More information
To avoid unnecessary backups, do not back up files with a .tmp extension.	Directives	"Task 5: Set up directives for special processing" on page 81
The same files and folders are backed up on each accounting computer.	Client resource	"Task 6: Create a backup Client resource" on page 81
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 79 "Task 6: Create a backup Client resource" on page 81

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 Sun StorageTek EBS servers and clients in a stand-alone or cluster environment.

Note: 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 Sun StorageTek EBS server and client privileges, or have root (UNIX) or Administrator (Windows) privileges.

  - Have Configure Sun StorageTek EBS privileges on the Sun StorageTek EBS server where the scheduled backup is to be configured.
- The Sun StorageTek EBS 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 910 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 75 are met. Also check the NetWorker daemon log for additional details. "Viewing log files" on page 904 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 77
- "Task 2: Set up a group for backup clients" on page 78

- "Task 3: Set up policies for quick access and long term storage" on page 79
- "Task 4: Set up a label template to identify volumes" on page 80
- "Task 5: Set up directives for special processing" on page 81
- "Task 6: Create a backup Client resource" on page 81
- "Task 7: Set up a pool to sort backup data" on page 84

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 609 provides information about getting permissions.

Note: Appendix F, "MAC OS X Support" provides information about backing up Sun StorageTek EBS clients on Mac OS X

Task 1: Set up a schedule for backups

A schedule can be applied to each client backup. Chapter 4, "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 Sun StorageTek EBS server.
- Sort data to specific backup volumes. To sort data, groups are used in conjunction with backup *pools*.

Chapter 4, "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 624.
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 Sun StorageTek EBS 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 206 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 Sun StorageTek EBS server creates a unique label for each volume by applying a label template. Chapter 7, "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 6, "Directives" provides information about directives.

Note: 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 230 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 /tmp, type:

```
<< "/" >>
skip: *
```

7. Click **OK**.

Task 6: Create a backup Client resource

A client is both a physical computer with Sun StorageTek EBS client software installed on it and a Sun StorageTek EBS *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 Sun StorageTek EBS 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 671 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 Sun StorageTek EBS 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. From the **Directive** attribute, select a directive from the list, if applicable.
10. 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 /bin/log and all of the data under the directory named /accounting, the entries would look similar to:

```
/bin/log
/accounting
```

Type **All** to back up *all* client data. For Microsoft Windows operating systems, the SYSTEM or Volume Shadow Copy Service (VSS) SYSTEM save sets that determine the client system's state should be backed up on a regular basis.

The following sections provide information about the respective topics:

- "Save sets" on page 84 provides information on using the All save set and the SYSTEM or VSS SYSTEM save sets.
- "Backing up a mapped drive" on page 120 provides information of how backing up a mapped drive.
- "Backing up Windows mount points" on page 102 provides information on backing up mount points and nested mount points.
- "Scheduled backups of non-ASCII files or directories" on page 672 provides information on using non-English paths in the **Save Set** attribute.

Note: 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 230 provides information.

---

11. From the **Group** attribute, select a group from the list.
12. From the **Schedule** attribute, select a schedule from the list.
13. 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 121 provides more information about this feature.
14. 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 7, "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 242 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 Sun StorageTek EBS 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.

Note: A save set is defined when a Client resource is created. "Task 6: Create a backup Client resource" on page 81 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 Sun StorageTek EBS client.

For Microsoft Windows Server 2003 that has no VSS client license or has VSS disabled, as well as for Windows XP Professional and Windows 2000, these predefined save sets are available:

- All
- SYSTEM STATE (Windows only)
- SYSTEM DB (Windows only)
- SYSTEM FILES (Windows only)
- Automated System Recovery (ASR) (Windows Server 2003 with no VSS client license or with VSS disabled, and XP Professional only)
- SHAREPOINT (Windows 2000 only)

For Windows Server 2003 with VSS licensed and enabled, and for Windows Vista or Windows Server 2008, these predefined save sets are available:

- All
- VSS SYSTEM BOOT
- VSS SYSTEM FILESET
- VSS SYSTEM SERVICES
- VSS USER DATA
- VSS OTHER
- VSS ASR DISK (Windows Server 2003 only)

Note: To properly protect Sun StorageTek EBS 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.

Note: Windows SYSTEM save sets do not support deduplication.

For information about:

- SYSTEM, VSS SYSTEM, SHAREPOINT, and ASR save sets, see Appendix A, "SYSTEM and VSS SYSTEM Save Sets."
- Sun StorageTek EBS support for ASR, see Chapter 23, "Support for Microsoft Automated System Recovery."

The All save set

The All save set is the default save set used when a client is created:

- For UNIX operating systems, the All save set backs up all locally mounted file systems.
- For Windows operating systems, the All save set backs up all locally mounted file systems, plus operating system files that determine the client system's state.

Table 10 provides a list of the save sets included in the All save set.

Table 10 Components in the All save set(1 of 2)

Operating system	Files
UNIX	Backs up all file systems listed in these locations: Solaris: /etc/vfstab HP-UX, HP Tru64, Linux, and SGI: /etc/fstab AIX: /etc/filesystems **Note:** If a save set's name includes a symbolic link, recovery by save set ID is not supported.
Windows XP Professional, Windows Server 2003 with no VSS client license or with VSS disabled	SYSTEM STATE SYSTEM DB SYSTEM FILES SHAREPOINT ASR All local, physical drives

Table 10 Components in the All save set(2 of 2)

Operating system	Files
Windows Server 2003 with VSS licensed and enabled, and for Windows Vista or Windows Server 2008	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
Windows 2000	SYSTEM STATE SYSTEM DB SYSTEM FILES SHAREPOINT All local, physical drives
Mac OS X	All local and mounted volumes

Using the save set All to back up particular file systems

When there are multiple file systems locally mounted on the Sun StorageTek EBS host, you can use a keyword to back up only one type of locally mounted file system.

To back up only those locally mounted file systems of a particular type, type **All-***filesystem* in the client's Save Set attribute. For example, to back up all locally-mounted ZFS filesystems on a Solaris host, type the following in the client's Save Set attribute:

```
All-zfs
```

Using the save set All to back up Solaris zones

One or more Sun StorageTek EBS Client resources can be set up for each global and nonglobal zone on a physical Solaris host. You can specify the All-local keyword in the Client resource's Save Set attribute to back up Solaris global and nonglobal virtual zones as follows:

- If the All-local keyword is specified for a global zone client, all nonglobal zones on the physical host will be backed up.
- If the All-local keyword is defined for a nonglobal zone client, only the nonglobal zone client on the physical host will be backed up.

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** command. "Performing a manual backup from the command prompt" on page 91 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.

Note: You can also start a scheduled backup manually. "Starting a group immediately" on page 549 provides information about starting a scheduled backup group manually.

Performing a manual backup on Windows

Note: If performing a Sun StorageTek EBS User backup on a Sun StorageTek EBS server, see "Excluding file type devices from a manual backup on Windows" on page 89.

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

   Note: 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 90 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 Sun StorageTek EBS log file.

   Note: Certain kinds of corrupt files or errors on computer disk volumes are not detected. Sun StorageTek EBS 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 Sun StorageTek EBS User backup on a Sun StorageTek EBS 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 Sun StorageTek EBS User backups, create a local directive on the Sun StorageTek EBS 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 228 provides more information about the *networkr.cfg* file. Chapter 11, "Device Operations" 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 and VSS SYSTEM Save Sets" provides more information.

Note: 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 Sun StorageTek EBS 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 91 and Figure 8 on page 91.

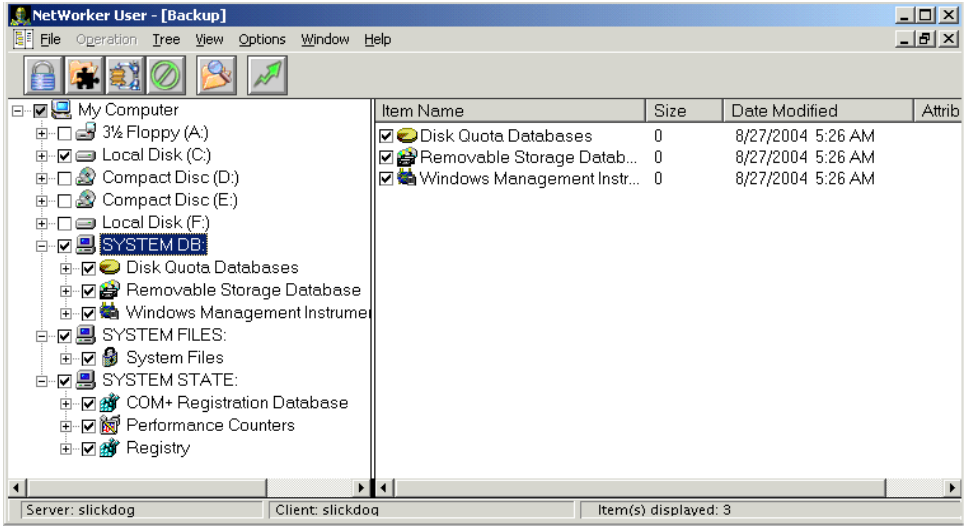


Figure 7 **SYSTEM save sets in Sun StorageTek EBS User backup window**

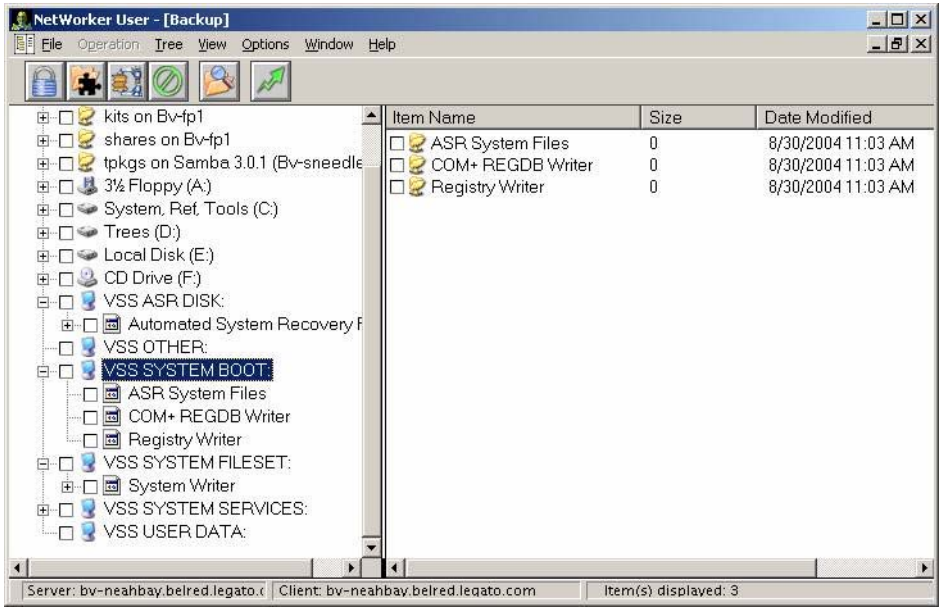


Figure 8 **VSS SYSTEM save sets in Sun StorageTek EBS User backup window**

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 Sun StorageTek EBS 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.

Note: 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 Sun StorageTek EBS server defined in the /nsr/res/servers file that comes first in alphabetical order.

Note: You can also manually back up the bootstrap and indexes for a group by using the **savegrp** command 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 no VSS licence or VSS disabled

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

- Windows Server 2003 with no VSS license or VSS disabled
- Windows XP Professional or Windows 2000

To back up all components of the Windows system state:

```
save [-s Sun StorageTek EBS_server_name] "SYSTEM
STATE:"
save [-s Sun StorageTek EBS_server_name] "SYSTEM
FILES:"
```

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

```
save [-s Sun StorageTek EBS_server_name] "SYSTEM DB:"
save [-s Sun StorageTek EBS_server_name] "SHAREPOINT:"
save [-s Sun StorageTek EBS_server_name] "ASR:"
```

With VSS enabled

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

To back up all components of the Windows system state:

```
save [-s Sun StorageTek EBS_server_name] "VSS SYSTEM
BOOT:"
save [-s Sun StorageTek EBS_server_name] "VSS SYSTEM
FILESET:"
```

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

```
save [-s Sun StorageTek EBS_server_name] "VSS SYSTEM
SERVICES:"
save [-s Sun StorageTek EBS_server_name] "VSS USER
DATA:"
save [-s Sun StorageTek EBS_server_name] "VSS OTHER:"
save [-s Sun StorageTek EBS_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.

Note: 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:"
```

Note: 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 Sun StorageTek EBS 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 Sun StorageTek EBS 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

Note: This feature is not available on UNIX clients.

Use the Sun StorageTek EBS Verify feature to ensure that backup data on the Sun StorageTek EBS 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 Sun StorageTek EBS server alerts you to any changes that have occurred to your data since the backup. Verification also determines whether a hardware failure kept the Sun StorageTek EBS server from completing a successful backup. The Verify feature provides a way to test the ability to recover data.

Note: 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.

Probe-based backups

The NetWorker server schedules probe-based backups that are based on user-defined events for clients and NetWorker modules, not just on time.

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

1. Create a user defined client probe, as described in "Creating a client probe" on page 655.
2. Associate a probe with a Client resource, as described in "Associating a probe with a Client resource" on page 656.
3. Create a probe group, as described in "Creating and scheduling a probe group" on page 169.
4. Run the probe group the same as you would a save group, as described in "Scheduled backups" on page 74, or "Manual backups" on page 88.

Note: Running **savegrp -g** bypasses probing when running a probe group backup from the command line.

Deduplication backups

Sun StorageTek EBS clients can be configured to take advantage of EMC Avamar deduplication technology to dramatically decrease the amount of time, network bandwidth, and disk capacity required to back up client data.

All typical client operations, such as adding, deleting, and drag-and-drop, work the same for deduplication clients as they do for other types of clients.

A deduplication save set is treated as a regular NetWorker save set, although it actually has two parts: metadata (hash info) and the

backed-up client data. Only the metadata is stored on a NetWorker storage node. The backed-up client data is stored on a deduplication node (Avamar server). The metadata can be cloned in the usual manner, while a copy of the backed-up data can be replicated on another deduplication node.

Consider the following points:

- Bootstrap and index data is not deduplicated.
- Windows SYSTEM save sets do not support deduplication.
- 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.

For deduplication backup information, see:

- "Creating a deduplication node and a replication node" on page 860
- "Deduplication backups" on page 856

BMR backups

Sun StorageTek EBS clients can be configured to take advantage of EMC HomeBase bare metal recovery (BMR) technology to ensure that a client can be restored or migrated from one hardware type to another. When a backup is run on a client, a HomeBase profile is automatically generated and saved to a specified HomeBase Server. This process is enabled with a HomeBase Agent license.

Chapter 27, "Support for HomeBase" provides detailed information about configuring and enabling a Sun StorageTek EBS client for BMR support.

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.



IMPORTANT

Do not use AES 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.

Note: AES data encryption is supported in Sun StorageTek EBS release 7.3 and later. For compatibility with earlier NetWorker releases, PW1 password protection and PW2 encryption are supported.

Set the Datazone pass phrase for a Sun StorageTek EBS 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**.

Note: 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 encryption. apply the Encyrption global directive to individual clients by using the Directive attribute of the Client resource. For information about editing a Client resource, see "Editing a client" on page 650.

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.

Note: 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. Chapter 6, "Directives" provides information about creating and applying directives.

Applying compression to a manual backup

To compress data for a manual backup on Windows:

1. Mark the data to be compressed.
2. Select the **Special Handling** option from the **File** menu in the NetWorker **User** program. "Special data handling for Sun StorageTek EBS clients on Windows" on page 98 provides more information.

To compress data for a manual backup on UNIX, you must use the compressasm in a local directive file. Chapter 6, "Directives" provides more information.

Special data handling for Sun StorageTek EBS 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 handing 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 option.

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

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

Backing up Sun StorageTek EBS Console Management data

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

Note: The **savepsm** backup command backs up the Console 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 Sun StorageTek EBS Console database

If a Sun StorageTek EBS server was specified during the setup of the Console server in the Console Configuration Wizard, a Client resource was created to back up the Sun StorageTek EBS Console database on a scheduled basis. If a Client resource was created, tailor 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 Sun StorageTek EBS Console database, create a Client resource as follows:

1. From the **Console** window, click **Setup**.
2. From the **Setup** menu, select **Set Database Backup Server**.

3. In the **Sun StorageTek EBS server** attribute, type the name of the Sun StorageTek EBS server that will back up the Console server database.
4. Select the **Create client** resource on this server attribute.
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 Sun StorageTek EBS server to back up the Console server database through the Console Configuration Wizard. "Accessing the Console Configuration Wizard" on page 582 provides more information.

"Setting up a scheduled backup" on page 75 provides information about creating and tailoring a Client resource.

Note: Only full, incremental, and skip backup levels are supported. All other backup levels (1 through 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> (Linux)

The environment variable to set varies by platform:

- Solaris/Linux: LD_LIBRARY_PATH
- AIX: LIBPATH
- HPUX: SHLIB_PATH

To perform a manual backup of the Sun StorageTek EBS Console database, type the following:

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

where <*Console_install_dir*> is the installation directory for the Console server. The default installation directory is /opt/SUNWebsg.

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 database is performed.

To truncate the transaction log file without performing a backup, type the **savepsm** command as follows:

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

Consideration for managing the Console database log file

If the transaction log file is truncated manually, the next backup of the Console database 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 database is a full backup:

1. Open the backup schedule for the Client resource that is associated with the Console 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 187 provides formation about editing 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 100 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 Server 2003, Windows XP Professional, and Windows 2000.

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.

Note: Sun StorageTek EBS 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 74 provides for more information about setting up a scheduled backup. "Directory specifications" on page 233 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 88 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*.

   Note: 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.

Note: 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 2000 or later

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.

Note: 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 Sun StorageTek EBS 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 Sun StorageTek EBS software to skip CIS catalog backups

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

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

Chapter 6, "Directives" 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 Sun StorageTek EBS SYSTEM DB save set. However, you can use the procedures in this section to configure Sun StorageTek EBS software to protect these databases.

Note: If VSS is licensed and 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 74 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 Sun StorageTek EBS 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 Sun StorageTek EBS client that is the WINS server.

Note: 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 106
- "Enabling hard link support" on page 107
- "Recovery and case-sensitivity" on page 107
- "Failed backup and recovery attempts" on page 108
- "Granting full permissions for backup of Disk Quota database" on page 108
- "Security settings for logging operations performed by backup operator" on page 109

Enabling short filename support

On Windows Server 2003 and Windows XP Professional, Sun StorageTek EBS 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 Sun StorageTek EBS 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**.

Note: Due to operating system limitations, Sun StorageTek EBS clients cannot save or restore these system-assigned short filenames when running on Windows 2000.

Enabling hard link support

The Sun StorageTek EBS 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 Sun StorageTek EBS 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**.

Recovery and case-sensitivity

The Sun StorageTek EBS 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.

Failed backup and recovery attempts

The Sun StorageTek EBS 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 904 provides information about viewing log files with the **nsr_render_log** program.

Granting full permissions for backup of Disk Quota database

Sun StorageTek EBS 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 Sun StorageTek EBS client that uses the Windows Disk Quota feature, during SYSTEM DB or VSS SYSTEM SERVICES backup, Sun StorageTek EBS 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 Sun StorageTek EBS 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.

Note: 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.

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 <Sun StorageTek EBS_install_path>\logs directory.

To enable Sun StorageTek EBS logging for Backup Operators, modify the security settings on the <Sun StorageTek EBS_install_path>\logs directory. For example, on a Sun StorageTek EBS client that is running Windows 2000, update the security settings as follows:

1. In **Windows Explorer**, navigate to the <Sun StorageTek EBS_install_path>\logs directory.
2. Right-click the <Sun StorageTek EBS_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**.

Sun StorageTek EBS 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 Sun StorageTek EBS 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 Sun StorageTek EBS 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. Executes 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 110 provides more information.
   - By typing **savepnpc** in the Backup Command attribute of the Client resource. This way, the client backup invokes the **savepnpc** program instead of the **save** program. The first time the client is backed up, **savepnpc** creates a default backup program file, which you can then customize for future backups of the client. "Using the savepnpc command with a customized backup program" on page 116 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 will be executed 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 **savepnpc**). If you specify a save set value of All, the program is executed 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 will be executed for each save set listed.

Note: 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 Sun StorageTek EBS **save** command.
- Always specify the full path of the **save** command in the script.
- The Sun StorageTek EBS **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 be successfully executed. Otherwise, the Sun StorageTek EBS server cannot complete the remaining instructions.
- On UNIX and Linux, when you invoke the Sun StorageTek EBS **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 Sun StorageTek EBS **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 Sun StorageTek EBS **save** command resides.

   Note: 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 Sun StorageTek EBS **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 executes pre-backup commands, runs the Sun StorageTek EBS **save** command, and then executes 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: Executes Sun StorageTek EBS commands to run the required 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 <Sun StorageTek EBS_install_path>\logs directory.

"Monitoring window" on page 545 provides information about viewing execution file details from the **Monitor Groups** tab.

"Viewing log files" on page 904 provides information about viewing log files.

This is an example backup script:

```
@ECHO OFF
SETLOCAL
ECHO =======START BATCH FILE================

ECHO =====Sun StorageTek EBS 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 =====Sun StorageTek EBS 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 Sun StorageTek EBS backup commands.

:save

REM Note: Enter correct path to your Sun StorageTek EBS
bin
REM directory (line below is default path)
C:\PROGRA~1\nsr\bin\save.exe %arg%

ECHO =====Sun StorageTek EBS 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 sevegrp.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 ===Sun StorageTek EBS 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 ===Sun StorageTek EBS 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 =====Sun StorageTek EBS 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 nsrlpr: 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 Sun StorageTek EBS **save** command (for example, on a Solaris system, the **save** program is installed in the /usr/sbin/nsr 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.

Note: 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
```

Using the savepnpc command with a customized backup program

As an alternative to using the **save** program with a custom script, use the **savepnpc** program. The **savepnpc** program differs from using a custom script with the **save** program in that preprocessing and postprocessing commands execute 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 **savepnpc** 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 **savepnpc** command.

To execute the **savepnpc** 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:

   savepnpc
6. Back up the client.

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

   - The /nsr/res (UNIX)
   - The *Sun StorageTek EBS_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: savepnpc;
   ```

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

Note: 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**.

Once 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 **savepnpc** 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 **savepnpc** 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 **savepnpc** 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 **savepnpc** 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 **pstcmd** 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: savepnpc;
precmd:
"V:\\usr\\sap\\PDB\\SYS\\exe\\run\\PDB-stop.cmd >
C:\\WINNT\\system32\\PDBStop.log 2>&1";
pstcmd:
"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 Sun StorageTek EBS 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 executed and no error message is 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:

     Sun StorageTek EBS_install_path\tmp*group-name*.res to *Sun StorageTek EBS_install_path*\tmp\ *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
   - *Sun StorageTek EBS_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 Sun StorageTek EBS client:

- UNIX: /nsr/logs
- Microsoft Windows: *Sun StorageTek EBS_install_path*\logs

The format of the **savepnpc** 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 904 provides information about viewing log files.

Considerations for backing up raw partitions

The Sun StorageTek EBS 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 238 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 Sun StorageTek EBS_server_name \\.\PhysicalDrive0
save -s Sun StorageTek EBS_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.

Support for backing up renamed directories

When a directory is renamed, the default behavior during a non-full level backup is to skip the unchanged files and folders under the renamed directory. 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.

In NetWorker 7.5 or greater, you can back up all of the files and folders under a renamed directory regardless of whether or not they have changed. The advantage of this feature is that data under the renamed directory does not appear to be lost when performing a recovery.

Be aware of the following when deciding whether or not to enable this feature:

- The backup renamed directory feature is enabled on a client basis. Consider enabling this feature only on clients that are prone to issues with recovering renamed directories.
- Even if no renamed directories are encountered, enabling this feature will increase the backup time for a client.
- Enabling this feature can further increase the backup time and result in increased media usage, depending on the level at which a renamed directory is encountered. A full backup will be performed on all sub-directories and files under a renamed directory.
- The enterprise environment and the NetWorker server and client operating system platforms have a large impact on the overhead associated with enabling this feature.

To enable the backup renamed directory feature:

- Select the Backup renamed directories attribute in the Client resource. "Editing a client" on page 650 provides information about editing a Client resource.

Note: If this option is selected and 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.

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 78
2. In the backup group properties Advanced tab, select **Index only** from the **Options** attribute.
3. Click **OK**.

3 Configuring Libraries and Devices

This chapter covers these topics:

- Devices and libraries 124
- Storage nodes 126
- Storage node configuration 127
- Deduplication nodes and replication nodes 132
- Dynamic drive sharing 132
- Autodetection of libraries and devices 141
- Add (configure) libraries 145
- Library reconfiguration 150
- Add (configure) stand-alone devices 152
- Support for cloud backup devices 155
- Deleting libraries and devices 162
- Library slots 163
- Troubleshooting autoconfiguration failure 164

Devices and libraries

Sun StorageTek EBS software automatically recognizes devices that are being used for backups and recoveries. The maximum number of configured devices for any Sun StorageTek EBS server and storage node combination is 512. The maximum number varies depending on the particular server that is being administered. The maximum number also does not include read-only devices created for advanced file type devices. "File type and advanced file type devices" on page 322 provides information about file type devices, advanced file type devices, and the Sun StorageTek EBS DiskBackup option.

The Sun StorageTek EBS Console permits extensive handling of devices, including the ability to autoconfigure:

- SCSI devices
- NDMP devices
- Silos

The details that describe a particular device can be viewed, and often changed, in the device's Properties.

To view details:

1. Select **View** > **Diagnostic Mode** to see the full range of settings and options available in the **Properties** window.
2. In the **Properties** window, click **Field Help** for a description of the various fields and attributes.

These options are available from many of the menus throughout the **Devices** task:

- Configure all Libraries
- Scan for Devices

If these options are started from the server folder instead of from the Storage Node folder, then all storage nodes on the Sun StorageTek EBS server are automatically selected for configuration in the wizard, or for scanning, respectively.

As with other Console functions, users can view and work with only those Sun StorageTek EBS servers for which they have access permission.

Support for LTO-4 hardware-based encryption

The use of LTO-4 hardware-based encryption is supported by Sun StorageTek EBS 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 Sun StorageTek EBS application can read from and write to LTO-4 devices that use hardware-based encryption. The use of this encryption is transparent to Sun StorageTek EBS. Neither the encryption nor the key management process is managed by the Sun StorageTek EBS application. This includes the ability to turn encryption on or off within Sun StorageTek EBS, and the management of encryption keys.

Storage nodes

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

From the Sun StorageTek EBS server, typical storage tasks can be performed, such as:

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

Only users who have the Configure Sun StorageTek EBS privilege can add to or change the configuration of the Sun StorageTek EBS server, media devices, and libraries. "Sun StorageTek EBS User Groups" on page 613 provides more information.

Prerequisites

To operate the Sun StorageTek EBS 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 order listed:
  1. Sun StorageTek EBS client software
  2. Sun StorageTek EBS storage node software
  3. (Optional) EMC AlphaStor software.

     This enables multiple Sun StorageTek EBS 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 Sun StorageTek EBS client and storage node software.

Licensing

Chapter 28, "Licensing" provides information on Sun StorageTek EBS licensing support for storage nodes.

Storage node configuration

After the required software and enabler codes are installed, each storage node device must be configured (by adding them to the Sun StorageTek EBS server). This procedure is similar to the tasks required for other devices. "Configuring a storage node" on page 127 provides more information.

Configuring a storage node

To configure a Sun StorageTek EBS storage node:

1. In the server's **Administration interface**, click the **Devices** button.
2. Right-click **Storage Nodes** in the navigation tree, and select **New**. The **Create NSR Storage Node** window appears, with the **General** tab shown.
3. In the **Identity** area:
   a. Type the fully-qualified domain name or short name of the storage node in the **Name** field.
   b. Check a box (NDMP, SCSI, or silo) in the **Type of Storage Node** field. The default is SCSI.

   Note: If NDMP is selected, the remote username and password must be entered. If configuring a silo for use with NDMP, select silo rather than NDMP. Also verify that the silo robot arm is undetected by the NDMP storage nodes.

4. In the **Device Management** area, configure:
   a. Select a **Device Sharing Mode**. The default is **Server Default**.

b. Select whether to search all LUNs (logical unit numbers). The default is **No**.

c. Name SCSI targets to skip, if any, and if the storage node type is set to SCSI.

d. Enter only one SCSI target per line.

5. In the **Remote Host** area, type the name of any remote user who is to be allowed to access the remote host from the storage node, along with the password for this remote user.

Note: Only one remote user with password is allowed per storage node. This applies to NDMP and advanced file type devices (AFTDs). "File type and advanced file type devices" on page 322 provides more information about AFTDs.

---

6. Click **OK** when done. The new storage node appears in the navigation tree.

Modifying timeout attribute for storage node operations

An attribute named **Nsrmmd Control Timeout**, which is set during Sun StorageTek EBS server configuration, configures the amount of time a Sun StorageTek EBS 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 **Nsrmmd Control Timeout** is five minutes.

Other attributes involved in storage node timeouts include:

- Nsrmmd polling interval, which determines the number of minutes between storage node polls.
- Nsrmmd restart interval, which determines the number of the minutes the Sun StorageTek EBS software waits before restarting the nsrmmd process. A value of zero for the **Nsrmmd 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 Sun StorageTek EBS server and select **Properties**.

4. Select the **Media** tab.
5. Modify the attributes as appropriate and click **OK**.

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**

Set the **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**.

Storage node affinity

The choice of which Sun StorageTek EBS 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 Sun StorageTek EBS 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 Sun StorageTek EBS server.

If a Client resource is created after a storage node is created, and the client is to back up to that storage node, the name of the storage node must be entered in the client's **Storage Nodes** attribute *above* the default **nsrserverhost**. Storage node names can be added to this list at any time. The client directs its data to the first storage node in the list that has an enabled device and is capable of receiving the data.

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**.

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 Sun StorageTek EBS server. A bootstrap cannot be backed up to a remote device, but a bootstrap clone can be written 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.

This bootstrap information also applies to NDMP devices.

Dedicated storage nodes

In Sun StorageTek EBS release 7.0 and later, 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.

Note: 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 Sun StorageTek EBS release 7.6, Sun StorageTek EBS 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. Sun StorageTek EBS 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 Sun StorageTek EBS data zone.

Tips for 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.

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 Sun StorageTek EBS side.

"Creating a deduplication node and a replication node" on page 860 provides information on how to create a Sun StorageTek EBS deduplication node.

Dynamic drive sharing

Dynamic Drive Sharing (DDS) is a feature that provides Sun StorageTek EBS software with the ability to recognize shared drives. DDS enables Sun StorageTek EBS software to do the following:

- Skip the shared drives that are in use.
- Route the backups or recoveries to other available shared drives.

Introduction to DDS

DDS controls application requests for media and allows the Sun StorageTek EBS 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.

Note: 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 Sun StorageTek EBS system provides these benefits:

- Reduces storage costs — A single tape drive can be shared among several storage nodes. In fact, since Sun StorageTek EBS 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 9 on page 134 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.

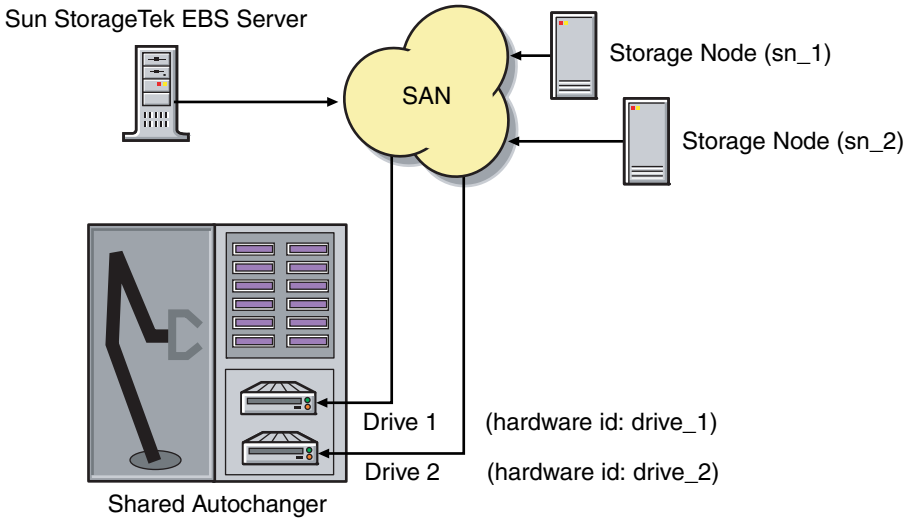


Figure 9 **Dynamic Drive Sharing**

Note: Ensure that all applicable devices can be seen from each storage node by running the **inquire -l** command locally on each storage node.

Block-size compatibility

With DDS enabled, drives can be shared between storage nodes on different platforms, such as UNIX and Microsoft Windows. For Sun StorageTek EBS 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.

"Compatible block size for UNIX and Microsoft Windows" on page 349 provides information about how to set block sizes for individual drives or tapes on different platforms.

Block-size incompatibility between UNIX and Microsoft 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 Sun StorageTek EBS 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.

SCSI Reserve/Release

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 Sun StorageTek EBS 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. Sun StorageTek EBS 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 359 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. It is also not supported on IRIX.
- It is supported only on a specific platform set and configuration. For details, refer to the *Sun StorageTek EBS Release Notes*.
- 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 Sun StorageTek EBS Reserve/Release feature to work.
- Even if all of the enterprise's Sun StorageTek EBS 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. Ensure, however, that the NDMP nodes are homogeneous.

For example, DDS can be enabled in these configurations:

- EMC Celerra® to EMC Celerra
- EMC IP4700 to EMC IP4700
- 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.

Figure 10 on page 138 illustrates a basic DDS configuration with NDMP.

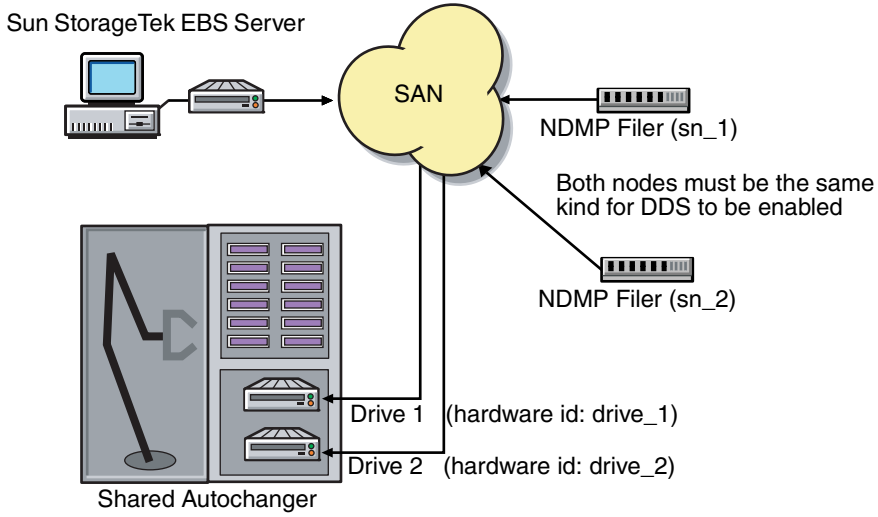


Figure 10 DDS with NDMP

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 11 on page 140 lists this attribute's allowed values and their descriptions.

Table 11 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 Advanced tab of a device's Properties when in diagnostic mode. The default value is 0 (zero) minutes. Select a time that is appropriate for the system configuration.

Note: Selection of zero means that the device never times out and the tape must be ejected manually.

High availability and DDS

The Sun StorageTek EBS 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 Sun StorageTek EBS 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 5 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 Sun StorageTek EBS 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 Sun StorageTek EBS software be able to access the same path through which the volume was initially loaded.

Licensing

Chapter 28, "Licensing" provides information about Sun StorageTek EBS licensing support for DDTS.

Autodetection of libraries and devices

This section provides information on scanning for libraries and devices.

Scanning for libraries and devices

Devices already known to the Sun StorageTek EBS server can be seen in the enterprise hierarchy in the navigation tree. Use the **Scan for Devices** option described here to find devices that were not yet known to the Sun StorageTek EBS server. It should also be noted 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.

To scan for available devices:

1. In the **Console** window, click **Enterprise**.
2. In the navigation tree, select a Sun StorageTek EBS server.

3. In the **Name** column of the **Host detail** table, double-click **Sun StorageTek EBS**. The **Sun StorageTek EBS Administration** window for the selected server opens. Note that while multiple **Sun StorageTek EBS 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 Sun StorageTek EBS 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 Sun StorageTek EBS 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.

Note: 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 Sun StorageTek EBS datazones.

Configuring the Sun StorEdge Network Foundation HBA/Driver

This section applies only to storage nodes that run Solaris 9 or earlier. The StorEdge Network Foundation host bus adapter (HBA) requires a special configuration to bind the device World Wide Port Names that are used by the EMC logical unit specification (lus) driver.

The **/usr/sbin/nsr/lus_add_fp_devs** script adds the proper entries to the lus.conf files. Before the script is run, however, ensure that all devices are powered on and connected to the HBAs.

Note: The **/usr/sbin/nsr/lus_add_fp_devs** script should be used only to configure libraries over Fibre Channel.

These are examples of additional entries for the /usr/kernel/drv/lus.conf script:

```
name="lus" parent="fp" target=0 lun=0
fc-port-wwn="22000090a50001c8";
name="lus" parent="fp" target=2 lun=0
fc-port-wwn="22000090a50001c8";
```

Note: Each time a new device is added, or an old device is removed, the new lus configuration must be re-created and reinstalled. "How to add a library device" on page 146 provides information about how to add or delete a device. When the **/usr/sbin/nsr/lus_add_fp_devs** script runs, check whether the World Wide Port Names appear in the script output. If none appears, the Fibre Channel might be configured incorrectly. Check the configuration by using the Sun **luxadm** program.

Sample /usr/sbin/nsr/lus_add_fp_devs script output

This is an example of **/usr/sbin/nsr/lus_add_fp_devs** script output:

```
# /usr/sbin/nsr/lus_add_fp_devs
   Updating /usr/kernel/drv/lus.conf
Found StorEdge HBA device:

/devices/pci@1f,0/pci@1/pci@1/pci@f/SUNW,qlc@4/fp@0,0:de
vctl
```

```
Found StorEdge HBA device:

/devices/pci@1f,0/pci@1/pci@1/pci@f/SUNW,qlc@5/fp@0,0:de
vctl
Mapping tape devices starting at target 0
   Adding:
   port WWN: 500104f000428e48 as SCSI target 0
   port WWN: 500104f000428e49 as SCSI target 1
   port WWN: 500104f000428f44 as SCSI target 2
   port WWN: 500104f000428f45 as SCSI target 3
Mapping tape devices starting at target 4
   Adding:
      port WWN: 500104f000428e48 as SCSI target 4
      port WWN: 500104f000428e49 as SCSI target 5
      port WWN: 500104f000428f44 as SCSI target 6
      port WWN: 500104f000428f45 as SCSI target 7
Adding medium changer devices starting at target 100
   Adding:
      port WWN: 100000e002000000 as SCSI target 100
      port WWN: 100000e00201105a as SCSI target 101
Adding medium changer devices starting at target 102
   Adding:
      port WWN: 100000e002000000 as SCSI target 102
      port WWN: 100000e00201105a as SCSI target 103
The changes made to lus.conf will not take effect until
lus is reloaded.
If you are sure that no process is currently using lus to
control a tape
   library, then it may be safely unloaded.
```

Note: If lus is not currently loaded, proceed to the section in the script stating that lus is not currently loaded.

```
Would you like to unload lus? (y/n)  y
```

Note: If no is selected, this output appears and the program exits:

```
lus is currently loaded! It must be unloaded and reloaded
for any changes to lus.conf to take effect
```

Note: If lus is not currently loaded, respond to this prompt:

```
lus is not currently loaded
Would you like to run inquire? (this will cause lus to be
loaded using
the newly modified lus.conf) (y/n) y
```

Note: If no is selected at the previous prompt, the program exits.

```
scsidev@0.0.0:FUJITSU MAJ3182M SUN18G 0804|Disk,
/dev/rdsk/c0t0d0s2
   scsidev@2.0.0:STK 9840 1.28|Tape, /dev/rmt/0cbn
   scsidev@2.1.0:STK 9840 1.28|Tape, /dev/rmt/2cbn
   scsidev@2.100.0:ADIC Scalar DLT 448  0114|Autochanger
(Jukebox)
   scsidev@2.101.0:ADIC Scalar DLT 448  0114|Autochanger
(Jukebox)
   scsidev@3.0.0:STK 9840 1.28|Tape/dev/rmt/6cbn
   scsidev@3.1.0:STK 9840 1.28|Tape, /dev/rmt/8cbn
```

Add (configure) libraries

A device resource must be created for each library device, including silos, on a storage node. Because a Sun StorageTek EBS server is also a storage node, this procedure applies to a server and all its storage nodes. Users can choose between configuring libraries automatically through use of a configuration wizard, or configuring them manually and individually in the user interface.

Storage nodes must be created before devices can be configured to be used by them. "Storage nodes" on page 126 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 (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.

These device types can be configured automatically:

- NDMP
- SCSI
- Silo (except DAS silo)

Note: On Microsoft Windows systems, Removable Storage Manager (RSM) should be disabled prior to configuring a SCSI-attached library.

These device types must be configured by using the **jbconfig** command:

- AlphaStor devices.

- DAS silo
- Devices managed by Microsoft Removable Storage.
- 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.

How to add a library device

To configure a new library 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 Sun StorageTek EBS 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. 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.

How to add a silo device

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 Sun StorageTek EBS 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 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. If a DAS silo controller was chosen, enter the DAS client name (which should match the storage node name).

f. (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 Sun StorageTek EBS software with ACSLS silos

Note: 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: ACSLS on a Solaris or AIX host, Library Station on an MVS host, or Horizon Library Manager on a system running Windows NT or Windows 2000.

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 Sun StorageTek EBS 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 Sun StorageTek EBS 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.

Queuing device resources for AlphaStor

Because the Sun StorageTek EBS 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*.

Virtual Tape Library (VTL) configuration

During library configuration, the Sun StorageTek EBS software automatically detects that a library is a VTL, and updates the read-only Virtual Jukebox attribute to Yes. If the NetWorker software determines that the library being configured is not a VTL, it sets the Virtual Jukebox attribute to No.

Licensing

Chapter 28, "Licensing" provides information about Sun StorageTek EBS licensing support for a Virtual Tape Library.

Library reconfiguration

Users with Configure Sun StorageTek EBS privileges can reconfigure a library in order to add or remove access paths to the drives in a library. This includes access paths that allow libraries to be shared.

Be aware of the following considerations when reconfiguring a library:

- The reconfiguring 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 reconfigure library feature described in "Reconfiguring a library" does not support adding NDMP drives to a non-NDMP library if both the NDMP server and the Sun StorageTek EBS storage node are on the same host. Instead, use the **jbedit** command, which is described in "Using the jbedit command" on page 151.

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

The **jbedit** program can be used as a fallback means of editing library configurations if, for some reason, the autoconfiguration program is not working. The command can be run on a Sun StorageTek EBS server, storage node, or client (if the client is a storage node). It works without disrupting any backup or recovery operations on the library.

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 150. The **jbedit** options provide selection lists that make it easy to find drives or devices to be added or deleted.

Note: Running the **jbedit** program requires Configure Sun StorageTek EBS privileges.

Table 12 on page 152 lists the most commonly used **jbedit** program options.

Table 12 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.

Add (configure) stand-alone devices

A Device resource must be created for each stand-alone 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 126 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 **Sun StorageTek EBS 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 **Sun StorageTek EBS 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.

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.

"File type and advanced file type devices" on page 322 provides instructions and restrictions on backing up to a file type device.

Configuring Sun StorageTek EBS DiskBackup

The Sun StorageTek EBS DiskBackup solution greatly reduces the time to save and recover data when compared to the use of tape. Sun StorageTek EBS DiskBackup enables data to be saved to a computer's local or network-attached disk, rather than by using a tape device to write the data to storage media.

Sun StorageTek EBS software offers two variations of this method:

- File Type Device (local only)
- Advanced File Type Device (can be a networked drive)

"File type and advanced file type devices" on page 322 provides information about the Sun StorageTek EBS DiskBackup solution.

Support for cloud backup devices

This section describes how to configure the Sun StorageTek EBS 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. Both ATMOS On Premise option (onsite) and ATMOS Online option (offsite) cloud configurations are supported.

Cloud backups enable you to simplify backup operations while lowering costs. Backup disks can be maintained in an on premise cloud configuration or offloaded to the ATMOS Online option. In this way, your organization can simplify and isolate, or eliminate, the burden of maintaining backup infrastructure.

Cloud backup devices compared to other device types

Sun StorageTek EBS 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 13 on page 155 lists the major similarities with other backup device types as well as the unique features of a cloud storage device.

Table 13 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 prerequisites and considerations

The following conditions must be met before you can backup to the cloud:

- The Sun StorageTek EBS Cloud Backup Option must be licensed and enabled. "Licensing Sun StorageTek EBS software components" on page 886 provides information about entering licence enablers.
- The NetWorker Cloud Backup Option is supported on Windows and Linux storage nodes only.
- An ATMOS On Premise or ATMOS Online 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 your ATMOS online 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:

- Sun StorageTek EBS Avamar deduplication storage nodes do not support cloud backups.
- For NDMP, only a Data Server Agent (DSA) is supported for cloud backups.
- Sun StorageTek EBS 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. For additional best practices information, refer to the *Best Practices for Cloud Backup with NetWorker*

Fast Start document, which is available in the NetWorker Whitepapers and Technical Notes section on EMC Powerlink.

Directing Sun StorageTek EBS 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 241 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 157
- "Label and mount the cloud storage device" on page 160

Create the cloud storage device

To create a cloud storage device:

1. In the server's **Sun StorageTek EBS 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: Unlike other devices, 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 Sun StorageTek EBS 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 Sun StorageTek EBS 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.

    Note: If the Sun StorageTek EBS 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 Sun StorageTek EBS AES 256 bit encryption and is selected by default. If desired a NetWorker datazone passphrase 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 passphrase, refer to "Encrypting backup data" on page 96.

Note: If encryption is already enabled for the Sun StorageTek EBS 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 on 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 bandwith. 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**.

   Note: 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 526
- **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 10, "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 9, "Cloning"provides more information.

Deleting libraries and devices

There are slight differences between how to delete a library and how to delete a device of any type.

How to delete a library

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 error 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 Sun StorageTek EBS 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.

How to delete a device

To delete a stand-alone device, file type device, advanced file type device, or a device in a library (after the library's definition has been deleted) from a storage node:

1. In the server's **Administration** interface, click **Devices**.
2. Select **Devices** in the navigation tree. The Devices detail table appears.
3. Right-click the device entry in the detail table, and select **Delete**.
4. Click **Yes** to confirm the deletion.

Note: Any attempt to remove a device from a library before the library itself (the library's definition) has been removed will result in an error message.

Library slots

The available slots feature controls which volumes the Sun StorageTek EBS 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.

Note: Ensure that volumes have been placed in all the available slots of the library so that the Sun StorageTek EBS 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.

How to configure library slots

To configure which slots are available in a library:

1. In the server's **Sun StorageTek EBS Administration** interface, select **View>Diagnostic Mode** from the menu bar.
2. Click **Devices**.
3. Open the **Libraries** folder in the navigation tree. The **Libraries** detail table appears.
4. 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**. The **Properties** window appears.
5. Select the **Advanced** tab of the **Properties** window.
6. 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.

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.

4

Backup Groups and Schedules

This chapter covers these topics:

- Overview of Sun StorageTek EBS scheduling 166
- Backup groups 166
- Managing backup groups 176
- Backing up open files 179
- Schedules 180
- Backup levels 190
- Save set consolidation 198

Overview of Sun StorageTek EBS 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, 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.

  Note: 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 7, "Sorting Backup Data" provides more information.

The Sun StorageTek EBS 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 6 **Using groups to balance client backups**

Figure 11 on page 168 illustrates how the Sun StorageTek EBS server uses two time-based backup groups to back up multiple client save sets. In Figure 11, 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 Sun StorageTek EBS 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**
Daily Start Time: **24:00 (Midnight)**

Client: mars
Save Set: All
Schedule: Full on Monday
Incremental, all other days

Client: jupiter
Save Set: All
Schedule: Full on Tuesday
Incremental on all other days

Client: saturn
Save Set: All
Schedule: Full on Wednesday
Incremental on all other days

Group: **Accounting**
Daily Start Time: **19:00 (7:00 P.M.)**

Client: mercury
Save Set: All
Schedule: Full on 1st of month
Level 5 on 15th of month
Incremental on all other days

Client: venus
Save Set: /usr/home
Schedule: Full on 1st and 15th of month
Incremental on all other days

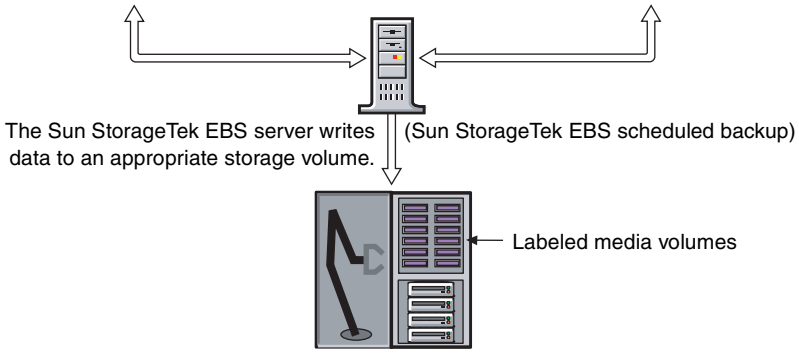


Figure 11 **How Sun StorageTek EBS groups are used to back up multiple clients**

Creating and scheduling a probe group

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.

To configure the probe-enabled group:

1. In the server **Administration interface**, click **Configuration**.
2. Right-click **Groups**, and select **New**.
3. In the **Advanced** tab of the **Create Groups** window, complete the fields in the **Probe** section as described in Table 14 on page 169.

Table 14 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	• 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.

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 655.

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.

Preconfigured groups

The Sun StorageTek EBS product ships with a single preconfigured group named Default. To ensure that all data is backed up, the Sun StorageTek EBS server automatically adds all clients to the Default group. However, you must enable the Default group for the Sun StorageTek EBS 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 15 displays some of the key attributes for the Group resource.

Table 15 Group attributes (1 of 3)

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.

Table 15 Group attributes (2 of 3)

Attribute	Description
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 174 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 Sun StorageTek EBS 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. This value cannot be greater than the Interval attribute.
Client Retries	When the Sun StorageTek EBS 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 will be terminated and the savegrp aborted. 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 invoke a retry on the save set if 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.

Table 15 Group attributes (3 of 3)

Attribute	Description
Probe interval	This attribute indicates how often probes (in minutes) should be run. Default value is 60 minutes.
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 set to 0, time since the last successful backup will not matter: **savegrp** always runs probes. If interval is specified and reached, **savegrp** will run the backup no matter what the result of the rest of the probes. The probes are run, so that the probe state data can be updated. Default is 0 days.
Time of the Last Successful Backup	Set to the time of the last successful backup by **savegrp**. 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 604 provides information about savegrp parallelism.

Aborted backup groups

If the backup of a save set fails, then the Sun StorageTek EBS 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 Sun StorageTek EBS backup hangs until it times out. The Sun StorageTek EBS server marks the save set "aborted."

How to create a group

"Task 2: Set up a group for backup clients" on page 78 provides information about creating a group.

How to edit a group

Note: You cannot change the name of an existing backup group.

To edit a group:

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

Note: 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, the new functionality is only available if a Sun StorageTek EBS group is selected in the configuration window in the Sun StorageTek EBS 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 Sun StorageTek EBS 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 172.
2. Select the **Advanced** tab.
3. In the **Interval** attribute, type a value in the hh:mm format.

   Note: 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 Sun StorageTek EBS 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 Forced 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 Sun StorageTek EBS 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 172 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 174 provides information about how to configure backup groups to occur more than once every 24 hours.

Managing backup groups

Table 16 on page 176 lists backup group management tasks and where to find more information about these tasks.

Table 16 Backup group management

Group Management Task	For more information
Start a group immediately	"Starting a group immediately" on page 549
Stop a group immediately	"Stopping a group" on page 549
Restart a group	"Restarting groups" on page 549
Preview a backup group	"Previewing a backup group" on page 176

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 Sun StorageTek EBS 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 172.
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 <Sun StorageTek EBS_install_dir>\logs directory. Information about viewing log files is available in"Viewing log files" on page 904.

Moving clients between groups

Multiple clients may be moved from one group to another by selecting the clients and dragging them to another group.

Note: 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 172.
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 *Sun StorageTek EBS_install_dir*\logs directory. "Viewing log files" on page 904 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 14, "Reporting" 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 *Sun StorageTek EBS_install_dir*\logs directory.

- "Viewing log files" on page 904 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 *EMC NetWorker Disaster Recovery Guide*.

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 904 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.

Sun StorageTek EBS 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 Sun StorageTek EBS software, from writing to or reading from the open file.

The Sun StorageTek EBS 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 Sun StorageTek EBS 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 Sun StorageTek EBS 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 88.

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 Sun StorageTek EBS releases 7.2 and later, if a VSS client license is present, 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 24, "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 Sun StorageTek EBS server what level of backup (for example, full or incremental) 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 190 provides information about backup levels.

Schedules for clients

Backups must be scheduled to avoid the node's read-only periods, when such cron jobs as checkpoint and garbage-collection are run. The Avamar server documentation provides more information.

With , backups require significantly less time, once the initial full backup has run.

"Backup levels for deduplication clients" on page 857 provides backup level and schedule information that is specific to clients.

Preconfigured Sun StorageTek EBS schedules

The Sun StorageTek EBS 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 17 on page 181 describes the preconfigured schedules.

Table 17 Preconfigured Sun StorageTek EBS schedules(1 of 2)

Schedule name	Sun StorageTek EBS 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.

Table 17 Preconfigured Sun StorageTek EBS schedules(2 of 2)

Schedule name	Sun StorageTek EBS backup operation
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.
Consolidate Every Friday	Completes a consolidated backup every Friday. Completes incremental backups on all other days.
Consolidate on First Friday of Month	Completes a consolidated backup on the first Friday of the month. Completes incremental backups on all other days. This schedule cannot be modified.
Consolidate on First of Month	Completes a consolidated backup on the first calendar day of the month. Completes incrementals on all other days of the month.
Consolidate Quarterly	Completes a consolidated backup on the first day of each quarter. Completes incremental backups on all other days of the quarter. 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 7 Weekly backup cycle

Figure 12 on page 183 illustrates a weekly backup cycle. In this example, a full backup is performed on a client each Sunday, and incremental backups are performed on the other days of the week.

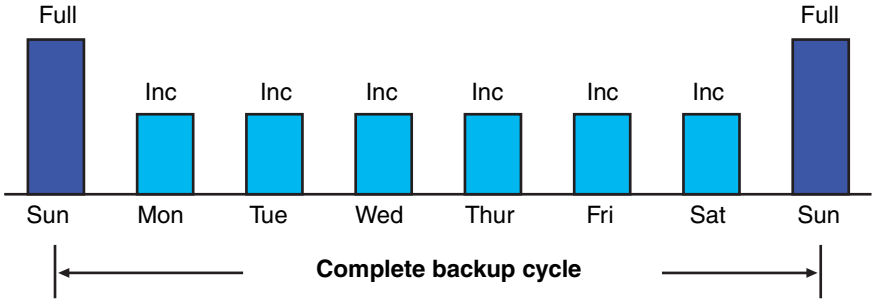


Figure 12 Weekly backup cycle

Use backup schedules to balance and stagger the load on a Sun StorageTek EBS 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 Sun StorageTek EBS server to perform full backups on Sunday, and incremental backups the rest of the week.

Example 8 Default schedule for multiple clients

Figure 13 on page 183 illustrates how the Default schedule works for three clients.

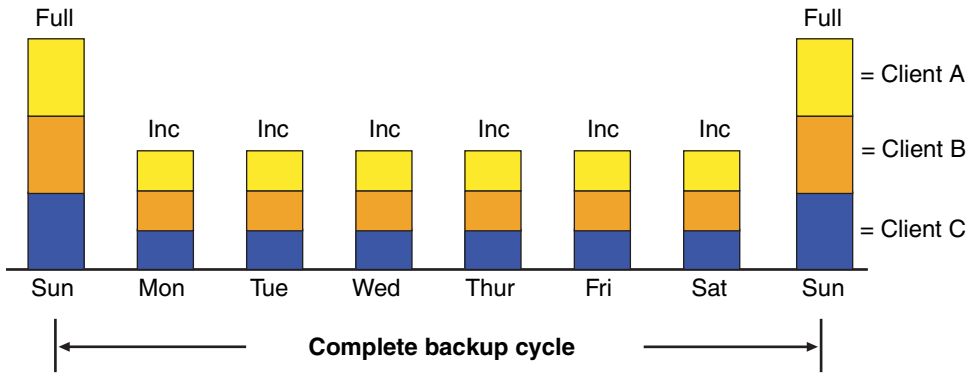


Figure 13 Default schedule for multiple clients

Note: If you have a short backup window period and need to create a full backup, you can use the consolidation backup. A consolidation backup uses the most volume space because it holds both a level 1 and a consolidated full backup. However, data recovery is faster because of the full backup. "Save set consolidation" on page 198 provides 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 9 **Staggered weekly schedules for multiple clients**

Figure 14 on page 184 illustrates how to use a staggered backup schedule for multiple clients.

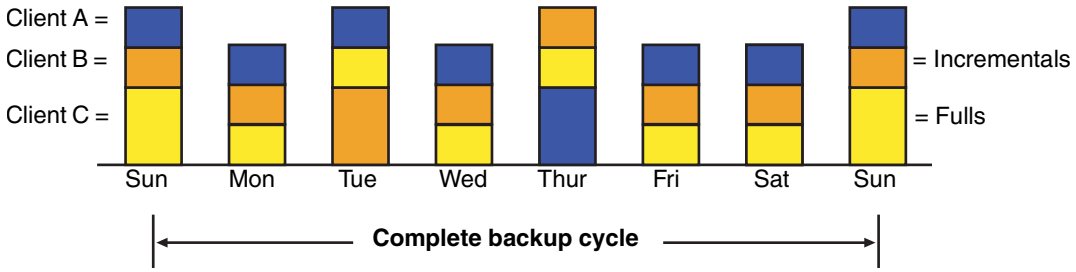


Figure 14 **Staggered weekly schedules for multiple clients**

By balancing and staggering the load, and using different start times for different groups of clients, you can increase the efficiency of a Sun StorageTek EBS server.

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 Sun StorageTek EBS server is determined by the browse and retention policies associated with each client. Chapter 5, "Browse and

Retention Polices" provides information about browse and retention policies.

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 (/fs1) with a single backup schedule in one group. Then, configure the second Client resource to back up another file system (/fs2) with a second backup schedule in another group.

Note: 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 18 on page 186 describes the key components of a Schedule resource.

Table 18 Key components of a schedule

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.
Calendar	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: • Full • Incremental • Level (1 – 9) • Consolidated • Skip 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. "Backup levels" on page 190 and "Save set consolidation" on page 198 provides more information on how backup levels work.

Note: The Force Incremental attribute in the backup group, determines the level used by the Sun StorageTek EBS 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 175 provides more information.

Working with schedules

This section provides information on how to edit, delete, and copy schedules.

Note: "Task 1: Set up a schedule for backups" on page 77 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

Note: 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 650 provides more information.
2. Clear the **Scheduled Backup** attribute and click **OK**.

Client timezone adjustment for Windows clients

The Sun StorageTek EBS software enables you to schedule backups over long-distance networks. In some cases, scheduling can occur over multiple time zones.

To incorporate time zones into your backup schedules, adjust each client time zone setting by editing the client's AUTOEXEC.BAT file.

Adjusting the client time zone

To adjust the time zone for backups of a remote client:

1. Open the AUTOEXEC.BAT file by using a text editor.
2. Type these lines in the AUTOEXEC.BAT file:

```
TZ = <local_standard_timezone><+/- gmt_differences>
<local_daylight_savings_timezone>
```

Table 19 on page 189 describes the time zone strings that can be adjusted in the AUTOEXEC.BAT file.

Table 19 Time zone strings

String	Description	Example
local_standard_timezone	The current standard time zone. All three characters are required.	PST is the three-character abbreviation for Pacific Standard Time.
+/- gmt_differences	The difference in the local standard time zones from Greenwich Mean Time (GMT) in hours. Positive numbers adjust westward from GMT and negative numbers adjust eastward.	The number 8 would represent Pacific Standard Time, while -1 would represent the central part of Europe.
local_daylight_savings_timezone	(Optional) The current daylight savings time zone.	PDT is the three-character abbreviation for Pacific Daylight Time.

A complete example of a time zone string is:

```
PST8PDT
```

If the TZ environment string is not present in the AUTOEXEC.BAT file, the Sun StorageTek EBS software presumes a default setting of TZ = EST5EDT.

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 20 on page 190 describes the five kinds of backup levels:

Table 20 Backup levels

Backup Level	Function
Full	Backs up all files, regardless of whether or not they have changed.
Level [1 – 9]	Backs up files that have changed since the last backup with a lower-numbered backup level. For example: • 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. **Note:** The Sun StorageTek EBS software ignores any incremental-level backups when determining what files should be backed up.
Incremental	Backs up files that have changed since the last backup, regardless of level.
Consolidated	Backs up all data that has changed since last full backup and then merges these changes with the last full backup. Information is available in "Save set consolidation" on page 198.
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 "Backup levels for deduplication clients" on page 857. Chapter 26, "Support for Avamar Deduplication," provides overview information about deduplication nodes.

How Sun StorageTek EBS 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 6, "Directives."

The following three examples illustrate the how backup levels affect the requirements for data recovery.

Example 10 Backup levels (part 1)

As shown in Figure 15 on page 191, 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.

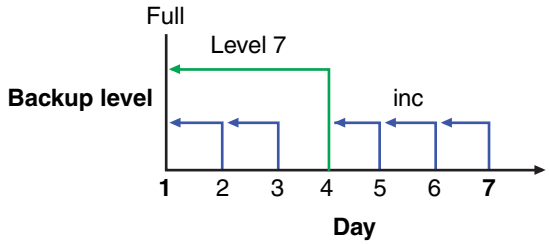


Figure 15 Backups for October 1 through October 7

Example 11 **Backup levels (part 2)**

Figure 16 on page 192 continues the example illustrated in Figure 15 on page 191 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 volume. 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.

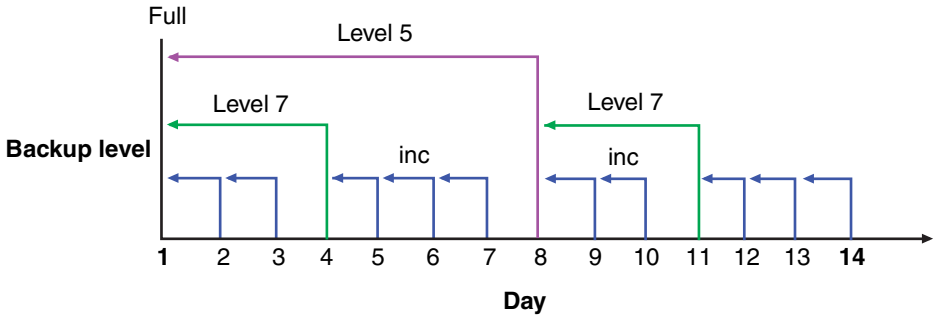


Figure 16 **Backups for October 1 through October 14**

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 12 **Backup levels (part 3)**

Figure 17 on page 193 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.

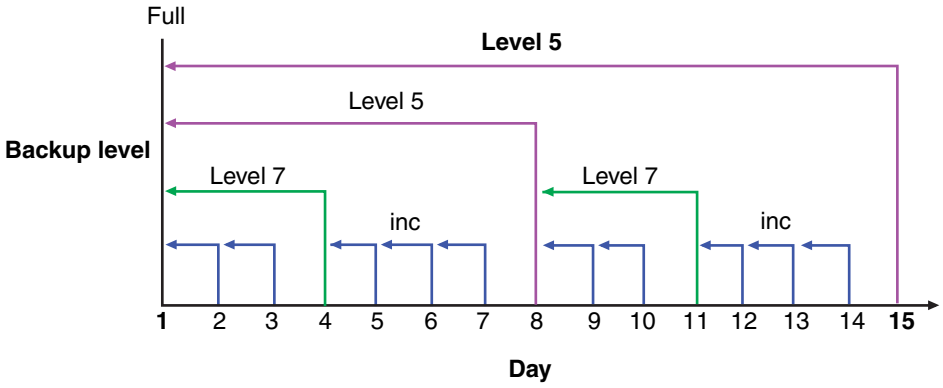


Figure 17 Backups for October 1 through October 15

The Sun StorageTek EBS 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.

- Consolidated backups provide the same benefits at the same cost as do full backups. The difference is that consolidated backups are less taxing on the network and client because only a level 1 backup is performed. The server's performance, however, is slower because the server consolidates the changed data with the most recent full backup.

Table 21 on page 194 lists advantages and disadvantages of each backup level.

Table 21 Advantages and disadvantages of backup levels

Backup level	Advantages	Disadvantages
Full	• Faster recovery	• Slow backup • High server load • High load on client and network • Uses more volume space
Level	• Faster backup • Low load on server • Uses least volume space	• Slow recovery • Data can spread across multiple volumes
Consolidated	• Faster backup (from the client's perspective only) • Faster recovery • Low load on client and network	• Longest high load on server • Requires at least two volume drives • Uses most volume space

Online indexes and backup levels

The online indexes are backed up in this way:

- The client file index is backed up every time the client is backed up, and is backed up at the same level as the client.
- The server index is backed up every time the server is backed up, and it is backed up at the same level as the server.
- The media database is backed up whenever the server is backed up, or after every scheduled backup if the server is not in an active group.

For example:

- If the Sun StorageTek EBS server is backed up at the level full, the backup levels of the server index and media database are also full.
- If the Sun StorageTek EBS server's backup is a level 5, the backups of the server index is also a level 5.

However, when the server's backup level is incremental, the backups of the server index is level 9. For a consolidated backup, the backups of the indexes are level full backups. The server does not perform a level 1 backup for this data.

Backup levels for the Sun StorageTek EBS online indexes

The Sun StorageTek EBS 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. Table 22 on page 195 compares the level at which the Sun StorageTek EBS server backs up regular files and the online indexes.

Table 22 Regular file and index backup levels

Regular files	Online indexes
Full	Full
Level 1–9	Level 1–9
Incremental	Level 9
Manual (using the User program)	Not saved

The Windows Archive attribute

The Windows file Archive attribute is used by the Sun StorageTek EBS software to help determine if the file should be backed up. The Sun StorageTek EBS software backs up a file if the Archive attribute is set. This functionality can be disabled by setting the environment variable NSR_AVOID_ARCHIVE to any value other than No. To set this as a system environment variable, use the Control Panel's System applet.

Note: You must restart the Sun StorageTek EBS Backup and Recover Server service for the environment variable to take effect.

- After backing up the file, the Sun StorageTek EBS software turns off the Archive attribute.
- After restoring the file, the Sun StorageTek EBS software turns on the Archive attribute.

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 Microsoft Windows system files. More information about these save sets is available in Appendix A, "SYSTEM and VSS SYSTEM Save Sets."

Note: A backup or restore of any SYSTEM or VSS SYSTEM save set automatically includes all eligible components of that save set.

Note: 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 2000, 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 Sun StorageTek EBS 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 Sun StorageTek EBS 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 always backed up at level full, even if another backup level is requested. There are no individual components to back up or restore.

Save set consolidation

Save set consolidation merges a new level 1 backup with the last full backup of a save set to create a new full backup. Consolidation is a level full backup. The resulting save set is the same as a level full save set. Although a consolidated backup took place, there is no such thing as a consolidated save set.

Save set consolidation eliminates the need to perform full backups at regular intervals. After scheduling a save set consolidation backup, you only perform one full backup during the first scheduled consolidated backup cycle. Afterward, all subsequent backups of the consolidated save set are incremental backups.

Save set consolidation shifts the burden of backups to the Sun StorageTek EBS server, which reduces Client resource use and network traffic. This shift to the server also results in more frequent level saves, which results in fewer tapes to process in the event of a full recovery.

When working with large files, save set consolidation shortens the backup window by doing incremental saves, thus reducing the number of tape drives that are required during a recovery.

Note: If a save set consolidation is performed on a save set that has no existing full backup, the consolidation backup defaults to a level full backup and the browse and retention policy is set to one week. Once a full backup exists, subsequent consolidation backups are performed as a level 1 followed by consolidation with the existing full backup.

Requirements

Save set consolidation requires at least three attached devices: two for reading and one for writing.

Save set consolidation usage

Use save set consolidation if these conditions are present:

- 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.
- Either the network bandwidth is small, or large backups over the network are cost-prohibitive.
- You need to back up large file systems with little incremental data.
- The server has the necessary resources (a minimum of two volume drives and preferably three or more volume drives) and the workload capacity to consolidate full backups locally.

Note: Under the appropriate conditions, save set consolidation can free network bandwidth and Client resources. Used inappropriately, save set consolidation can lower performance, since it uses tape-to-tape transfer, which might be slower than the disk-to-tape data transfer that is used by backups.

Limitations of save set consolidation

These are some of the limitations that can cause the save set consolidation process to fail:

- Save sets cannot be consolidated when either the full or Level 1 save set contains a file larger than 2 GB.
- Raw disk file partitions are not eligible for consolidation.
- Database systems cannot be consolidated.
- Renamed and deleted directories cannot be consolidated. If incremental data contains renamed or deleted directories, the save set consolidation process will detect it and abort.

Additionally, do not use save set consolidation if any of the following conditions are present:

- The client is connected to the server over a fast network or is a local client, and the network data traffic generated by full backups is not a problem. In this case, save set consolidation will not produce a measurable benefit.
- The file systems being backed up are either not very large or they contain a large number of small files that are modified often.
- The incremental data usually contains a large amount of data, and the number of files that have changed since the last full backup is large compared to the total number of files in the file system.
- It is cost-prohibitive to allocate three tape drives for the exclusive use of the server while it consolidates the full backup.

Note: If a consolidated backup cannot be completed because of a system crash, tape drive errors, or other problem, data integrity is assured. The transactional log assists save set consolidation in putting back changes made to the online index.

Performing save set consolidation

You can schedule save set consolidation from the Administration window, or perform a manual consolidation from the command prompt.

Setting up save set consolidation on a schedule

To set up save set consolidation on a schedule:

1. Open the schedule to which save set consolidation will be applied. More information is available in "Editing a schedule" on page 187.
2. Select the day the save set consolidation is to occur.
3. Right-click and select **Consolidate** from the **Set Level** menu.

Note: Renamed or deleted directories cannot be consolidated. If incremental data contains renamed or deleted directories, the save set consolidation aborts.

Perform a manual save set consolidation

You can invoke save set consolidation through two different command line programs:

- The **savegrp** program

  The **savegrp** program must be used with the **-lc** option to indicate that the backup level is consolidated, or with a schedule that has a level consolidate on the schedule. This program automatically performs a level 1 backup, followed by the actual consolidation process.

- The **nsrssc** program

  The **nsrssc** program completes the consolidation process. For the **nsrssc** program to be successful, a level 1 save set and a level full save set must already exist.

Note: When **nsrssc** is run from the command prompt, bootstrap information is not written to the media used for consolidation. As a result, you still need the original media for disaster recovery.

Both programs also offer other options to maximize the flexibility of save set consolidation. For more information about the **savegrp** and **nsrssc** commands, refer to the *EMC NetWorker Command Reference Guide* or the UNIX man pages.

Directing data from a consolidation backup to a specific pool

By default, save sets from a consolidation backup are written to whatever media is mounted for the group that was most recently backed up.

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 78.
2. Create a **Pool** for consolidation backups, as described in "Creating a media pool" on page 251.
3. In the **Pool** resource, select the group that was created in step 1 as the value for the **Groups** attribute.
4. Edit the **Group** attribute in each **Client** resource that is to have consolidated backups, and assign the group that you created in step 1 .

Example 13 Consolidation of level 1 and full backups

On October 16, the level consolidated performs a level 1 backup, which backs up all the data that has changed since the last full backup. The level 1 backup is merged with the last full backup and builds a new level full backup. To recover from a disk crash on October 17, use the volume that was created on October 16.

Figure 18 on page 203 illustrates how the Sun StorageTek EBS server merges a level 1 backup with an existing full level backup to create another level full backup.

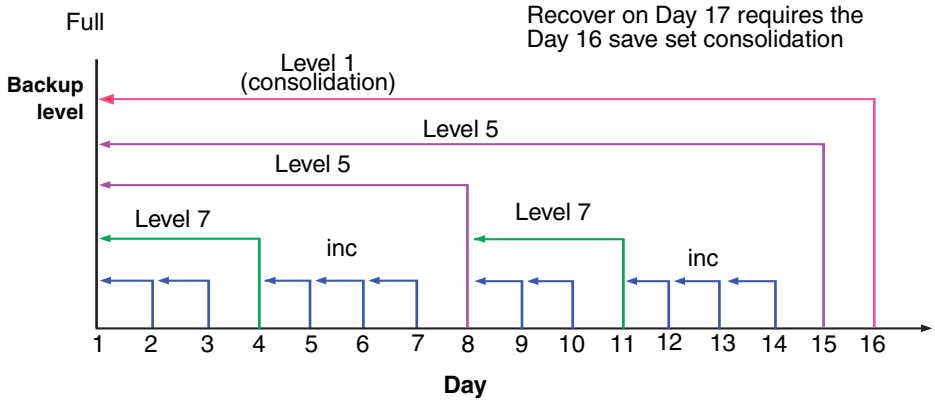


Figure 18 Consolidated backup

5

Browse and Retention Polices

This chapter covers these topics:

- About browse and retention policies 206
- Managing the data lifecycle..... 213
- Browse and retention policies for manual backups..... 219
- Modifying the browse and retention policy on a save set 220

About browse and retention policies

The *browse policy* determines how long files are maintained in the client's file index on the Sun StorageTek EBS server. During the period of the browse policy, users can browse backed-up data from the Sun StorageTek EBS 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 Sun StorageTek EBS 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 Sun StorageTek EBS 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 Sun StorageTek EBS client computer. Users can browse and mark files, and initiate data recovery. The Sun StorageTek EBS 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 180 provides information about schedules and "Backup levels" on page 190 provides information about backup levels

Example 14 One-week browse policy

In Figure 19 on page 207, 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.

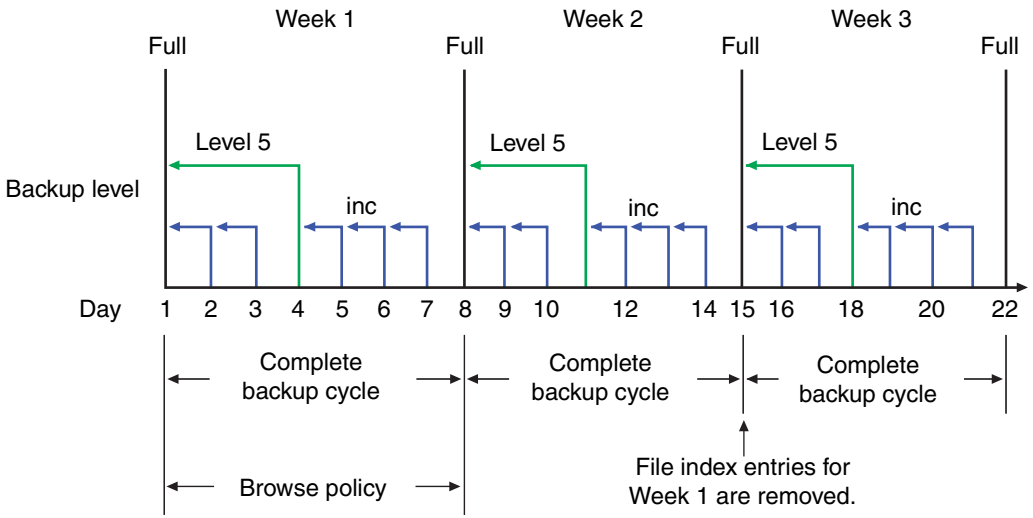


Figure 19 One-week browse policy

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 19 on page 207, entries from the Week 1 backup cycle are removed from the client file index on October 15.

Example 15 Two-week browse policy

In Figure 20 on page 209, 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.

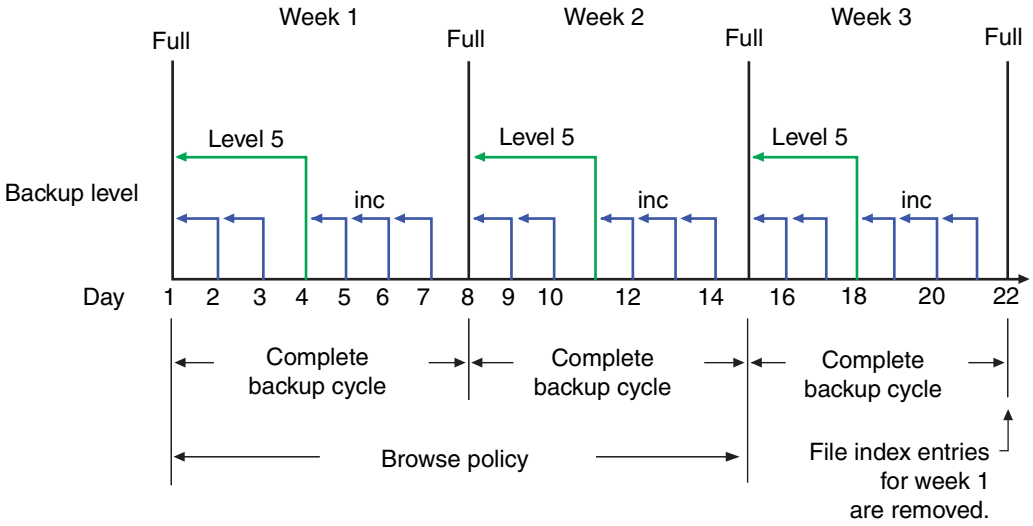


Figure 20 Two-week browse policy

Retention policies

The Sun StorageTek EBS 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 Sun StorageTek EBS 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 policy. 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 Sun StorageTek EBS 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 386 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 Sun StorageTek EBS 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.

"Restore a save set entry in the online indexes" on page 454 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 16 **Three-week retention policy**

Figure 21 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.

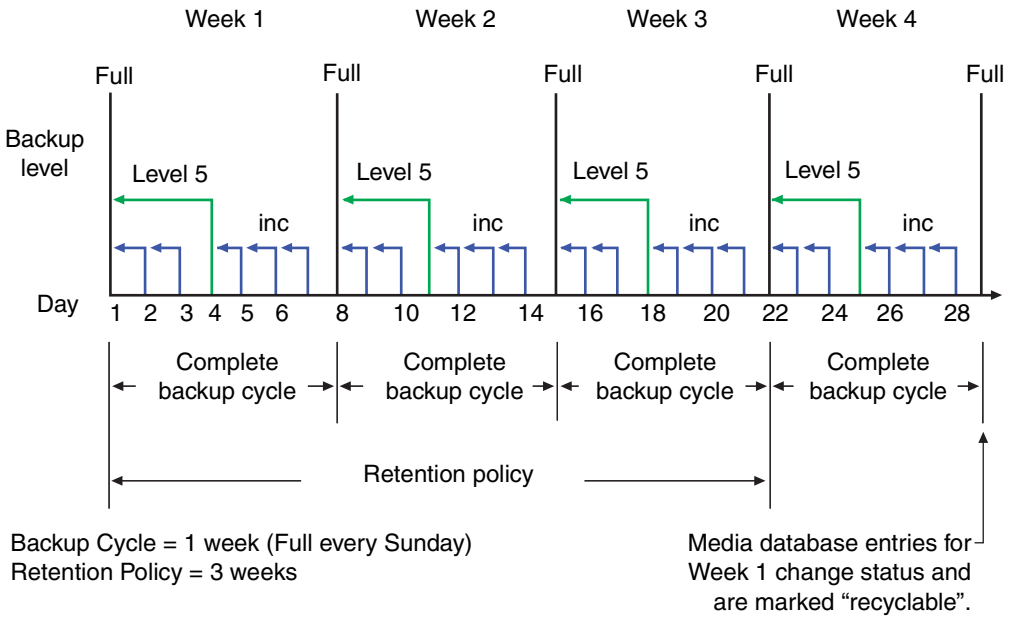


Figure 21 **One-week backup cycle and three-week retention policy**

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.

Note: 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 390
- "Schedules" on page 180
- "Backup levels" on page 190

Retention policies for client file index save sets

The client file indexes that reside on the Sun StorageTek EBS 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 Sun StorageTek EBS client to which the client file index corresponds. For example, if Sun StorageTek EBS 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 Sun StorageTek EBS server. This ensures that if a Sun StorageTek EBS 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 Sun StorageTek EBS software will 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 retention policies for clone data" on page 290 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 219 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 22 on page 214 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 Sun StorageTek EBS 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 Sun StorageTek EBS client computer. "Recovering by save set selection" on page 436 and "Restore a save

set entry in the online indexes" on page 454 provide more information.

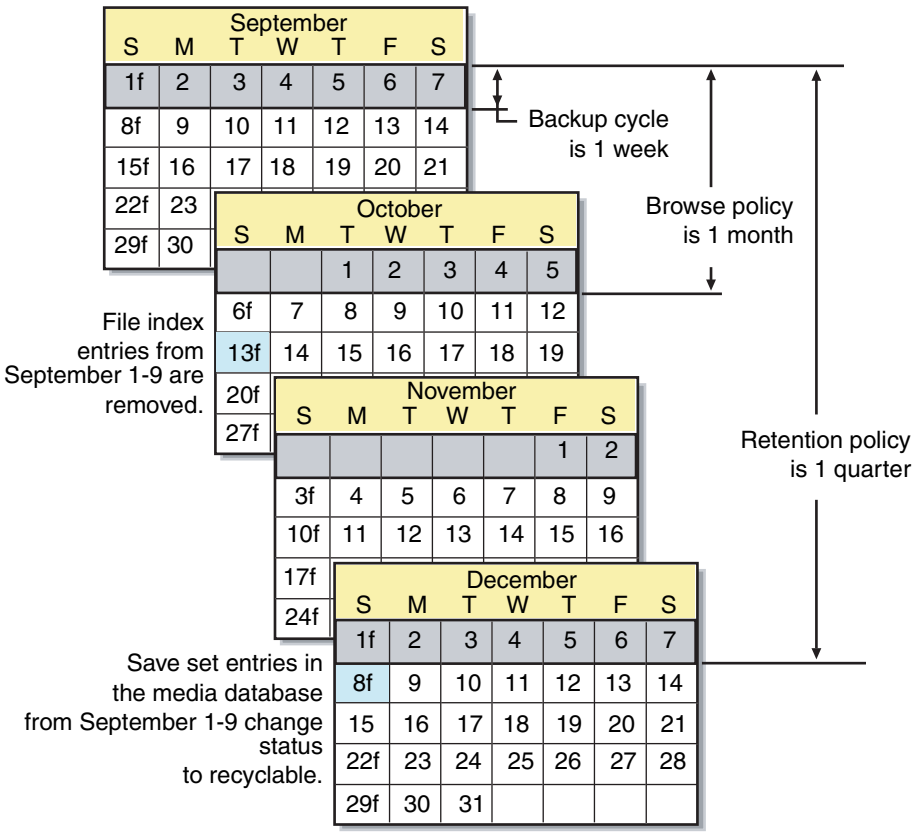


Figure 22 Data lifecycle in the client index and the media database

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.

Note: 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 Sun StorageTek EBS 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 Sun StorageTek EBS 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 17 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 policies

Sun StorageTek EBS 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 79 provides information about creating a policy.

Editing a policy

Note: You *cannot* change the name of a policy. To rename a policy, delete the current policy and create a new one.

To edit a policy:

1. In the **Administration** window, click **Configuration**.
2. In the expanded left pane, select **Policies**.
3. In the right pane, select the 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 policy

Note: Preconfigured policies cannot be deleted. For more information, see "Preconfigured policies" on page 215.

To delete a policy:

1. In the **Administration** window, click **Configuration**.
2. In the expanded left pane, select **Policies**.
3. In the right pane, select the 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 PowerSnap™ Module software. 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 NetWorker PowerSnap Module software and creating a snapshot policy, refer to the EMC PowerSnap Module documentation. 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 Sun StorageTek EBS software provides two preconfigured policies that can be used with the PowerSnap Module software:

- Serverless

  With the Serverless snapshot policy, a single snapshot is taken per day. The data is then backed up to traditional tape 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 PowerSnap Module documentation.

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 Sun StorageTek EBS host, the Client resource with the longest browse time is adopted. For example, if there are three Client resources for the Sun StorageTek EBS 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 retention policies for clone data" on page 290 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 18 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**.

6 Directives

This chapter covers these topics:

- Directives overview 224
- Types of local and global directives 224
- Creating a global directive resource 225
- Editing a global directive resource 225
- Deleting a global directive resource 226
- Copying a global directive resource 226
- Applying directives 226
- Local directives within the NetWorker User program 228
- Preconfigured global directive resources 230
- Format of directive statements 232

Directives overview

Directives are resources that contains special instructions that control how the Sun StorageTek EBS server processes files and directories during backup and recovery. Sun StorageTek EBS administrators can create directives to customize the Sun StorageTek EBS software to your specific needs, maximize the efficiency of backups, and apply special handling to individual files or directories.

Note: If deduplication is used for a client, directives cannot be applied.

Types of local and global directives

There are three types of directives.

Global directives

Administrators can create global directives by using the **Sun StorageTek EBS Administrator** program. These directives are stored as resources on the Sun StorageTek EBS server, and can be selectively applied to individual clients by using the Directive attribute of the Client resource.

Sun StorageTek EBS 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. The directives these files 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 232 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 Sun StorageTek EBS Client resource. Example 19, "Applying a global directive" provides more information.

Editing a global directive resource

Note: 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

Note: 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**.

Applying directives

This section contains some basic examples of global and local directives, and describes how to apply them to Sun StorageTek EBS clients.

Note: If deduplication is used for a client, directives cannot be applied.

Example 19 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 Sun StorageTek EBS 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 Sun StorageTek EBS 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 Sun StorageTek EBS client, all files that match the *.tmp pattern in the specified directories will be skipped.

Example 20 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 Sun StorageTek EBS 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 Sun StorageTek EBS 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 228 provides information on creating local directives.

Example 21 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 Sun StorageTek EBS client. When a scheduled backup is performed on the Sun StorageTek EBS 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 Sun StorageTek EBS 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 Sun StorageTek EBS 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 a Sun StorageTek EBS **save** 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 96 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\Legato

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 Sun StorageTek EBS software comes with a number of preconfigured global Directive resources. All preconfigured Directive resources can be modified, but they cannot be deleted.

Table 23 lists the preconfigured directives and their descriptions.

Table 23 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 96.
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 23 Preconfigured directives (2 of 3)

Directive resource	Description
Mac OS standard	Contains a set of directives used to back up standard Mac OS clients. Applies these ASMs: • The **skip** ASM is applied to these files and directories: /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 **swapasm** 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 23 Preconfigured directives (3 of 3)

Directive resource	Description
UNIX standard	Contains a set of directives used to back up standard UNIX clients. Applies these ASMs: • 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 **swapasm** 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	Contains the same set of directives as the UNIX standard directive, along with applying the **compressasm** ASM to all files. **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.
VCB directives	VCB directives are valid for backing up virtual machines using the VCB methodology. This directive is supported in the following scenarios: • When file level incremental backups are performed instead of FULL image level backups. • When FULL file level or incremental file level backups are performed when the save set is ALLVMFS. The vcb directive skips the following files and folders: • pagefile.sys • hiberfil.sys (Hibernation file) • WINDOWS\system folder • WINDOWS\System32 folder

Format of directive statements

The format of a directives uses three primary types of statements:

- Directory specifications
- Application-Specific Module (ASM) specifications
- Save environment keywords

Comments can be included in directives as well. Any text after a hash (#) character is treated as a comment.

Note: For more information about directive formats, refer to the nsr and nsr_directive information in the *EMC NetWorker Command Reference Guide* or the UNIX man pages.

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 234 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 24 on page 235 lists the three save environment keywords.

Table 24 Save environment keywords

Keyword	Instruction	Example

forget	Instructs the Sun StorageTek EBS server to no longer apply inherited directives (those directives that begin with a +).	To skip all *.o files in the directory /SRC, except those in the /SRC/SYS directory, type:
		`<</SRC>>`
		`+skip: *.o`
		`<</SRC/SYS>>`
		`forget`
		This uses the **skip** ASM to instruct the Sun StorageTek EBS 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.

ignore	Instructs the Sun StorageTek EBS server to ignore all directives applied to the subdirectories below the current directory.	To override any local directives set in users' home directories, type:
		`<</HOME>>`
		`ignore`

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 /HOME/DOC directory to be applied, type:
		`<</HOME>>`
		`ignore`
		`<</HOME/DOC>>`
		`allow`

Application Specific Modules (ASMs)

Directives use Application Specific Modules (ASMs) to process files and directories. ASMs are programs that operate within the Sun StorageTek EBS environment to perform various tasks during backup and recovery. For example, the **compressasm** program instructs the Sun StorageTek EBS 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 25 describes the Sun StorageTek EBS ASMs.

Table 25 Sun StorageTek EBS ASMs (1 of 3)

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 96 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	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 uses considerable amounts of CPU resources, so its benefits could be limited on low-powered systems. If your storage device compresses data, you can also apply this ASM. However, the additional compression might actually result in slightly more data being written to tape.	`+compressasm: .`
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 Sun StorageTek EBS 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`
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`

Table 25 Sun StorageTek EBS ASMs (2 of 3)

ASM name	Description	Example
nsrindexasm	Used to recover from Sun StorageTek EBS file index backups performed by Sun StorageTek EBS 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 Sun StorageTek EBS 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 Sun StorageTek EBS 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 238 provides more information.	`rawasm: /dev/oracle1`

Table 25 Sun StorageTek EBS ASMs (3 of 3)

ASM name	Description	Example
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 120 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 234 provides more information.

7

Sorting Backup Data

This chapter covers these topics:

- Media pools 242
- Label templates 260

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 Sun StorageTek EBS 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 Sun StorageTek EBS 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 260 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 Sun StorageTek EBS 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 Sun StorageTek EBS server writes data to the volume.
- If there is no correctly labeled volume mounted on a storage device, the Sun StorageTek EBS 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 Sun StorageTek EBS 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 Sun StorageTek EBS 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
```

Sun StorageTek EBS media pool types

Sun StorageTek EBS software contains preconfigured media pool types to keep different types of data separate. Sun StorageTek EBS 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 251 provides information on customizing media pools.

Sorting data with media pools

When you configure the Sun StorageTek EBS server, you can create additional media pools and sort data by media pool type and any combination of the following:

- Group (backup group)
- Sun StorageTek EBS 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 4, "Backup Groups and Schedules" provides information about groups and backup levels.

Example 22 Using media pool configurations to sort data

Figure 23 on page 245 illustrates how the Sun StorageTek EBS 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 Sun StorageTek EBS server looks for a media pool configuration that matches the group named Accounting and the level full. When the Sun StorageTek EBS 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.

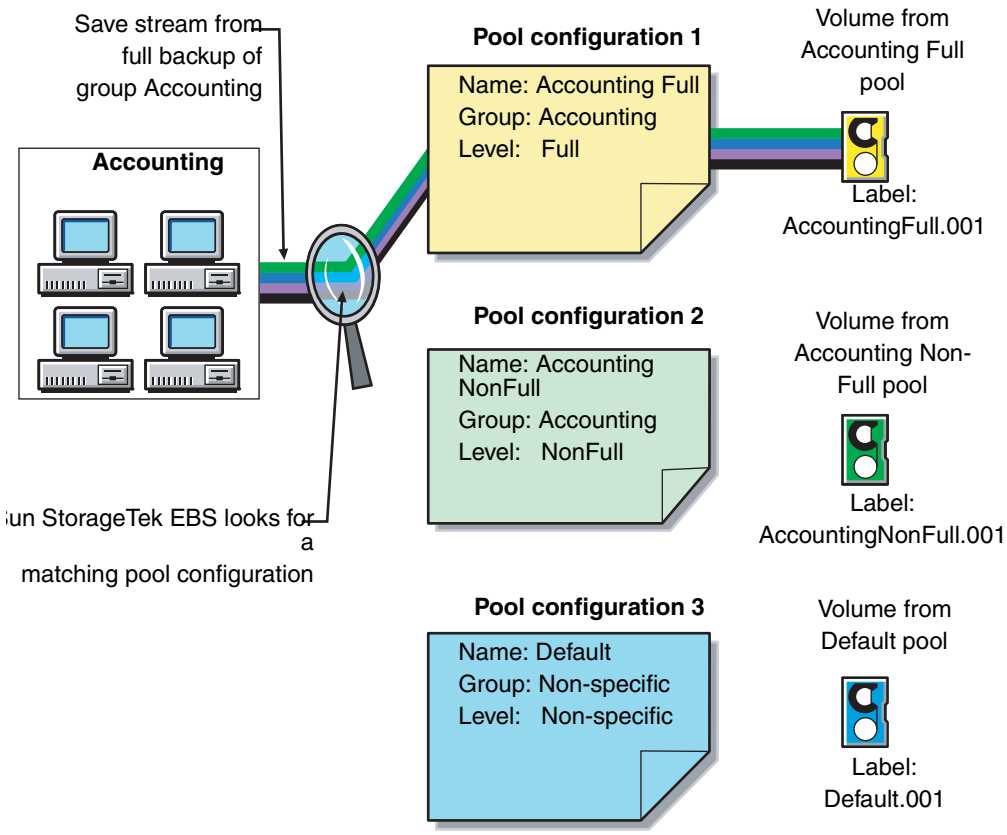


Figure 23 **Using media pool configurations to sort data**

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 23 **Sending bootstrap information and all the client file indexes to the same media pool**

In this example, the client file indexes are in index. To send the Sun StorageTek EBS server's bootstrap and all the client file indexes from this file system to the same media pool, create a media pool (in the

Media Pool resource) with the values listed for the attributes in Table 26 on page 246.

Table 26 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 Sun StorageTek EBS 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 172 provides information.
2. Configure a **Media Pool** resource for consolidated backups."Creating a media pool" on page 251 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 Sun StorageTek EBS 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 27 on page 247 details the hierarchy that the Sun StorageTek EBS 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 27 Sun StorageTek EBS 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		
	x		x	x

Table 27 Sun StorageTek EBS 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	
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 Sun StorageTek EBS 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 Sun StorageTek EBS 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 Sun StorageTek EBS 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 <Sun StorageTek EBS_install_dir>\logs directory. "Viewing log files" on page 904 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 Sun StorageTek EBS 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 Sun StorageTek EBS 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 Sun StorageTek EBS 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 9, “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 8, “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 266 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 Sun StorageTek EBS server notifies you that it will create a label template for the media pool.

To have the Sun StorageTek EBS server create the label template, click **OK**. "Using label templates" on page 260 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.

   Note: 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 253.
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 253.

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 **Sun StorageTek EBS 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 Sun StorageTek EBS server verifies data written to volumes from this media pool.

Data is verified by repositioning the 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.

Supporting WORM and DLTWORM tape drives

Sun StorageTek EBS 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 28 on page 254 describes the WORM devices that are supported by the Sun StorageTek EBS software.

Table 28 WORM supported devices (1 of 2)

Device	Description
HP LTO Ultrium 3	Unique to HP Ultrium-3 • Inquiry VPD page 0xb0, byte 4 bit 0 indicates WORM capable • Read attribute # 0x0408 bit 7 to indicate WORM media present
Quantum SDLT600, DLT-S4, and DLT-V4 (SCSI and SATA)	Any drive with product inquiry data of "*DLT*" tape drive that reports WORM capability the way these drives do. ("Quantum" not required in the vendor inquiry data.) • Inquiry data VPD page 0xc0, byte 2, bit 0 to indicate WORM capable • Read attribute # 0x0408 bit 7 to indicate WORM media present
Sony AIT-2, AIT-3, AIT-4, and SAIT	Any drive with "Sony" in the vendor inquiry data that reports WORM capability like these drives do: • Mode sense page 0x31, byte 5 bit 0 indicates WORM capable • Mode sense byte 4 bit 6 indicates WORM tape present

Table 28 WORM supported devices (2 of 2)

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

The WORM and DLTWORM attributes determine whether or not the Sun StorageTek EBS 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 Sun StorageTek EBS 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 24 on page 256, a (W) is appended to the volume names displayed in the Sun StorageTek EBS Administrator window. If a volume is both read-only and WORM, an (R) is appended to the volume name.

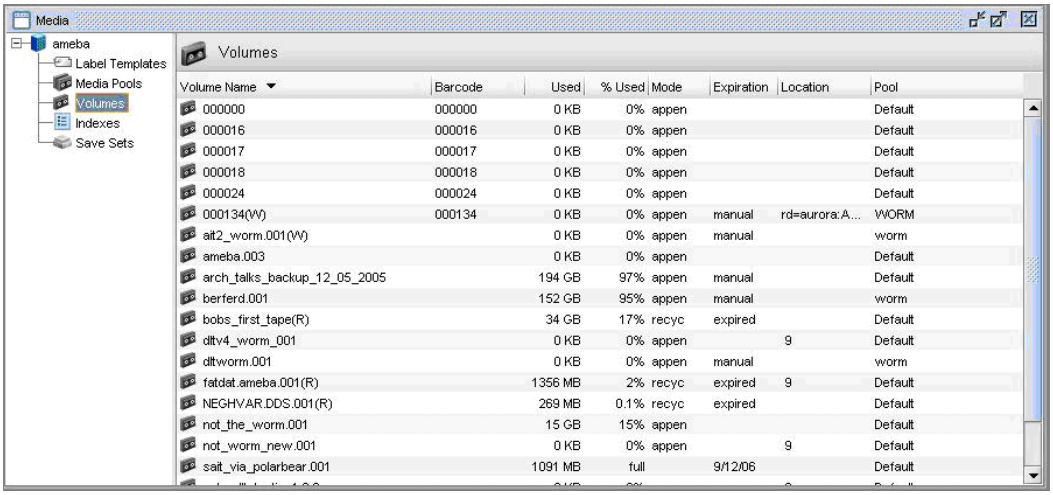


Figure 24 Identifying WORM tapes in the Sun StorageTek EBS Console

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 Sun StorageTek EBS 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. Sun StorageTek EBS 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 29 WORM/DLTWORM attributes (1 of 2)

Attribute	Description
WORM pools only hold WORM tape	By default, the Sun StorageTek EBS 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, Sun StorageTek EBS only lets you label WORM tapes into WORM pools. Clear this option when: 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 29 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	If selected, before the Sun StorageTek EBS software labels a tape in a drive capable of creating DLTWORM volumes, Sun StorageTek EBS 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. **Note:** Refer to the Quantum web site for information on which tapes can be converted to DLTWORM tapes. Not all firmware revisions for all of these devices support WORM operation. Check the tape drives website to make sure that your drive has up-to-date firmware.

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 Sun StorageTek EBS 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 251 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 Sun StorageTek EBS 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 Sun StorageTek EBS 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 Sun StorageTek EBS server uses the preconfigured Default media pool and corresponding Default label template.

Figure 25 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.

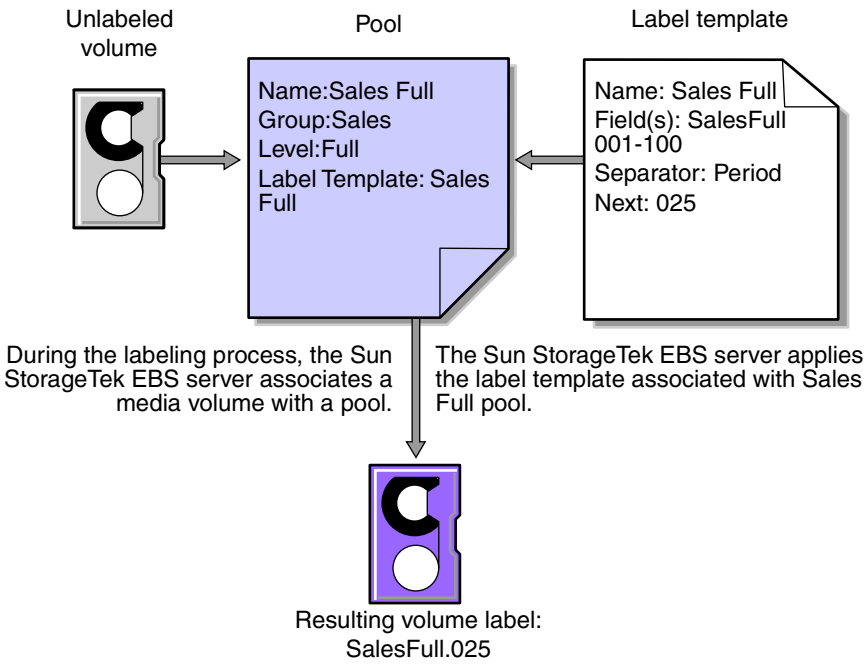


Figure 25 Labeling a volume by using a label template

Using preconfigured label templates

The Sun StorageTek EBS 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 Sun StorageTek EBS 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 30 describes how to complete the key attributes for this resource. "Creating a label

template" on page 266 provides more information on creating a label template.

Table 30 Key label template attributes

Attribute	Guidelines
Name	Keep the label name consistent with the media pool name, so that the label name reflects how the data is organized. For example, a label template named "AcctFull" would identify volumes that belong to a media pool called "Accounting Full." Do not use these characters in label template names: / \ * ? [ ] ( ) $! ^ ; ' " ' ~ < > & | { } : - . _
Fields	A label template is made up of one or more fields. Each field, or component, provides a layer of specificity to your organizational structure. There can be any number of components, but it is best to keep the template simple with as few as necessary. The label cannot exceed 64 characters. You can use four types of components: • Range of numbers (for example, 001-999) • Range of lowercase letters (for example, aa-zz) • Range of uppercase letters (for example, AA-ZZ) • Character string (for example, Accounting) Each range includes a start value, a dash (-), and an end value. The start value and the end value must have the same number of characters. For example, use 01-99 (not 1-99) or aaa-zzz (not aa-zzz). The order in which you enter each component of the Field attribute is important. The Sun StorageTek EBS server applies each component in a left-to-right order, starting with the first one entered. Table 31 on page 264 illustrates how label templates use components to create a number sequence for volume labels.
Separator	Choose the symbol to appear between component entries. Use the period, dash, colon, or underscore to separate each component of the label template. If label components do not have separators (for example, AA00aa), the labels can be difficult to read.
Next	Choose the next sequence number to write on the label that the Sun StorageTek EBS server places on a volume (according to the template). • To force a label to start the label scheme at a particular point, type a start label value. The server continues to generate labels from that point on, according to the rules of the template. • To have the Sun StorageTek EBS server generate the first label, leave this attribute blank. When the Sun StorageTek EBS 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.

Table 31 lists examples of number sequences for volume labels.

Table 31 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 Sun StorageTek EBS 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 Sun StorageTek EBS 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 32 illustrates several ways you can use components to organize labels.

Table 32 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 Sun StorageTek EBS 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.

Note: 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 Sun StorageTek EBS 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 Sun StorageTek EBS. For example, if you physically label a volume mars.1, its Sun StorageTek EBS 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 30 on page 263.

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.

8 Archiving

This chapter covers these topics:

- Overview of archiving 270
- Permissions for archiving 273
- About archive pools 274
- Creating custom Archive pools 275
- Archiving data procedures 275
- Retrieving archived data 281
- Archive request management 285

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 Sun StorageTek EBS software components. Chapter 28, "Licensing" provides more information on licensing procedures.

Benefits of using the Sun StorageTek EBS 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.

Note: Sun StorageTek EBS software does not support archiving of SYSTEM or VSS SYSTEM save sets.

Archive requirements

Before Sun StorageTek EBS archive feature, ensure that you have:

- A device, either stand-alone or in an autochanger or silo, connected to a Sun StorageTek EBS 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. Chapter 28, "Licensing" provides more information.

How the Sun StorageTek EBS server archives data

The Sun StorageTek EBS 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 Sun StorageTek EBS server writes data to storage volumes that belong to an Archive pool. Information about archive data is tracked in the Sun StorageTek EBS 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 Sun StorageTek EBS 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 Sun StorageTek EBS Administrator program.

Figure 26 on page 272 illustrates how the Sun StorageTek EBS software archives data.

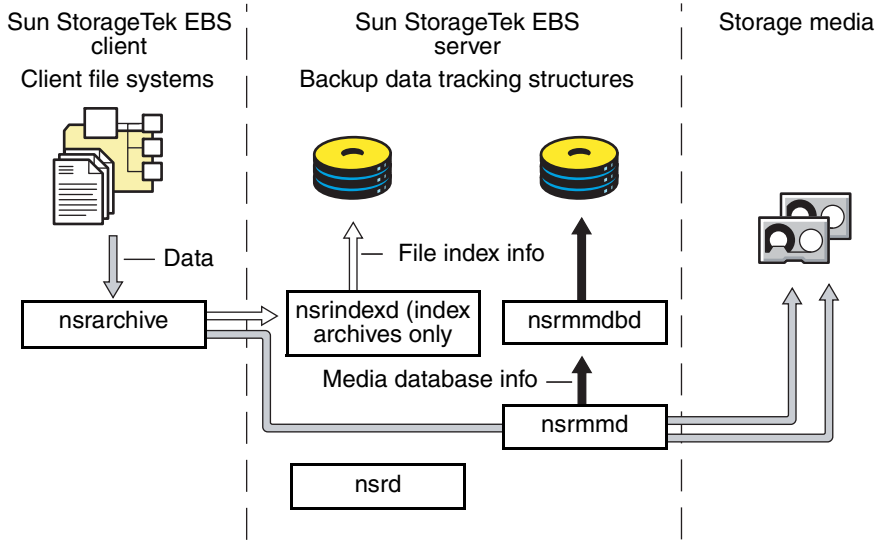


Figure 26 Overview of archive operation

Indexed and nonindexed archiving

The Sun StorageTek EBS 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 Sun StorageTek EBS server, all clients of that server are enabled for the Sun StorageTek EBS 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 Sun StorageTek EBS server, ensure that the Archive Services attribute is enabled in the Client resource for the server. "Enabling archive services for a Sun StorageTek EBS client" on page 276 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 Sun StorageTek EBS client software are installed on the same computer and both back up to the same Sun StorageTek EBS 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 attribute in the Client resource specifies the users who are allowed to archive data. If no users are listed in this attribute, only members of the Sun StorageTek EBS Administrators group with read permissions to the data can archive the data. To enable users who are *not* members of the Administrators group to archive data, include their usernames in this attribute. These users must have read permissions to the data and Backup Local Data privileges to archive

or retrieve the data. "Sun StorageTek EBS User Groups" on page 613 provides more information about the Administrators group and the Backup Local Data privilege.

Note: Only members of the Administrators group or users with the Change Security Settings permission enabled can change the Archive Users attribute.

Users can only retrieve data that they own. If other users need to retrieve data they do *not* own, public archives access must be enabled.

Enabling public archive access

To allow users listed in the Archive Users attribute of the Client resource to retrieve any archived files from that client:

1. In the **Administration** window, click **Configuration**.
2. In the left pane of the **Configuration** window, select the Sun StorageTek EBS 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 Sun StorageTek EBS 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 275 provides information on creating custom Archive pools.

If you do not specify a pool to store archived data, the Sun StorageTek EBS 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 Sun StorageTek EBS 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 242 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 Sun StorageTek EBS client

To enable archive services for a Sun StorageTek EBS 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. In the **Archive Users** attribute, type all appropriate users. "Enabling or restricting archive access" on page 273 provides details about setting this attribute.
5. 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 278 provides instructions.

Note: If the Sun StorageTek EBS client is set up for encryption with the **aes** ASM, archive data will also be encrypted. "Encrypting backup data" on page 96 provides information about setting up encryption for a Sun StorageTek EBS 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 Sun StorageTek EBS 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 278.

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 Sun StorageTek EBS server to check the integrity of the data on the storage volume, select **Verify**.
   - To instruct the Sun StorageTek EBS 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 Sun StorageTek EBS server appears in the **Archive Status** window, which monitors the progress of the archive. When the Sun StorageTek EBS 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 Sun StorageTek EBS software deletes archived files from the local disk.

Perform a manual archive from a Sun StorageTek EBS client on UNIX

To perform a manual archive from a UNIX client, use the **nsrarchive** 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 Sun StorageTek EBS client" on page 276 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 6, "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 **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 Sun StorageTek EBS server to use after completing the archive. "Notifications" on page 619 provides details.
14. Select the appropriate response for these options:
    - To instruct the Sun StorageTek EBS server to remove the archive files from the disk, select **Remove** from the **Grooming** list.
    - To instruct the Sun StorageTek EBS 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 286 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 285 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 Client resource's Archive Users attribute 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 Sun StorageTek EBS administrator or the owner of the file can retrieve the data.

"Permissions for archiving" on page 273 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 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 Sun StorageTek EBS 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. "Recovering by save set selection" on page 436 provides more information.

Retrieving nonindexed archives from a client on UNIX

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 **nwrecover** program. Type:

   `nwrecover -s` *servername*

where *servername* is the name of the server to use when recovering save sets. The /servers file, located in /nsr/res/servers, contains an entry for each available server.

If the **-s** option is not entered and there is only one server detected, that server will be connected to automatically. If there are no servers detected, or if there is more than one server available, the Change Server dialog box appears, which allows you to choose the server.

After the server is specified, the **nwrecover** program displays.

3. From the **Options** menu, select **Recover Archived Save Sets**.
4. Mark the save sets to be recovered by selecting the save sets and clicking the **Mark** button.
5. Click the **Start** button on the toolbar.

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**.

Recovering indexed archive data from a client on UNIX

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.

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 **nwrecover** program. Type:

   `nwrecover -s` *servername*

   where *servername* is the name of the server to use when recovering the data. The /servers file, located in /nsr/res/servers, contains an entry for each available server.

   If the **-s** option is not entered and there is only one server detected, that server will be connected to 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.

   After the server is specified, the **nwrecover** program appears.

2. From the **Options** menu, select **Recover Archived Files**.
3. Mark the files to be recovered by selecting the files and clicking **Mark**.
4. Click **Start**.

Retrieving archive data from the command prompt

One can also retrieve archive save sets by using the **nsrretrieve** command. For more information, refer to the *EMC NetWorker Command Reference Guide* or the UNIX man pages.

To identify an archive for retrieval, use the **-A** option to specify an annotation for the archive. However, when you use the **-A** option you must specify a string that uniquely identifies one annotation.

Example 24 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.

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.

9 Cloning

This chapter covers these topics:

- Overview of cloning 288
- Cloning compared to a basic copy 289
- Volume cloning 291
- Save set cloning 292
- Manual clone operations 301
- Recovering cloned data 301
- Cloning archived data 303
- Directing clones to a special storage node 304
- Using file type devices for clone operations 308
- Backup-to-tape for deduplication clients 309
- Using the nsrclone command 309

Overview of cloning

Cloning allows for secure offsite storage, transfer of data from one location to another and verification of backups. To perform cloning operations, you must be a Sun StorageTek EBS administrator. 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 Sun StorageTek EBS message file and the savegrp log file, which are located in the <Sun StorageTek EBS_install_dir>\logs directory. "Viewing log files" on page 904 provides information about viewing log files.

Note: Cloning works differently for deduplication backups. "Cloning and deduplication" on page 867 provides more information. It is also possible to send the backup data of deduplication nodes to tape volumes. "Backup-to-tape for deduplication clients" on page 868 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 Sun StorageTek EBS 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 Sun StorageTek EBS Administrators group. "Sun StorageTek EBS User Groups" on page 613 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.

Cloning performance

In general, a write occurs at the same speed whether it is part of a backup operation or a cloning operation. However, if a cloning operation is automatically requested as part of a scheduled backup, you may experience performance degradation in subsequent scheduled backups.

The Sun StorageTek EBS server attempts to complete one group's scheduled backup before initiating a scheduled backup for another group. However, the Sun StorageTek EBS server considers a group backup finished when the backup operations are complete, not when automatic cloning is complete. If another group starts its backup while the previous group's clone operation is under way, there may be contention for **nsrmmd** resources or specific volumes.

To avoid this problem, refrain from automatic cloning. Instead initiate a single cloning operation by using the **nsrclone** command as part of a scheduled task that runs at a nonpeak time after all backups are complete. "Using the nsrclone command" on page 309. provides more information about the **nsrclone** command.

Cloning compared to a basic copy

Cloning leaves traceable entries in both the client file index and the media database. Being able to track cloned data is an important difference between a cloning and a basic copy operation.

A single save set may be contained on one or more volumes. When a volume is cloned, all save sets that begin on the volume are cloned completely. As a result, it is possible that more than one volume is required for a single volume clone.

Sun StorageTek EBS software does not recognize a copied volume, because it does not have a valid Sun StorageTek EBS label. If Auto Media Management is enabled and a copied volume is left in the autochanger, the unlabeled volume is eligible for labeling and reuse during a scheduled backup. Chapter 12, "Media Management" provides more information about Auto Media Management.

It is possible to make an exact copy of a Sun StorageTek EBS archive volume. However, a duplicate of an archived save set does not include annotation information, which is used to locate an archive for recovery.

Specifying retention policies for clone data

The retention policy for clone data can be specified independently of the original save set. In this way, clone data can be recovered after the retention policy of the original save set has expired.

To specify the retention policy for clone data, do one of the following:

- Specify a retention policy in the Clone pool.
- Specify a retention policy from the command prompt.

Specify a retention policy for a Clone pool

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 and then click **OK**.

"Configuring media pools" on page 249 provides information about editing or creating a pool.

Specify a retention policy from the command prompt

To specify a retention policy from the command prompt, do 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 **nsrmm -e** command.

Note: Retention policies that are specified from the command prompt, override the retention policy specified for the clone pool.

You can also specify a browse policy for a clone save set by using 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 and all other clones of that save set.

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. Save sets can be automatically cloned as part of any of these operations, or manually cloned at another time.

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** and 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**.

Save set cloning

Save sets can be cloned automatically or manually. One can clone a single save set, or configure the software to clone full backups for an entire group.

When you specify that a particular save set is to be cloned, the Sun StorageTek EBS software first determines whether that save set already has a clone. If multiple clones of the save set exist, use the command-prompt to specify the save set clone as the source.

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 292 provides more information.

Cloning a save set manually

To manually clone a save set, first query the media database to locate the correct save set. After querying 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.

6. 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 Sun StorageTek EBS Console.

7. Use the **Status** attribute to limit the search to save sets that have a particular status. Table 33 lists the values that can be selected.

Table 33 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: • 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). • Scanned-in: Select if the save set was scanned in by using the **scanner** program. • 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.

8. 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.

9. 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** and then click **Yes** on the confirmation screen.

"Viewing manual clone history" on page 301 provides information on viewing the status of a manual clone operation or to cancel a clone operation that is in progress.

Cloning a save set automatically

Save set data can be automatically cloned when a scheduled group backup is run. Chapter 2, "Backing Up Data" provides information on configuring groups.

Automatic volume cloning is performed after all backup operations are complete. The savegroup completion report indicates the success or failure of the cloning operation.

The location of the devices where the clone data is written can be specified in the Client resource's Clone Storage Nodes attribute. "Directing clones to a special storage node" on page 304 provides detailed information about how devices are chosen for a read or write operation.

Examples of manual and automatic clone operations

These examples assume that there are two storage devices which contain mounted volumes: one for the source volume and another for the destination (clone) volume.

Example 25 Manual cloning of a save set

In this example, a user has requested that the Sun StorageTek EBS 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 Sun StorageTek EBS administrator must have this information:

- Sun StorageTek EBS 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.

Example 26 Automatic cloning of one client

In this example, a system administrator will clone data automatically, on a daily basis, for the Sun StorageTek EBS client **jupiter**, where the company product documentation resides. Client **jupiter** performs an incremental backup every night of the week except Friday, when it performs a full backup.

The system administrator must perform these tasks:

- Create a backup group named **Documents** to which the client jupiter will belong.
- Make the appropriate selections for the client jupiter, which includes selecting the group **Documents**.

To automatically clone jupiter's data, the system administrator must:

1. Create a group named **Documents**. "Task 2: Set up a group for backup clients" on page 78 provides more information on creating groups.

   Select the following choices for the group:

   - For the **Autostart** attribute, select **Enabled**.
   - For the **Clones** attribute, select **Yes**.
   - For the **Clone pool** attribute, select **Default Clone** (or a custom clone pool).

2. Create the client **jupiter**. "Task 6: Create a backup Client resource" on page 81 provides information on creating clients.

   Select the following backup choices for **jupiter**:

   - For the **Schedule** attribute, select **Full Every Friday**.
   - For the **Group** attribute, select **Documents**.

With this set of instructions, the administrator creates a group for just one client, so that only that client's data is cloned on a daily basis.

Example 27 Automatic cloning of several clients

Figure 27 on page 298 illustrates the level of control you can achieve with cloning. You can create an unlimited number of automatic cloning configurations by using Groups, Schedules, and the Client resources.

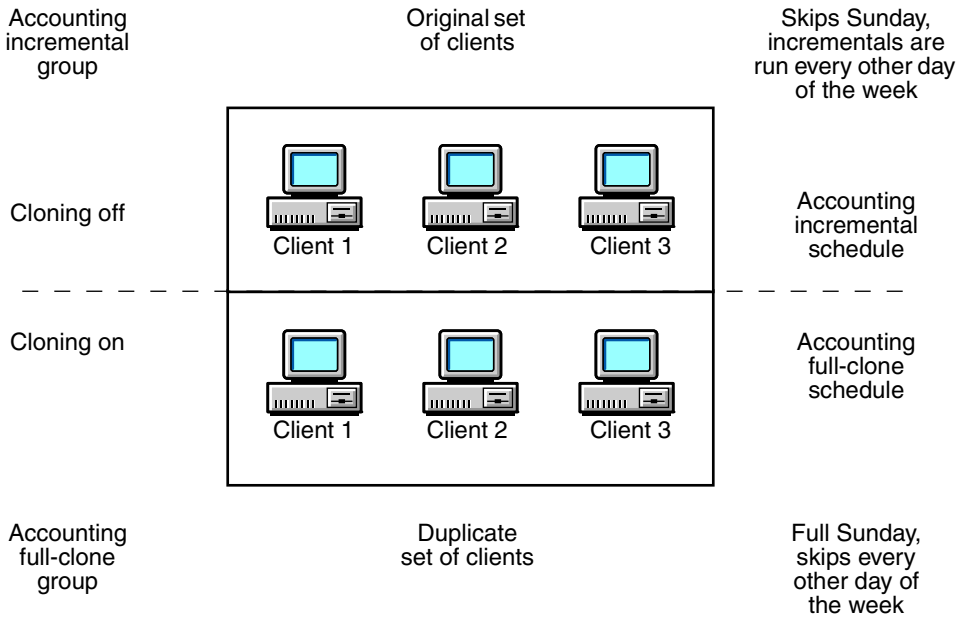


Figure 27 **Automatic cloning scheme**

In this example, the system administrator wants to automatically clone the data for a small group of Sun StorageTek EBS clients that contain accounting data. However, only clones of full backups are needed to store in an offsite vault.

Note: To keep full backups and incremental backups on separate volumes, the administrator must create two pools.

The system administrator must:

- Create two schedules: one to perform incremental backups for the clients and one to perform full backups and clones:

- Create one schedule named Accounting Incremental, as shown in Table 34.

Table 34 Accounting Incremental schedule

Sun	Mon	Tue	Wed	Thu	Fri	Sat
skip	incr	incr	incr	incr	incr	incr

- Create another schedule named Accounting Full-Clone, as shown in Table 35.

Table 35 Accounting Full-Clone schedule

Sun	Mon	Tue	Wed	Thu	Fri	Sat
full	skip	skip	skip	skip	skip	skip

Note: By creating two schedules and selecting them for the appropriate set of clients, the system administrator avoids cloning all of the incremental backups.

- Create two groups: one named Accounting Incremental and another named Accounting Full-Clone.
- Create two clients with the same name for each client host. In this example, one set of clients will use the Accounting Incremental schedule, and the other set will use the Accounting Full-Clone schedule.

To automatically clone data, the system administrator must:

1. Create two schedules:
   a. The Accounting Incremental schedule to perform incremental backups every day except Sunday, when it skips the backup. See Table 34 on page 299.
   b. The Accounting Full-Clone schedule to perform full backups every Sunday and skip backups the rest of the week. Table 35 on page 299 provides details.

   "Task 1: Set up a schedule for backups" on page 77 provides details on creating schedules.
2. Create two groups:
   a. The Accounting Incremental group with these attributes:

- For the **Autostart** attribute, select **Enabled**.
- Verify that the **Clones** attribute is unselected (or set to "No").

b. The Accounting Full-Clone group with these attributes:

- For the **Autostart** attribute, select **Enabled**.
- Select the **Clones** attribute (or set it to "Yes").
- For the **Clone pool** attribute, select **Default Clone**.

"Task 2: Set up a group for backup clients" on page 78 provides details on creating groups.

3. Create two sets of identical clients, and select the appropriate schedule and group for each client.

   "Task 6: Create a backup Client resource" on page 81 provides details on creating clients.

4. Select the first set of clients, one client at a time, and assign them to the Accounting Incremental group. Make the selections in the **Client Properties** dialog box. The required attributes are:

   - For the **Schedule** attribute, select Accounting Incremental
   - For the **Group** attribute, select Accounting Incremental

   After making these choices for each client, click **OK**.

5. Select the second (duplicate) set of clients, one client at a time and select **Accounting Full-Clone** for the **Schedule** and **Group** attributes.

6. After making these choices for each client, click **OK**.

7. To sort the clone and non-clone backups onto separate volumes:

   a. Open a Media Pool resource for non-clone backups and in the **Data Source** attribute, select the Accounting Incremental group.

   b. Open a Media Pool resource for clone backups and in the **Data Source** attribute, select the Accounting Full-Clone group.

      You can also assign backup data to media pools by using additional criteria. Chapter 7, "Sorting Backup Data" provides more information.

When the Accounting Full-Clone group runs its full backup on Sunday, the data is automatically cloned after the backup has

completed. The incremental backups (skipped every other day of the week) are *not* cloned.

Manual clone operations

This section covers operations that can be performed on both volumes and save sets that have been manually cloned. For information about manually cloning a volume or save set . "Volume cloning" on page 291 or "Save set cloning" on page 292 provide more information.

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.

Recovering cloned data

No special operation is required to recover cloned data. The volume, either a clone or original volume, that is selected for a recovery operation is determined as follows:

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 420 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 **nwrecover** (UNIX) or Sun StorageTek EBS User (Windows) programs. "Viewing the volumes required for data recovery" on page 435 provides more information.

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. "Restore a save set entry in the online indexes" on page 454 provides more information.

Recovering a cloned save set from the command prompt

Recovery of a cloned save set can be performed at the command prompt, for example:

```
recover -S ssid/cloneid
```

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

1. 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.
   ```

2. 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 notrecycable** 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

To clone an existing archive volume:

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 **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 292 provides more information.
5. 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 <ln Link>step 1 is ignored and the selection criteria behaves as though the volume is not mounted as described in <ln Link>step 2.
   - When cloning is used in a Virtual Tape Library (VTL) environment such as a CLARiiON Disk Library (CDL), the Sun StorageTek EBS 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 28 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 904 provides information about viewing log files.

Storage node selection criteria for writing the clone

The following criteria are used to determine the storage node to which the clone data will be written (write source):

1. The **Clone Storage Node** attribute of the read source storage node's Client resource is used as the write source.
2. If the read source host does not have a Client resource, the **Clone Storage Node** attribute of the Sun StorageTek EBS server is used as the write source. Otherwise, the **Storage Nodes** attribute of the Sun StorageTek EBS server is used as the write source.

No matter where cloned data is directed, the client file index and media database entries for the cloned save sets still reside on the Sun StorageTek EBS server. This ensures that browse and retention policies are handled consistently regardless of where the clone data was directed.

Directing a clone from one storage node to another storage node

To direct a clone from one storage node to another:

1. Open the **Client** resource for the read source storage node. "Editing a client" on page 650 provides more information.
2. Click the **Globals (2 of 2)** tab.
3. In the **Clone Storage Nodes** attribute, add the hostname of the storage node to which the clone data will be written. The first entry in the list that has a functional, enabled device will receive the cloned data from the read source storage node.

The **Clone Storage Nodes** attribute applies only to Sun StorageTek EBS servers and storage nodes. Leave this attribute blank for Sun StorageTek EBS clients that are not also Sun StorageTek EBS servers or storage nodes.

Directing clones from all storage nodes to a single storage node

To direct clones from all storage nodes to a single storage node:

1. On the **Client** resource for each storage node, leave the **Clone Storage Nodes** attribute blank.

1. Open the **Client** resource for the Sun StorageTek EBS server. "Editing a client" on page 650 provides more information.

2. Click the **Globals (2 of 2)** tab.

3. In the **Clone Storage Nodes** attribute, add the hostname of the storage node to which all cloned data is to be written.

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 <ln Link>step 1 is ignored and the selection criteria behaves as though the volume is not mounted as described in <ln Link>step 2.

   When cloning is used in a Virtual Tape Library (VTL) environment such as a CLARiiON Disk Library (CDL), the Sun StorageTek EBS 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.

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 deduplication clients

"Backup-to-tape for deduplication clients" on page 868 provides information on tape backups of deduplication client data.

Using the nsrclone command

For cloning, the **nsrclone** command is often used rather than automatic cloning. Using automatic cloning at the end of savegroup sessions can lead to concurrent device contention with other in-progress savegroup backups and preceding savegroup sessions, staging, and so on. Automatic cloning can thus be problematic and complicated.

However, there are limitations with the **nsrclone** command. Complex **mminfo-nsrclone** wrapper scripts are often used to extend functionality, such as enhanced filtering of candidate save set or clone objects on several selection / deselection criteria. Cloning scripts can be complex because they involve running mminfo reports, parsing the results, and so on.

In NetWorker 7.5, the need for **mminfo** scripting is greatly reduced. By using the enhanced **nsrclone** command, you now have 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. The selection criteria generally reflect the query semantics of the **mminfo** command. You can also use the **nsrclone** command to change the browse time of a save set, if the retention of a copy or clone of the save set is being changed.

The number of valid copies selection option facilitates retrying aborted nsrclone sessions. You can now simply retry the **nsrclone** command with **-C N** to only clone save sets that do not yet have a sufficient number *N* of copies in the desired target clone pool, where *N = 1* implies only clone save sets currently without any valid clones in the target clone pool. Note that existing save set copies or clones that are marked aborted, incomplete, suspect, or are past their retention period are not considered valid. Similarly, copies or clones for save sets on the read-only pseudo mirror of an Advanced File Type Device (AFTD) volume are not counted as physically valid. Complementing this retry feature is more extensive reporting at the end of a user-aborted **nsrclone** operation when run in extra verbose **-vvvv** mode.

New options **-C**, **-l**, and **-N** are supported by the enhanced **nsrclone** command, and existing options **-c** and **-g** have their functions extended.

Table 36 on page 310 provides the descriptions of the various options that can be used with the nsrclone command.

Table 36 List of nsrclone options and their descriptions

Options	Description
-C *less than copies in target pool*	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. 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 *level or range*	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.
-N *save set name*	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 *client name*	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 *group name*	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. This example assumes that no copies existed prior to that aborted 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 assuming that no copies existed prior to that aborted session, and with extended retention and browse periods:

```
nsrclone -S -e now -C 1 -y 12/12/2010 -w 12/12/2009
```

10 Staging Backups

This chapter covers these topics:

- Save set staging 314
- Working with staging policies..... 314

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 11 page 322 provides information on file type (FTD), and advanced file type device (AFTD) 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**.

   Note: 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.

   Note: The adv_file device and its corresponding _AF_readonly device will both be selected automatically, even if only one device was selected as the source of 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.

   Note: The Default volume can only be staged to the Default or Default Clone pool. Similarly, the Default Clone volume can only be staged to the Default or Default Clone pool and Archive data can only be staged to the Archive Clone pool. The other volume types can be staged to any pool. If

the Clone pool you have selected is restricted to storage node devices, you will also need to modify Clone Storage Node attribute of the Client resource for the Sun StorageTek EBS server to include the storage node name.

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).

   Note: 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.

    Note: 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.

    Note: 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.

    Note: 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**.

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.

   Note: 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

Note: 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.

Staging and cloning from the command prompt

Staging a save set from the command prompt works differently than staging a save set from the Sun StorageTek EBS Console. When staging from the Sun StorageTek EBS 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 Sun StorageTek EBS 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.

11

Device Operations

This chapter covers these topics:

- Overview of device operations 322
- File type and advanced file type devices 322
- Compatible block size for UNIX and Microsoft Windows 349
- Library device maintenance 351
- Device parameter settings 355
- Common device interface 359
- Device ordering 361
- Nonrewinding tape device usage (UNIX/Linux only) 364
- Block-size mode (UNIX/Linux only) 365
- Silos 365
- Miscellaneous device operations 372
- Tips for using libraries 376

Overview of device operations

Chapter 12, "Media Management" provides information on operations involving storage media.

You can view and access all libraries in the enterprise by selecting the Libraries option in the Console window. The Libraries option displays a hierarchy of Sun StorageTek EBS servers, configured Sun StorageTek EBS storage nodes, and devices that exist in the enterprise. The library information displayed is updated periodically or when a user requests it. Library information is also updated when Sun StorageTek EBS servers are added to, or deleted from, the enterprise hierarchy.

The details describing a particular device can be viewed, and often changed, in the device's Properties.

To see the full range of settings and options available in the Properties window, select **View** > **Diagnostic Mode**. For a description of the various fields and attributes on a given Properties tab, click **Field Help**.

As with other Console functions, users can view and work with only those Sun StorageTek EBS servers for which they have access permission.

File type and advanced file type devices

The Sun StorageTek EBS DiskBackup solution reduces the time it takes to save and recover data when compared to the use of tape. Sun StorageTek EBS DiskBackup allows data to be saved to a computer's local or network-attached disk, instead of using a tape device to write the data to storage media. The DiskBackup solution requires a DiskBackup option license; which are tiered based on disk size.

Sun StorageTek EBS software offers two variations of this solution:

- "File type device" on page 323 (available on Windows only)
- "Advanced File Type Device (AFTD)" on page 324

To prevent the file system from becoming full when backing up data to file type and AFTDs, a staging policy must be created to move the data off the disk as soon as possible. To make space for additional backups, perform one of these activities:

- Set up a staging policy.
- Review and, if required, modify the retention policy of the save sets.

Save set staging is helpful when used with either of these device types. "Save set staging" on page 314 provides more information.

File type device

Configure the Sun StorageTek EBS server to back up data to a file type device by creating a new Device resource in the same way as for any other storage device.

The following restrictions apply to the configuration of a file type device:

- Multiple file type devices can be supported on a system. Each device must, however, have a different name.
- Dynamic Drive Sharing is not supported.
- To use multiple file type devices on the same disk, partition the disk and create only one file type device per partition.
- For most device types, the Volume Default Capacity attribute is an arbitrary value that is an estimate of what the volume capacity is likely to be. It has no impact on how much data is actually written to a volume. If the value is not set correctly, the percent used calculation will not be correct. For file type devices, the Volume Default Capacity is a hard limit to the amount of data that can be written to the device.



IMPORTANT

In order to avoid accidentally filling a file system, set the Volume Default Capacity attribute to restrict the size of the file type device. For example, if a Volume Default Capacity of 100 MB is set for a file type device, then the device will be marked full when the capacity of 100 MB is reached. This is true for any type of file system.

If the Volume Default Capacity of a volume is changed, the changes will not take effect until the volume has been relabeled, the directory contents have been deleted, and the file device has been re-created.

In the Administration interface, the Volume Default Capacity appears in the Configuration tab of the Properties window of a device when Diagnostic Mode is on. To turn on Diagnostic Mode, select View>Diagnostic Mode. The value for this attribute must be less than 4 TB.

Note: The advanced file type device ignores the Volume Default Capacity value to allow dynamic expansion of disk space.

- Save sets from file type devices can be cloned to an advanced file type device, to take advantage of advanced file type device features.
- The following restrictions apply to file type devices created on a UNIX or Linux network file system (NFS):
  - The file system used for the file type device must not be used for any other data.
  - There must be one file type device per NFS system.
  - The Volume Default Capacity attribute for the file type device must be set to a size that is less than 100 percent of the total capacity of the file system.
- 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.
- When a file type device is used as a backup solution, the upper limits of save set size depend on either of these:
  - The upper limits supported by the operating system
  - The file size specified by the disk device vendor

"Differences between file type devices and AFTDs" on page 337 describes the functional differences between the file type device and the advanced file type device.

Advanced File Type Device (AFTD)

The advanced file type (adv_file) is designed for large disk devices that use a volume manager to dynamically extend available disk space if the disk runs out of space while being accessed as an AFTD during backup. The *EMC Software Compatibility Guide* provides a list of supported volume managers. "Differences between file type devices and AFTDs" on page 337 provides more information.

When an AFTD runs out of disk space, the current backup is suspended. This message appears:

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

Immediately after this message appears, the action that is associated with the Filesystem Full — Recover adv_file Space notification takes place. By default, the action for this notification uses the **nsrim** command to delete expired save sets. If enough space has been cleared, the backup continues. If the recycle setting for the volume is manual, however, 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 has not been cleared after a 10-minute period from the start of the initial deletion of the expired save sets, the action that is associated with the Filesystem Full—Waiting for adv_file Space notification takes place. By default, the action for this notification sends an email to the root user on the Sun StorageTek EBS server on UNIX and Linux systems, and logs a message in the media log file, located in <Sun StorageTek EBS_install_path>\logs file on Windows systems. "Viewing log files" on page 904 provides information about viewing log files.

At this point, the backup ceases until more space has been made available. The notifications listed are default notifications that pertain specifically to AFTDs. Examples of two customized adv_file [full event] notifications based on the two default adv_file notifications follow. The custom notifications allow users to run their own scripts or invoke other programs (for example, to extend or increase the disk capacity) so that a backup can continue. These custom notifications would be used in addition to, rather than as replacements for, the default notifications.

Note: Managing large disk devices can involve complex workflow management once various events are triggered. This workflow can be handled by using notifications. The Sun StorageTek EBS software provides the two preconfigured notifications described previously. Users can also customize their own notifications, they can change and expand on the way the Sun StorageTek EBS software behaves when an AFTD full notification occurs, or even invoke custom scripts and other programs that can be used to expand the capacity of existing AFTDs, among other things. "Notifications" on page 619 and "Configuring NetWorker SNMP notifications" on page 808 provides information about using notifications.

Example 29 Create a custom notification to extend disk space

While the Sun StorageTek EBS 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 diskspace, for example: /mybin/my_first_custom_script.
7. Click **OK**.

Example 30 Create a Custom Notification for Insufficient Disk Space

The Sun StorageTek EBS 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 environments

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.

Note: Dynamic Drive Sharing is not supported with the AFTD.

As part of the validation process, conduct these tests:

- 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

Note: Some versions of Network File System (NFS) or Common Internet File System (CIFS) drop data s when a file system becomes full. Be sure to use versions of NFS, CIFS, and operating systems that fully interoperate and handle a full file system in a robust manner.

On some disk devices, the volume labeling process may take longer than expected. This extended labeling time depends on the type of disk device being used and does not indicate a limitation of the Sun StorageTek EBS 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.

Users should not edit device files and directories. This action is not supported, and such editing can cause unpredictable behavior and make it impossible to recover data.

AFTD allowed directories configuration

A new AFTD Allowed Directories attribute provides control over the directories where an AFTD may be created on a storage node. Only a system administrator with security privileges can set up an allowed

directory list for AFTD creation. If there is a conflict due to change of allowed directories, the change of allowed directories will be rejected.

Before creating an AFTD, set up a AFTD Allowed Directories to control access to storage node file systems. If this attribute is populated it will restrict the creation of an AFTD so that the device path must be under the list of directories. If this attribute is left empty, then there will be no restrictions on the path of an AFTD when it is created, with the exception that certain well known system level directories are always rejected.

When creating a new AFTD by using the Sun StorageTek EBS Console, a verification dialog box will appear after the user commits the new AFTD device. The verification will ask whether to have the Sun StorageTek EBS software verify the device path that was entered to ensure that it exists on the file system of the storage node where the AFTD will be created. If you accept the verification, the Sun StorageTek EBS software ensures that the path exists. If it does not exist, you will be notified and the AFTD will not be created. You then can type a different path or quit the creation of the AFTD.

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. "Differences in the cloning process" on page 339 provides more information.

Figure 28 on page 329 shows a system configuration where both Sun StorageTek EBS client A and Sun StorageTek EBS client B can use NFS to recover files concurrently from AFTD1.

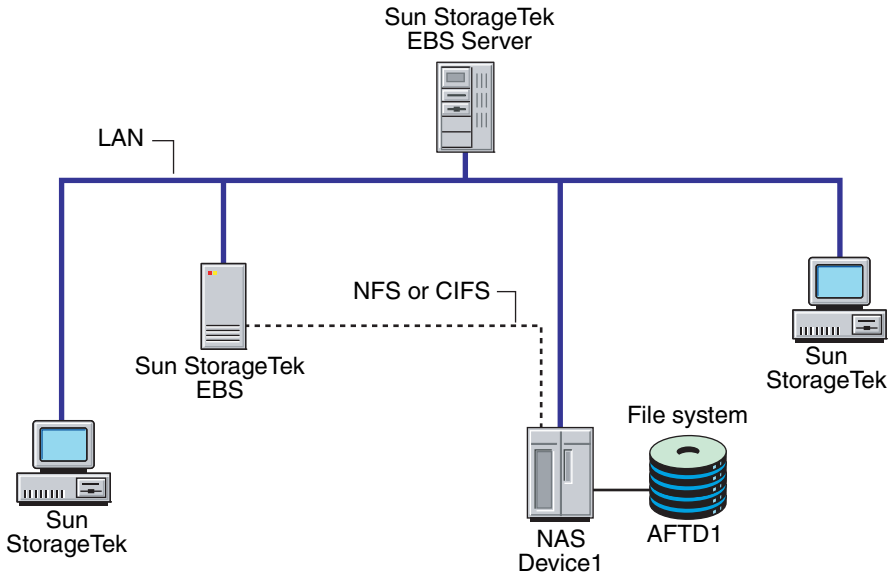


Figure 28 **Configuration showing an AFTD**

The creation of an AFTD yields two devices. When the primary (read-write) advanced file type device is labeled, the Sun StorageTek EBS software creates a secondary device with read-only accessibility. The Sun StorageTek EBS software then creates a volume with an .RO suffix in its name, and then automounts it in this device.

Note: As indicated by the name, the secondary device is read-only. This enables concurrent operations, such as reading from the read-only device. Do not make any changes to the _AF_readonly Device resource. If, however, changes are made to the read-write volume or save set, these changes can be propagated to the read-only volume or save set.

In Table 37, when the AFTD is labeled, the _AF_readonly device is created during the labeling process. The Sun StorageTek EBS software

then labels and mounts a volume with an .RO suffix (for Read Only) on this device.

Table 37 AFTD names

Device	AFTD name	Type	Volume
UNIX and Linux format	/mercury/dev1	adv_file	af_1
	/mercury/dev1/_AF_readonly	adv_file	af_1.RO
Windows format	d:\file_devices\dev1	adv_file	af_1
	d:\file_devices\dev_AF_readonly	adv_file	af_1.RO
Common Internet File System (CIFS)	(rd=storage_node:)\\<CIFS_host>\sharepoint\dev1	adv_file	af_2
	(rd=storage_node:)\\<CIFS_host>\sharepoint\dev1_AF_readonly	adv_file	af_2.RO

Table 37 also illustrates that when backing up data to a device, it is possible to read concurrently from the device from the device labeled with the _AF_readonly suffix.

Note: The terms _AF_readonly name and the .RO suffix are reserved for AFTDs. Do not use these terms when naming devices or labeling volumes.

Table 38 on page 330 describes which changes must be propagated from the read-write device, volume, or save set to the read-only device, volume, or save set. For example, when creating and labeling the read-write AFTDs, this information would apply.

Table 38 also shows which Sun StorageTek EBS operations can be performed concurrently within a single storage node.

Table 38 Concurrent operations

Backups	Recoveries (from the recover command line program only)	Manual cloning	Automatic or manual staging
X	X		
X		X (1 clone)	
X			X (1 stage)

This table also applies to SAN-attached storage nodes when sequentially sharing a single AFTD.

Allow one operation to complete before starting another operation that is not supported to run concurrently. Only these operations work concurrently with the AFTD option:

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

There may be a need 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.

Limitations with concurrent AFTD recovery operation

AFTD concurrent recovery currently has these limitations:

- Not available to the Windows recover interface (winworkr). Use the **recover** command. See the *NetWorker Command Reference Guide* or the **recover** man page for more information.
- Not available to nonfile recoveries, such as NDMP and Sun StorageTek EBS database modules.
- Sun StorageTek EBS release 7.2*x* clients on Windows recover data from AFTD storage nodes sequentially.
- Concurrent recoveries must be performed from the command line using the **recover** command, either by using multiple **-S** options to identify multiple save sets, or executing multiple **recover** commands concurrently.

Sharing an AFTD between storage nodes in a NAS or SAN environment

In a Network Attached Storage (NAS) Environment, Sun StorageTek EBS operations can be performed concurrently on two storage nodes that share an AFTD. When sharing an AFTD, one storage node can save to a writable volume, while the other storage node either recovers or clones from a read-only volume.

In a Storage Area Network (SAN), Sun StorageTek EBS operations are performed sequentially when two storage nodes share an AFTD. Only one storage node at a time can use the shared AFTD.

Any staging occurs from the storage node that has the read-only volume mounted.

Certain configuration changes must be made to share an AFTD between storage nodes. "Configuring a shared AFTD in a NAS environment" on page 333 and "Configuring a shared AFTD in a SAN environment" on page 334 provide more information.

Once the configuration has been done, however, no change is required in how to start these operations:

- Backup
- Recovery
- Cloning
- Staging

Before configuring shared AFTDs

Review the following preparations before sharing AFTDs:

- Ensure that OS permissions (directory and file) and sharing are set up properly between storage node hosts for the root user or Windows administrator, to enable proper sharing of AFTDs on the file system.

  Use OS commands to create, copy, or erase directories or files on the sharing disk (such as NAS, JBOD in SAN), to ensure that sharing is possible at file system level between machines. If the operating system does not permit such sharing, then the Sun StorageTek EBS software cannot change that.

- To share an AFTD between storage nodes, ensure that:
  - The proper Storage Node and Clone Storage Node attributes of the Client resource have been set.
  - The proper Device attribute (the one with the read-only volume mounted) has been specified for the Staging resource.
- To share an AFTD in NAS on Windows storage nodes, ensure that:
  - The **nsrexecd** (Sun StorageTek EBS Remote Exec) service has been started by a Windows Administrator account on the storage node.
  - The CIFS-mapped AFTD was created with UNC pathnames that have the appropriate remote user and password specified. An example of UNC path syntax is:

rd=sn_a:\\nas1\path\shared_aftd

"CIFS-mounted AFTD on Microsoft Windows servers and storage nodes" on page 345 provide more information.

Configuring a shared AFTD in a NAS environment

When sharing an AFTD between two storage nodes in a NAS environment, these conditions apply:

- For UNIX or Linux platforms, NFS mount the NAS file system for AFTD on the sharing storage nodes.
- For Windows platforms, CIFS mapping to the NAS file system is optional on the storage nodes.

Perform the following steps to configure a shared AFTD in a NAS environment:

1. To create and label the AFTD on storage node A (*sn_a*):
   - On UNIX or Linux platforms: rd=sn_a:/shared_aftd
   - On Windows platforms: rd=sn_a:\\nas1\path\shared_aftd

   The remote user and password must be specified if this requirement has been set up in the NAS file system.
2. To create the same AFTD on storage node B (*sn_b*):
   - On UNIX or Linux platforms: rd=sn_b:/shared_aftd
   - On Windows platforms: rd=sn_b:\\nas1\path\shared_aftd
3. In the Sun StorageTek EBS software, mount the AFTD (but do not label it) on storage node B, to create the read-only volume.
4. In the Sun StorageTek EBS software, unmount all volumes of the shared AFTD on storage nodes A and B.
5. To write (back up) to storage node A in the Sun StorageTek EBS software:
   a. In the Sun StorageTek EBS software, unmount the read-write volume from storage node B (if it is mounted).
   b. Mount the read-write volume on storage node A.
   c. Configure the **Storage Nodes** attribute of the Client resource (*sn_a*).
   d. Ensure that the Sun StorageTek EBS server's local storage node has a volume available for a bootstrap backup.
6. To read (recover, clone, or stage) from storage node B:

a. In the Sun StorageTek EBS software, unmount the read-only volume from storage node A (if it is mounted).

b. Mount the read-only volume on storage node B.

c. Configure the Recover **Storage Nodes** attribute of the **Client** resource (sn_b).

d. Configure the **Clone Storage Nodes** attribute in the **Client** resource with the target of the clone (or stage) operation.

e. Configure the staging policy's **Devices** attribute to select rd=sn_b:/shared_aftd for autostaging from sn_b.

"Concurrent operations" on page 330 provides more information.

Configuring a shared AFTD in a SAN environment

When sharing an AFTD between two storage nodes in a SAN environment, use the appropriate operating system command or SAN command to unmount (disable) the sharing disk from any machine, and mount (enable) the same disk on storage node A.

Perform the following steps to configure a shared AFTD in a SAN environment:

1. Create, label, and mount the sharing AFTD on storage node A.
2. Perform the **save**, **recover**, or **clone to** operation from storage node A (with the Storage Nodes attribute of the Client resource configured as storage node A).
3. Perform the **stage** operation from storage node A. Choose the shared AFTD from storage node A in the Device attribute of the Stage resource.
4. In the Sun StorageTek EBS software, unmount the AFTD volumes from the sharing storage node A.
5. Use the appropriate operating system command or SAN command to unmount (disable) the sharing disk from storage node A, and mount (enable) the same disk on storage node B.
6. Create the AFTD to share in SAN for storage node B.
7. In the Sun StorageTek EBS software, mount the AFTD (but do not label it) on storage node B.
8. Perform the **save**, **recover**, or **clone to** operations from storage node B (with the Storage Nodes attribute of the Client resource configured as storage node B).

9. Perform the **stage** operation from storage node B. Choose the shared AFTD from storage node B in the Device attribute of the Stage resource.

Issues and limitations with sharing AFTDs

Shared AFTDs have these limitations:

- Read-only volumes might be automounted onto the writing storage node during saves. The workaround is to unmount all sharing instances of read-only and read-write AFTD volumes from all storage nodes, and then correctly remount them.
- Limitations on sharing AFTD between storage nodes in NAS/SAN:
  - Supports only homogeneous storage node platforms when sharing AFTD (such as Windows with Windows storage nodes, or UNIX with UNIX storage nodes).
  - The sharing read-write or read-only AFTD volume must be manually unmounted from one storage node before being mounted on another, in order to prevent a potential out-of-sync state for the volume.
  - All instances of the sharing read-write or read-only AFTD volume must be manually unmounted from all of the sharing storage nodes before relabeling the sharing AFTD, in order to prevent potential data loss.
  - For SAN only:
    - Only one storage node at a time can perform Sun StorageTek EBS operations.
    - Operating system or SAN commands must be issued to mount or enable the sharing disk on the second storage node after unmounting or disabling from the first storage node. This ensures that file and directory sharing of the same disk is supported (set up) in the SAN from the sharing storage nodes, which allows sequential sharing of the disk as an AFTD in the Sun StorageTek EBS software.

Changing block size for AFTDs

You can adjust the maximum potential block size for backups to an AFTD device. 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 via 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.



IMPORTANT

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

To set the maximum potential block size for a particular 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 or right-click the device and select **Properties**.
5. Select the **Advanced** tab. In the **Device block size** attribute, select a value from 128 to 256.
6. Click **OK**.

   Relabel the AFTD device to make the new setting take effect.

Device and volume display

When the AFTD is implemented, certain display and volume changes occur.

Device display

Both the read-write and read-only devices with the *_AF_readonly* suffix are displayed:

- Volumes corresponding to both the read-write and read-only devices appear.

- Volume corresponding to the *_AF_readonly* device appears with the *.RO* suffix.

Volume display

With the AFTD, read-only volumes (including cloned volumes) appear with the *.RO* suffix.

Propagating changes to the Enable attribute from read-write to read-only

When an AFTD is labeled, an *_AF_readonly* device is automatically created and added to the Sun StorageTek EBS device list. When data is backed up, additional entries are added to the Sun StorageTek EBS media index for the data backed up to that device.

When a change is made to the Enabled attribute for the user-created, AFTD read-write device, the same change must immediately be made to the *_AF_readonly* device.

To propagate the changes:

1. Access the read-write device and change the **Enabled** attribute.
2. Access the read-only device and change the **Enabled** attribute to match the change made to the read-write device.

Note: Always make attribute changes to the read-write device, and then make these changes to the read-only device.

Differences between file type devices and AFTDs

Table 39 lists the functional differences between the file type device and the AFTD.



IMPORTANT

Data loss will result if a full FTD is made appendable while there is a backup pending completion with a save set partially written to the full FTD. In this case, the partial save set (currently in "incomplete" state) will be overwritten.

Table 39 Advanced file type and file type device differences (1 of 2)

Function or operation	AFTD	File type device
Create a device	Select media type: adv_file UNIX/Linux: local or NFS Windows storage node: • Local or Common Internet File System (CIFS) Universal Naming Convention (UNC) path for CIFS: • Remote user • Password	Select media type: file UNIX/Linux: local or NFS. The Important Note on "AFTD environments" on page 327 provides more information. Windows storage node: local path only
Label a device	• Volume name must not exceed 60 characters. • Automatically creates an _AF_readonly device, then labels and automounts (enables) a volume with a *.RO* suffix.	• Volume name must not exceed 63 characters.
Concurrent operations: "AFTD concurrent operations" on page 328 provides more information.	Yes	No
Save set ID/clone ID: "AFTD concurrent operations" on page 328 provides more information.	• One save set ID with two clone IDs created for each save set. • A second (virtual) clone created on *_AF_readonly* device. • User can specify virtual clone ID for concurrent operations.	One save set ID created with one clone ID for each save set.
Reclaiming or recovering space	• 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.	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).

Table 39 Advanced file type and file type device differences (2 of 2)

Function or operation	AFTD	File type device
Volume default capacity for devices	Does not apply.	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.
When file system or volume is full	• 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.	• 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.
Save set continuation	No. Save sets that start on an AFTD must be completed on the same device.	Yes
Save set consolidation	Not supported.	Supported
Data format in device	Save stream (**uasm**) format (uses less space).	EMC Open Tape Format (OTF).

Differences in the cloning process

For both file type devices and AFTDs, save sets are cloned one at a time. Important differences exist, however, in the cloning process for the two types of devices:

- For file type devices, both automatic and manual cloning begins only after all the save sets in a savegroup have been backed up.
- For AFTDs, automatic cloning begins after the save sets have been backed up. Manual cloning of a save set can, however, begin as soon as it has finished its backup.
- For AFTDs, two simultaneous **clone** operations can be run from the command-prompt, as long as no **backup**, **recover**, or **stage** operations are running at the same time. The syntax is as follows:
  - **nsrclone -S** *[ssid/cloneid1]*
  - **nsrclone -S** *[ssid/cloneid2]*

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

Support for Sun StorageTek EBS DiskBackup configurations

Sun StorageTek EBS DiskBackup supports these configurations:

- A local disk on a Sun StorageTek EBS storage node is supported for both file type devices and AFTDs.
- A network-attached disk device that is NFS-mounted to a UNIX version of the Sun StorageTek EBS storage node. This is supported for both file type devices and AFTDs.
- A network-attached disk device that is CIFS-mounted to a Sun StorageTek EBS storage node running on Microsoft Windows. This is supported only on AFTDs.

Supported disk vendors

The Sun StorageTek EBS DiskBackup solution supports these disk vendors:

- Those NFS implementations that must behave as UNIX or Linux NFS.
- Those CIFS implementations that must behave as Windows CIFS.

Example 31 Sun StorageTek EBS DiskBackup within a system configuration

Figure 29 on page 341 illustrates how to use Sun StorageTek EBS DiskBackup within a system configuration. In this example:

- The AFTDs are created on Disk Device 1, Disk Device 2, and the Local Disk is created on storage node A.
- A file type device is created on Local Disk 2.
- Linux/UNIX storage node A writes its backups to one of the following:
  - The AFTD through an NFS connection to Disk Device 1.
  - The AFTD on Local Disk 1.
- The Microsoft Windows Storage Node B uses a CIFS connection to back up to the AFTD on Disk Device 2. The file type device cannot be used on a CIFS connection.
- The Microsoft Windows Storage Node C writes its backups to the file type device on the Local Disk 2.

Sun StorageTek EBS DiskBackup can be tailored to the specific configuration of the system.

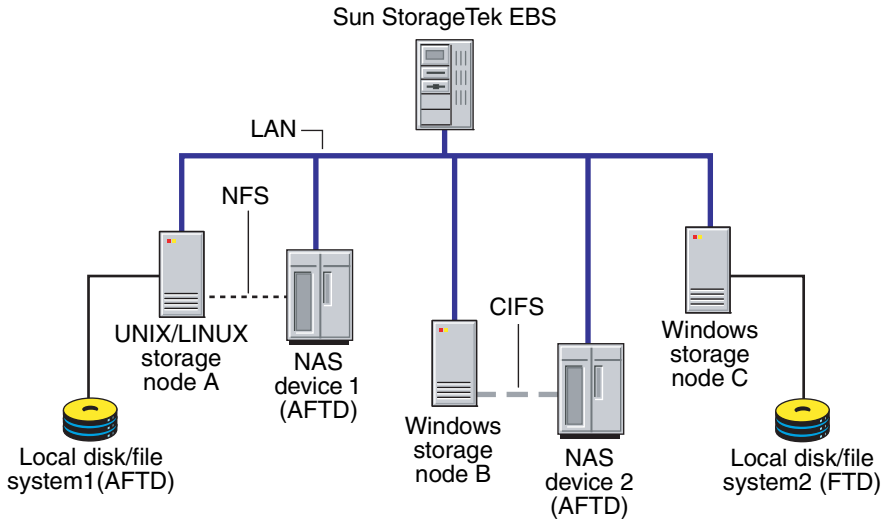


Figure 29 Sample Sun StorageTek EBS DiskBackup configuration

Sun StorageTek EBS DiskBackup configuration

This section explains how to configure a Sun StorageTek EBS DiskBackup Device, and includes:

- "File type device or AFTD on UNIX and Linux servers and remote storage nodes" on page 342 for UNIX or Linux servers and remote storage nodes with local or NFS-mounted disks.
- "File type device or AFTD on local disks of Microsoft Windows servers and storage nodes" on page 344 for Microsoft Windows servers and remote storage nodes with local disks.
- "CIFS-mounted AFTD on Microsoft Windows servers and storage nodes" on page 345 for Microsoft Windows servers and remote storage nodes with CIFS-mounted disks (AFTD only).

"File type and advanced file type devices" on page 322 provides specific instructions and restrictions on backing up to a file type device disk.

File type device or AFTD on UNIX and Linux servers and remote storage nodes

This section covers file type or advanced file type devices on UNIX and Linux systems. Configure Sun StorageTek EBS software to back up data to a file type device or AFTD by creating a new Device resource for the Sun StorageTek EBS storage node just as for any other storage device.

Note: Do not modify directories or files created by file type devices or AFTDs. Changes to the device path or its contents prevent the Sun StorageTek EBS software from locating the device.

Configure a file type device or AFTD on UNIX and Linux servers and storage nodes

To configure a file type device or AFTD on UNIX and Linux:

1. Create one directory for each disk (or partition) to be used for an AFTD or file type device. AFTD and file type devices both require creation of a directory in the disk file system that the Sun StorageTek EBS server or storage node recognizes as the device name (and the destination for the data).

   Note: Do not use the tmp directory for Sun StorageTek EBS DiskBackup devices because the data could be overwritten.

2. In the server's **Administration** interface, click **Devices**.
3. Right-click **Devices** in the navigation tree, and select **New**.

   The **Create Device** window appears, with the **General** tab selected, and a default device name in the **Name** field of the **Identity** area of the window.

4. Replace the default name with the complete path of the directory created. For example:

   `/`*directory_name*
   `rd=`*storage node*`:/`*directory_name*

   For an NFS-mounted file type device or AFTD, ensure the NFS mount point is available by either automounting or manually mounting it as required:

   a. If the device is configured on the server's storage node (file type or AFTD), the name is the simple device path, such as /backup/d0.

b. If the device is configured on a remote storage node, then the name must indicate that the storage node is remote, by including **rd=** and the name of the remote storage node in the device path. For example, if the remote storage node is neptune, then the device path might be rd=neptune:/backup/d0.

5. In the **Identity** area:

   a. Add a comment (optional) in the **Comment** field.

   b. In the **Media Type** field, select a type from the menu:

      – file, for a file type device.
      – adv_file, for an AFTD.

6. In the **Status** area, make sure that the **Auto Media Management** feature for AFTD or file type device 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 to set these attributes as required:

   - Target Sessions
   - Max Sessions
   - **Local Backup to a Dedicated Storage Node** (maintain the default No setting)

9. Click **OK**, when the configuration is complete.

10. Label and mount theSun StorageTek EBS DiskBackup Device in the same manner as other offline media used for Sun StorageTek EBS backups.

    - Since the AFTD is never marked full, there is no need to add another device to the pool.
    - Because the AFTD automatically affixes the *.RO* suffix to a volume name, limit the volume name to 60 characters when labeling the volume.

Note: Save set staging is especially useful when using either device type of the Sun StorageTek EBS DiskBackup feature.

Configuring AFTD on sub-directory of mount point if NAS device mounted on local directory

If a NAS device has been mounted on a UNIX mount point on the local directory, configure the AFTD to use a sub-directory of the mount point. Otherwise, if a NAS device becomes unmounted during backup and the AFTD is configured to use that mount point as the device path, the backup will not fail and NetWorker will continue writing to the local file system.

For example, if you mounted the NAS device on a local **/mnt** directory and configured **/mnt** as the AFTD, and the NAS device unmounts during backup, NetWorker continues to write to the local **/mnt** folder.

To ensure this does not occur, create a sub-directory on a mounted NAS file system and configure the AFTD on a sub-directory of the mount point (using the above example, configure AFTD as **/mnt/directory1** instead of **/mnt**). If an unmount occurs, the backup will abort because this directory is not part of the local file system.

File type device or AFTD on local disks of Microsoft Windows servers and storage nodes

Configure the Sun StorageTek EBS software to back up data to a file type device or AFTD by creating a new Device resource for the Sun StorageTek EBS server just as for any other storage device.

Note: Do not use the temp directory for Sun StorageTek EBS DiskBackup devices. The data could be overwritten. Because AFTDs are never marked as full, there is no need to configure another device in the same pool.

Configure a file type device or AFTD on local disks of Microsoft Windows servers and storage nodes

To configure a file type device or AFTD:

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

   Each disk (or partition) requires a directory that is used for file type device or AFTD storage. Each of these devices requires a directory in the disk file system that the Sun StorageTek EBS server recognizes as the device name (and the destination for the data).

2. In the server's **Administration** interface, click **Devices**.

3. Right-click **Devices** in the navigation tree, and select **New**. The **Create Device** window appears, with the General tab selected, and a default device name in the **Name** field of the **Identity** area of the window.
4. Replace the default name with the complete path of the directory created. For example:

   ```
   H:\device01
   rd=mercury:H:\device01
   ```

5. In the **Identity** area:
   a. Add a comment (optional) in the **Comment** field.
   b. In the **Media Type** field, select a type from the menu:
      – For a file type device, select **File**.
      – For an AFTD, select **Adv_file**.
6. In the **Status** area, make sure that the **Auto Media Management** feature for AFTD or file type device 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 to set these attributes as required:
   - Target Sessions
   - Max Sessions
   - **Local Backup to a Dedicated Storage Node (**maintain the default No setting)
9. Click **OK**, when the configuration is complete. Label and mount the Sun StorageTek EBS DiskBackup Device in the same manner as other offline media used for Sun StorageTek EBS backups.

Note: Because the AFTD automatically affixes the *.RO* suffix to a volume name, limit the volume name to 60 characters when labeling the volume. Save set staging is effective when used with either device type of the Sun StorageTek EBS DiskBackup feature. "Save set staging" on page 314 provides more information.

CIFS-mounted AFTD on Microsoft Windows servers and storage nodes

This section explains how to start the Sun StorageTek EBS service and configure the AFTD on a CIFS connection on a Microsoft Windows system. Note that the Sun StorageTek EBS service must be started by

a member of the Administrator group before the AFTD can be configured.

Start the Sun StorageTek EBS service on Microsoft Windows

To start the Sun StorageTek EBS services as a user with administrator rights:

1. Create an Administrators account. For example: NWservice.
2. Ensure that the user account created for the NetWorker **Remote Exec** service has Full control on the share where the AFTD is used.
3. Ensure that **User must change password at next logon** is cleared.
4. Ensure that **User cannot change password** is selected.
5. Ensure that **Password never expires** is selected.
6. Do not set the **Administrator** account to an account that has **login as a service** rights.
7. From the Windows Start menu, select **Settings > Control Panel > Administrative Tools > Services**.
8. Right-click **Services**.
9. Stop the appropriate service:
   - For a NetWorker server, stop the NetWorker **Backup and Recover** service and NetWorker **Remote Exec** service.
   - For a NetWorker storage node, stop the NetWorker **Remote Exec** service.
10. Right-click the NetWorker **Remote Exec** service that was stopped in step 9 and select **Properties**.
11. Select the **Log On** tab.
12. Select **This Account**.
13. Replace **Local System** with the Administrative user created (NWservice).
14. Type the administrator password in the appropriate field.
15. Click **OK**.

When a service is given a user account, it is automatically given login as a service rights. When the **NetWorker Backup and Recover** service is started, it will automatically start the NetWorker **Remote Exec** service.

Configure a CIFS-mapped AFTD on Windows server and storage node

To configure a CIFS-mapped AFTD:

1. Create one directory for each disk (or partition) to be used for AFTD. AFTD and file type devices both require creation of a directory in the disk file system that the Sun StorageTek EBS server recognizes as the device name (and the destination for the data).
2. In the server's **Administration** interface, click **Devices**.
3. Right-click **Devices** in the navigation tree, and select **New**. The **Create Device** window appears, with the **General** tab selected, and a default device name in the **Name** field of the **Identity** area of the window.
4. Replace the default name with the complete path of the directory created by using the Universal Naming Convention (UNC):
   - `\\`*CIFS_host*`\`*share-point-name*`\`*path*
   - **`rd=`***sn_host*`:\\`*CIFS_host*`\`*share-point-name*`\`*path*

   Note: To configure a Microsoft Windows server, do not include rd=<sn_host>: in the UNC.
5. For **Media Type**, select adv_file.
6. In the **Status** area, make sure that the **Auto Media Management** feature for AFTD or file type device 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 to set these attributes as required:
   - Target Sessions
   - Max Sessions
   - **Local Backup to a Dedicated Storage Node** (maintain the default No setting)
9. In the **Remote Host** area:
   - Leave NDMP unselected.
   - These are required if the CIFS file system has been set up to require password access:

a. Type the CIFS file system username in the **Remote User** field. For example, in the configuration illustrated in Figure 30 on page 348, type the administrator of the CIFS file system.

b. Type the CIFS password in the **Password** field. For example, in the configuration illustrated in Figure 30 on page 348, type the password for the CIFS file system.

10. Click **OK**, when the configuration is complete.

Note: If a new password for the CIFS AFTD has been provided, the Sun StorageTek EBS services must be restarted on both the server and storage node in order for the new password to take effect, and to reestablish the CIFS connection correctly.

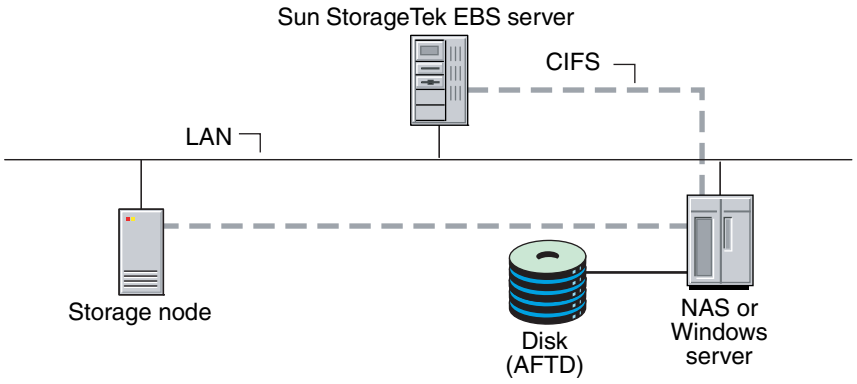


Figure 30 CIFS-mounted advanced file type device

Deleting file type devices and AFTDs

"Devices and libraries" on page 124 provides information about deleting file type and advanced file type devices.

Sun StorageTek EBS DiskBackup licensing

Chapter 28, "Licensing" provides information on licensing for the Sun StorageTek EBS DiskBackup option

Compatible block size for UNIX and Microsoft Windows

Different SCSI hardware limitations exist between UNIX and Microsoft Windows operating systems. This can lead to 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.

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 **Sun StorageTek EBS 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**.

Solutions to block-size compatibility problems

To solve problems with block-size compatibility for a given device:

1. In the server's **Sun StorageTek EBS Administration interface**, click **Devices**.
2. Select **View** > **Diagnostic Mode** on the menu bar.
3. Open the **Libraries** folder in the navigation tree and select the drive that contains the mounted volume with the block size being checked. The drive's detail table appears.

4. Right-click the drive in the detail table, and select **Properties**. The **Properties** window appears.
5. Select the **Advanced** tab. In the **Device Configuration** area, the currently configured **Device Block Size** value is displayed.
6. Select the appropriate **Device Block Size** value.
7. Click **OK**.

Note: It is also possible to solve problems with block-size compatibility by changing the block size for an entire device type. The change, however, must be made on each storage node where it is to be available. Once the block size is changed, it affects only those volumes that are labeled after the change. Volumes can be relabeled to use the new block size, but if they contain data that should be saved, be sure to clone the data beforehand to a volume that already uses the new block size.

Setting the block size for a device type

To set the block size for an entire device type:

- On a UNIX system, 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 Sun StorageTek EBS 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 Sun StorageTek EBS 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 Sun StorageTek EBS server in order for changed environment variables to take effect.

Library device maintenance

This section describes how to maintain storage devices.

Periodic maintenance

Clean a storage device periodically to keep it working correctly. The Sun StorageTek EBS server provides automatic cleaning of drives 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 Sun StorageTek EBS 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.

How to manually select a tape device for cleaning

To set the cleaning attributes of a library:

1. In the server's **Sun StorageTek EBS 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**.

Note: Do not enable automated cleaning for silos in the Sun StorageTek EBS 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.

How to delay tape device cleaning

It is sometimes 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 **Sun StorageTek EBS 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 359 provides more information about CDI.

Devices that are capable of TapeAlert perform constant self-diagnostics and communicate the diagnostic information via the **nsrmmd** program to logs that can be viewed in the Monitoring task.

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 40 describes the nature of the tape alert levels.

Table 40 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 Sun StorageTek EBS drive cleaning.

Warning messages

Warning messages indicate the following types of drive errors:

- Recoverable read or write errors occurred.
- Media is at end of life.
- Read-only tape format is in the drive.
- Periodic cleaning is required.

Critical messages

Critical messages are warnings that a drive might be disabled and requires immediate attention to avoid data loss:

- Unrecoverable read or write errors occurred.
- Tape is marked read-only.
- Drive require immediate cleaning.
- Drive is predicting hardware failure.

Informative and warning messages should clear automatically by **nsrmmd** once the reported issue is handled.

Critical messages about hardware errors are not cleared by **nsrmmd** because they might indicate intermittent hardware problems.

Service mode

Use the service mode setting to take a device offline temporarily. Service mode differs from the disabled state in that the **nsrmmd** 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 has been taken out of the collection of drives that the Sun StorageTek EBS 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 410 provides more information about how to set the maximum consecutive error count.

Note: The drive must be manually reset to Enabled for the Sun StorageTek EBS 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**.

Device parameter settings

Since Sun StorageTek EBS release 7.3, device parameters can be set in the Sun StorageTek EBS Administration interface for individual devices on a server. In earlier releases, environment variables were used to set device parameters for all devices on a given Sun StorageTek EBS server.

Environment variables can still be used to adjust several Sun StorageTek EBS device parameters, but the adjustment of environment variables should only be done by users who know the server environment and performance tuning (for example, an administrator who wants to fine-tune performance by changing a certain setting for all LTO devices on a particular Sun StorageTek EBS server).

The variables (and their equivalent names in the Administration interface) are described in these sections:

- "Device settings in the Sun StorageTek EBS Administration interface" on page 355
- "Device settings as environment variables" on page 356

Device settings in the Sun StorageTek EBS 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.
5. Select the **Advanced** tab. In the **Device Configuration** area, the device settings are the first fields shown. The following tables lists the fields and their corresponding environment variables:

Table 41 Device settings and environment variables

Device setting	Corresponding environment variable
Device Block Size	NSR_DEV_BLOCK_SIZE_*MEDIA_TYPE*
Device File Size	NSR_DEV_TAPE_FILE_SIZE_*MEDIA_TYPE*
Device Load Time	NSR_DEV_LOAD_TIME_*MEDIA_TYPE*
Device Eject Time	None
Device Poll Interval	NSR_DEV_LOAD_POLL_INTERVAL_*MEDIA_TYPE*
Device Min Load Tries	NSR_DEV_LOAD_TRY_LIMIT_*MEDIA_TYPE*
Device Default Capacity	NSR_DEV_DEFAULT_CAPACITY_*MEDIA_TYPE*
Device Tape Flags	None

When device parameters are set in this interface, it is not necessary to stop and restart the Sun StorageTek EBS 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 Sun StorageTek EBS server.

Note: The media type syntax must be all uppercase, with underscores (_) replacing blank spaces and hyphens. For example, a device displayed in the Sun StorageTek EBS software as "8mm Mammoth-2" would be listed as: `8MM_MAMMOTH_2`
To determine the media type, right-click the device an select the **General** tab. The **Media Type** attribute contains the media type that should be used in these environment variables.

Microsoft Windows users use the Control Panel's System applet to set these system environment variables. The Sun StorageTek EBS Backup and Recover Server service must be restarted in order for the environment variables to take effect.

UNIX/Linux users set these environment variables in the Sun StorageTek EBS startup script before this line:

echo 'starting Sun StorageTek EBS daemons:') > /dev/console

Type the variables in this format:

ENV_VAR_NAME = value

export *ENV_VAR_NAME*

The Sun StorageTek EBS server must be restarted in order for the environment variables to take effect.

Table 42 lists the location of Sun StorageTek EBS startup scripts on UNIX and Linux systems.

Table 42 Sun StorageTek EBS startup script location on UNIX/Linux systems

Operating system	Startup script location
Solaris, Linux, and IRIX	/etc/init.d/networker
HP-UX, HP Tru64	/sbin/init.d/networker
AIX	/etc/rc.nsr

NSR_DEV_BLOCK_SIZE_MEDIA_TYPE

NSR_DEV_BLOCK_SIZE_*MEDIA_TYPE* is organized in units of kilobytes. This environment variable will cause Sun StorageTek EBS 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. "Compatible block size for UNIX and Microsoft Windows" on page 349 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 Sun StorageTek EBS 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 Sun StorageTek EBS software writes asynchronous filemarks. This setting has a minimal effect on performance.

NSR_DEV_LOAD_TIME_MEDIA_TYPE

NSR_DEV_LOAD_TIME_*MEDIA_TYPE* is the number of seconds that **nsrmmd** polls and waits for a drive to become ready after the library inserts a tape into the device. **NSR_DEV_LOAD_POLL_INTERVAL**_*MEDIA_TYPE* is used to set the number of seconds **nsrmmd** waits between polls during load time.

If the value of **NSR_DEV_LOAD_TIME**_*MEDIA_TYPE* is too short, there could be unnecessary load failures. If it is too long, then labeling new tapes takes longer than necessary. The minimum allowable value is 10 seconds. The maximum value is 600 seconds. For example:

NSR_DEV_LOAD_TIME_DTL8000=300

NSR_DEV_LOAD_POLL_INTERVAL_MEDIA_TYPE

NSR_DEV_LOAD_POLL_INTERVAL_*MEDIA_TYPE* is the number of seconds that **nsrmmd** waits between each attempt to read a newly

inserted tape. The minimum allowable value is 1 second, the maximum value is 30 seconds. For example:

NSR_DEV_LOAD_POLL_INTERVAL_DLT=10

NSR_DEV_LOAD_TRY_LIMIT_MEDIA_TYPE

NSR_DEV_LOAD_TRY_LIMIT_*MEDIA_TYPE* is the number of times that **nsrmmd** will attempt to open a drive. The **nsrmmd** program will poll the drive until the limit set in **NSR_DEV_LOAD_TIME**_*MEDIA_TYPE* has been reached. After the limit has been reached, it will retry until the **NSR_DEV_LOAD_TRY_LIMIT**_*MEDIA_TYPE* has been 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 Sun StorageTek EBS server to send commands to tape devices.

CDI support can be set in the Sun StorageTek EBS Administration interface:

1. In the server's **Sun StorageTek EBS 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 Sun StorageTek EBS 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.

The CDI feature is not supported in these cases:

- Within an NDMP environment.
- On an SGI IRIX server or storage node.

The CDI feature is not supported for tape devices connected directly on SGI systems. It can, however, be used by CDI- supported storage nodes (for example, Solaris) controlled by a Sun StorageTek EBS server that is running on SGI.

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 Sun StorageTek EBS 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 Powerlink website.

Device ordering

The Sun StorageTek EBS 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 Sun StorageTek EBS software no longer match the names of the devices as recognized by the operating system.

If device reordering occurs, the Sun StorageTek EBS software will not be able to use any affected drives until the configuration is manually corrected.

Beginning with Sun StorageTek EBS version 7.4 Service Pack 3, the Sun StorageTek EBS 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 Sun StorageTek EBS server stops all tape operations on that device and an error message will be posted. CDI must be enabled for this functionality. "Common device interface" on page 359 provides more information about enabling CDI.

Persistent binding and naming

Some operating systems provide the option known as "persistent binding" to permanently bind logical and physical addressing so that the associations are retained. Further, guaranteeing 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 will always see 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. Your 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 Sun StorageTek EBS, because the Sun StorageTek EBS 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 Sun StorageTek EBS prior to enabling persistent binding on the host, delete existing devices from the library resource and in the 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 Sun StorageTek EBS Management Console.

If devices have already been configured in Sun StorageTek EBS 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 **sjisn** commands, then verify that the device configuration within your Sun StorageTek EBS 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 **sjisn** command to determine the current order of the devices:

   ```
   sjisn scsidev@bus.target.lun
   ```

   where *bus.target.lun* is the SCSI address of the robotic arm returned by the inquire command in step 1, for example, 1.2.0.
3. Match the serial numbers of the devices in the sjisn output to the device names that correspond to these serial numbers in the inquire -cl output. This will give you the current device order by device filename.
4. Execute the **nsrjb** command to determine the order of devices as configured in Sun StorageTek EBS. Drive entries towards the end of the nsrjb output list the device order as configured in Sun StorageTek EBS.
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 driver ordering changes

After a drive ordering change has taken place and the Sun StorageTek EBS software is no longer correctly communicating with devices, you can correct the problem within your Sun StorageTek EBS configuration by using the Sun StorageTek EBS Console or the **jbedit** command line program.

Using Sun StorageTek EBS Console to correct drive ordering changes

To correct drive ordering changes by using the Sun StorageTek EBS Console:

1. Ensure that you have a current backup of the resource database.
2. Delete the library resource in the Sun StorageTek EBS Console. "How to delete a library" on page 162 provides details.
3. Rescan the library. "Scanning for libraries and devices" on page 141 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 Sun StorageTek EBS configuration
2. Use the **jbedit** command with the **-a** option to add the devices again.

"Using the jbedit command" on page 151, or the UNIX man page for **jbedit** or the *NetWorker Command Reference Guide provides* more information on the **jbedit** command.

Nonrewinding tape device usage (UNIX/Linux only)

Tape drives used as storage devices must be accessed by nonrewinding device files. The Sun StorageTek EBS 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 Sun StorageTek EBS 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/0m*bn*, 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 Sun StorageTek EBS software (for example /dev/rmt2.1)

Note: Never change a device pathname from nonrewinding (/dev/rmt/0*cbn*) 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 would have been overwritten by later saves.

Block-size mode (UNIX/Linux only)

Ensure that the block-size mode for tape devices that are used with Sun StorageTek EBS 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.

Silos

This section describes silos and silo devices. Silos and libraries are managed similarly by Sun StorageTek EBS 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 Sun StorageTek EBS 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 411 provides information on silo-specific, media-management concerns.

Sun StorageTek EBS software interactions with a silo

A Sun StorageTek EBS server acts as a client of the silo management software, which resides on the silo server. The Sun StorageTek EBS server communicates with the silo through the Silo Tape Library Interface (STLI), which must be installed on the Sun StorageTek EBS server that uses the silo.

To access the volumes and devices in a silo, the Sun StorageTek EBS 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 Sun StorageTek EBS 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 Sun StorageTek EBS 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 Sun StorageTek EBS software:

1. Install the silo management software on the silo server.
2. Install the **STLI** on the Sun StorageTek EBS server, if required. For more information, refer to the documentation from the silo vendor.

   Note: For example, for a Sun StorageTek EBS 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 Sun StorageTek EBS software is installed:

   - IBM 3494 on Solaris and AIX
   - StorageTek on Solaris, AIX, and HP-UX
   - DAS on Solaris, AIX, and HP-UX
3. Ensure that the Sun StorageTek EBS server is properly connected to the media devices in the silo.
4. Add the silo. Chapter 3, “Configuring Libraries and Devices” provides more information.

   Note: When adding a DAS silo, your hostname configuration files (for example, etc/hosts and etc/inet/ipnodes) must use the shortname for the server.

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 Sun StorageTek EBS server can use. There is no way to get this information from the Sun StorageTek EBS 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 Sun StorageTek EBS 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 Sun StorageTek EBS software, or test to determine which silo drive name matches with each Solaris device name.

DAS

The silo management software for the DAS silos is a program called DAS (Dual Attach Station). DAS acts as a front end for the silo control program called AMU. When the silo is configured, the silo administrator gives each drive a symbolic name. The symbolic name can be any alphanumeric string.

Set up DAS to work with Sun StorageTek EBS software

To set up DAS to work with Sun StorageTek EBS software the silo administrator must:

1. Configure DAS to accept commands from the Sun StorageTek EBS server or storage node computer.
2. Perform one of these tasks:
   - Use the **dasadmin allocd** command to allocate one or more devices to the Sun StorageTek EBS server or storage node.
   - Configure the Sun StorageTek EBS server or storage node as an administrator, to be able to execute the **dasadmin allocd** command to allocate devices from the Sun StorageTek EBS server or storage node computer.

Note: The Console cannot be used to configure a DAS silo. To configure a DAS silo, use the **jbconfig** command.

Find the assigned names for DAS devices in the silo

To find the names assigned to the devices in the silo, use the **dasadmin** utility, which is installed with the Sun StorageTek EBS software.

To find the names:

1. Set these three environment variables:
   - DAS_SERVER, the hostname of the silo management server, which runs DAS.
   - DAS_CLIENT, the hostname of the Sun StorageTek EBS server or storage node.

- ACI_MEDIA_TYPE, one of these values: 3480, CD_THICK, CD_THIN, DECDLT, 8MM, 4MM, D2, VHS, 3590, CD, TRAVAN, DTF, BETACAM, AUDIO_TAPE, or DAS_MEDIUM, same value as ACI_MEDIA_TYPE.

2. Type this command to see a list of drives and the hostnames to which they are allocated:

```
dasadmin ld
```

Releasing a silo device

When a silo device is configured for use with a Sun StorageTek EBS server, it is possible to restrict silo access only to the Sun StorageTek EBS 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 Sun StorageTek EBS 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 DAS and StorageTek silos

Environment variables must be set for DAS and StorageTek silos. Table 43 lists the environment variables to set.

Table 43 DAS and StorageTek environment variables

Silo model	Environment variables
DAS	DAS_SERVER = *name_of_DAS_server* DAS_CLIENT = *name_of_system_as_defined_to_DAS_server* DAS_MEDIUM = *type_of_tape_drive_used* ACI_MEDIA_TYPE = *type_of_tape_drive_used* For DAS_MEDIUM and ACI_MEDIA_TYPE, use one of these values: • 3480 • OD-Thick • OD-Thin • DECDLT • 8MM • 4MM • D2VHS • 3590 • CD • TRAVAN • DTF • BETACAM • AUDIOTAPE To specify the cartridge-access port or range of ports for depositing or withdrawing volumes, use these environment variables rather than the **-P** *ports* option with **nsrjb**: • DAS_INSERT_AREA_NAME=*port* • DAS_EJECT_AREA_NAME=*port* The default values for DAS_INSERT_AREA_NAME is I01, and the default value for DAS_EJECT_AREA_NAME is E01.
StorageTek	**For UNIX systems:** • CSI_HOSTNAME = *name_of_ACSLS_system* The following commands should also be running on the system and can be included in the Sun StorageTek EBS startup script: • <ebs_binaries_path>/mini_el & • <ebs_binaries_path>/ssi & For Windows systems: 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. Once installed and configured, it starts on reboot.

Set the environment variables for UNIX systems

To set environment variables for DAS and StorageTek silos for UNIX systems:

1. Locate the Sun StorageTek EBS startup script.
2. Using a text editor, open the Sun StorageTek EBS startup script and add the environment variables after the lines:

```
(echo 'starting Sun StorageTek EBS daemons:') >
/dev/console
```

3. Type the environment variables by using the format:

```
ENV_VAR = value
export ENV_VAR
```

4. Restart the Sun StorageTek EBS server in order for the environment variables to take effect.

Table 44 lists the locations of the Sun StorageTek EBS startup scripts.

Table 44 Sun StorageTek EBS Startup Commands

Operating system	Startup command
Solaris, IRIX	/etc/init.d/networker
HP-UX, HP Tru64 UNIX	/sbin/init.d/networker
AIX	/etc/rc.nsr

Set the environment variables for Microsoft Windows systems

To set environment variables for DAS silos for Microsoft Windows systems:

1. From the **Start** menu, select **Settings** > **Control Panel** > **System** > **Advanced** > **Environment Variables** > **New** > **System Variables**.
2. Type the environment variable names.
3. Restart the system in order for the environment variables to take effect.

Miscellaneous device operations

This section covers various additional device operation topics.

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 Sun StorageTek EBS 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 **Sun StorageTek EBS Libraries** field, type the appropriate time, in hours.
4. Click **OK**.

Stopping a device operation

To stop a device operation that is in progress:

1. In the Sun StorageTek EBS **Administration** window, click **Monitoring**.
2. Select the **Operations** tab.
3. Right-click an operation, and select **Stop**.
4. Click **Yes** to stop the operation, or **No** to resume it.

Display device operations messages

To display device operations messages:

1. In the Sun StorageTek EBS **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.

Sharing libraries among Sun StorageTek EBS hosts

The Sun StorageTek EBS software permits different Sun StorageTek EBS hosts (a Sun StorageTek EBS 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 Sun StorageTek EBS host to control the library's robotic arm, while other Sun StorageTek EBS 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 Sun StorageTek EBS host. Figure 31 shows how multiple Sun StorageTek EBS hosts can share library devices.

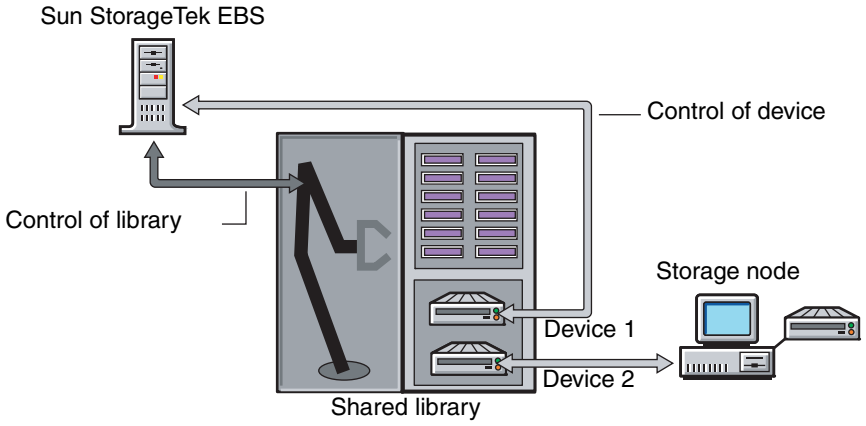


Figure 31 How library sharing works

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 Sun StorageTek EBS software automatically configures default sleep periods. Change these values only when troubleshooting a library's performance, or if a Sun StorageTek EBS 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 45.

Table 45 Library resource sleep attributes(1 of 2)

Attribute	Description	Default value
Load Sleep	Number of seconds that the Sun StorageTek EBS software waits for a library to complete loading a cartridge.	15 seconds
Unload Sleep	Number of seconds that the Sun StorageTek EBS software waits for a library to complete unloading a cartridge.	60 seconds
Eject Sleep	Number of seconds that the Sun StorageTek EBS software waits for a an eject operation to complete.	60 seconds
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

Table 45 Library resource sleep attributes(2 of 2)

Attribute	Description	Default value
Cleaning Delay	Number of seconds that the Sun StorageTek EBS 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	Number of minutes that the Sun StorageTek EBS software waits for a device to become idle after a backup has completed.	0 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 Sun StorageTek EBS server will use the same value. Therefore, the Sun StorageTek EBS 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 Sun StorageTek EBS server uses notifications to send messages about Sun StorageTek EBS 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 619 provides more information about notifications.

The Sun StorageTek EBS 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 has been corrected, it might be necessary to mount a volume so the Sun StorageTek EBS server can continue to back up or recover files.

Recycling compared to adding more volumes

The Sun StorageTek EBS 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 Sun StorageTek EBS server overwrites the data with new backups, but maintains the existing labels. "Changing a volume's mode" on page 392 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.

Resetting a library

A library must be reset each time the library and the Sun StorageTek EBS software become out of sync. Example 5, "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 Sun StorageTek EBS 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 (**sji** commands) and tape drives (**cdi** commands). These commands should only be used by the most knowledgeable of Sun StorageTek EBS 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 242 provides more information.

Device calibration

For information about the frequency and method for calibrating the loading mechanism for the device, refer to the library manufacturer's documentation.

Reordering tape drive numbers (Microsoft Windows only)

If more than one tape drive is attached to the Sun StorageTek EBS server when both the server and drives are shut down, restart all of the tape drives, either before or immediately after the Sun StorageTek EBS 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 Sun StorageTek EBS 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**.

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 Sun StorageTek EBS 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 Sun StorageTek EBS 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 406 has information about inventorying libraries.

Note: When a library has been partitioned, the Sun StorageTek EBS software does not become aware of the partitioning. This means that the entire physical library will be disabled, not just one partition.

12

Media Management

This chapter covers these topics:

- Storage management operations 384
- Auto Media Management 386
- Volume operations 389
- Media handling errors 410
- Media management in a silo 411
- Volume save sets 418

Storage management operations

This section describes the components involved in the operation of storage volumes done through the Sun StorageTek EBS 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 Sun StorageTek EBS servers for which they have access permission.

How the Sun StorageTek EBS server uses volume labels

A volume label is a unique internal code, applied by the Sun StorageTek EBS 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 242 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 Sun StorageTek EBS server uses the media database records to determine which volumes are needed for backing up or recovering data.

When it labels a volume, the Sun StorageTek EBS 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.

     Note: 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 Sun StorageTek EBS server selects a volume

When a backup takes place, the Sun StorageTek EBS 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 Sun StorageTek EBS 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.

   Note: 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.

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 389 provides information about appendable and recyclable volumes.

Data recovery and volume selection

The Sun StorageTek EBS 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. 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 Sun StorageTek EBS server displays all the volumes, in the order needed. During the recovery process, the server requests the volumes, one at a time.

Note: Sun StorageTek EBS 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 Sun StorageTek EBS from unloading the volume, the device should be set to service mode while the command is being run.

Automatic volume relabel

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. The mode of a volume can also be manually changed to recyclable. "Changing a volume's mode" on page 392 provides information about changing the mode of a volume.

Auto Media Management

This section describes how Sun StorageTek EBS works with Auto Media Management.

Auto Media Management gives the Sun StorageTek EBS server automatic control over media loaded in the storage device. When

Auto Media Management is enabled during device configuration, the Sun StorageTek EBS server automatically:

- Labels the volume (recognizes EDM labels and does not overwrite them).
- Mounts the volume.
- Overwrites volumes it considers to be unlabeled.

  The Sun StorageTek EBS server considers a volume to be unlabeled under the following conditions:

  - Has no internal label.
  - Is labeled with information other than a Sun StorageTek EBS label.
  - Is labeled with a Sun StorageTek EBS 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 Sun StorageTek EBS volumes are shared among devices with different densities.
- When the Auto Media Management feature is not enabled, the Sun StorageTek EBS server ignores unlabeled volumes and does not consider them for backup.

Note: The Sun StorageTek EBS server considers volumes that were labeled by a different application to be valid relabel candidates if Auto Media Management is enabled. Once the Sun StorageTek EBS 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 Sun StorageTek EBS labels

When Auto Media Management is used with tapes that have Sun StorageTek EBS 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 Sun StorageTek EBS server waits for the full, verified volume to be unmounted.

  Note: Nothing happens until the full volume has been unmounted by the operator.

- 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 Sun StorageTek EBS 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 Sun StorageTek EBS server continues to wait for a writable volume to continue the backup.

  Note: Again, nothing happens until the full volume has been unmounted by the operator.

- If the volume is recyclable and is a member of the required pool, it is recycled the next time a writable volume is needed.
- If the volume is unlabeled, it is labeled when the next writable volume is needed for a save.

Note: If a partially full volume is unmounted, the Sun StorageTek EBS 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 Sun StorageTek EBS server. Information about storage volumes is available for each device on a Sun StorageTek EBS 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 Sun StorageTek EBS 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 Sun StorageTek EBS 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 Sun StorageTek EBS 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 46 on page 391 lists the volume-related categories displayed in the Volumes detail table.

Table 46 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: 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 47 on page 392 lists the Sun StorageTek EBS volume modes and their definitions. "Changing a volume's mode" on page 392 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 47 Volume modes

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.
(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 392 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 *EMC NetWorker Command Reference Guide*) 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 392 provides information on changing the volume mode.

Changing a volume's mode

To change the mode of a volume:

1. In the server's Sun StorageTek EBS **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**.

Note: 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 Sun StorageTek EBS 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. Sun StorageTek EBS software applies a label template to create a unique internal label for each volume. The Sun StorageTek EBS server uses label templates and pool configuration settings to sort, store, and track data on media volumes.

Note: 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 7, "Sorting Backup Data" provides information about label templates and pools.

Label templates

Several preconfigured label templates are supplied with the Sun StorageTek EBS software. These preconfigured label templates cannot be deleted. "Naming label templates" on page 264 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.

   Note: 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 393 provides more information.

6. To you want 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 Sun StorageTek EBS 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.

Note: When a volume is relabeled, that volume is initialized and becomes available for writing again.

Empty slots in label operations

Slots that have been intentionally left empty (such as bad slots) are skipped during labeling operations. The Sun StorageTek EBS 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 47 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 Sun StorageTek EBS server records and tracks the label in the media database. The Sun StorageTek EBS 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 Sun StorageTek EBS 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 Sun StorageTek EBS 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 398 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 Sun StorageTek EBS 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 Sun StorageTek EBS 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 Sun StorageTek EBS 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 Sun StorageTek EBS 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 Sun StorageTek EBS software:

- When Sun StorageTek EBS 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 Sun StorageTek EBS 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 Sun StorageTek EBS 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 Sun StorageTek EBS 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 Sun StorageTek EBS 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 Sun StorageTek EBS 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_nam*e*.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.

Note: 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**.

Note: To perform an unattended backup using a stand-alone device, ensure that a labeled volume is mounted in the device before it is left unattended.

Label and mount a volume in one operation (stand-alone tape drive)

When more than one storage device is connected to the Sun StorageTek EBS 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 Sun StorageTek EBS 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 Sun StorageTek EBS server storage device, or ensure that a volume of this type is already present for the Sun StorageTek EBS 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.

      Note: If the volume is unlabeled, the Sun StorageTek EBS 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 Sun StorageTek EBS 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 Sun StorageTek EBS server disregards the assigned browse and retention policies. Therefore, only an administrator can mark the volume recyclable.

Note: 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 393 provides more information.

d. The **Mount After Labeling** attribute is selected by default. The Sun StorageTek EBS 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 Sun StorageTek EBS 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 396 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 prelabeled without being mounted.

To label a volume without mounting, follow the same procedures as for labeling and mounting in one operation, but clear the **Mount After Labeling** attribute in the **Label** window.

Mount an uninventoried volume

To mount a volume that is not included in the library inventory, but which is a valid (properly labeled) Sun StorageTek EBS 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 555 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 Sun StorageTek EBS software then inventories the specified slot.
10. Mount the inventoried volume as described in Chapter 1, "Overview".

Note: Unlabeled tapes may not be mounted for inventorying. Unlabeled tapes can only be mounted to be labeled. An attempt to mount an uninventoried volume by using unlabeled media results in an I/O error. The volume will also be ejected.

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, which applies only to shared libraries or libraries that are managed by AlphaStor software, 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. This attribute can be useful in configuring shared drives. The attribute appears on the Advanced tab of a device's Properties when 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. Select a time that is appropriate for the system configuration.

Using libraries with a volume import and export capability

The Sun StorageTek EBS 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 Sun StorageTek EBS 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 Sun StorageTek EBS 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 Sun StorageTek EBS 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 Sun StorageTek EBS 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 Sun StorageTek EBS 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 Sun StorageTek EBS software automatically records the slot number in which each newly labeled volume is located.

The Sun StorageTek EBS 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 396 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 Sun StorageTek EBS 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 Sun StorageTek EBS 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 Sun StorageTek EBS 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 Sun StorageTek EBS 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

Note: 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 Sun StorageTek EBS 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 *volid*

- For stand-alone drives:

  nsrmm -o full *volid*

Where *volid* 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.

   Note: This is descriptive information. Filling in this field does not send a volume anywhere.

5. Click **OK**.

Cloning volumes

"Volume cloning" on page 291 provides information about volume clone operations.

Media handling errors

The architecture of device drivers can produce media handling errors. The Sun StorageTek EBS software automatically retries a failed operation such as a **mount** or **read** of a volume. The number of times the Sun StorageTek EBS 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.

Reenable 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 Sun StorageTek EBS 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 Sun StorageTek EBS binaries.
2. Type this command:

```
nsrjb -HE
```

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 Sun StorageTek EBS software from overwriting volumes used by other programs.

Numbering a silo slot

In a library, the Sun StorageTek EBS software specifies many functions by slot number. A library has a fixed number of slots, and Sun StorageTek EBS 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 Sun StorageTek EBS pool.

"Mounting and unmounting volumes" on page 399 provides more information.

Labeling a silo volume

The Sun StorageTek EBS labels for volumes in a silo include both a regular Sun StorageTek EBS 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 393 and "Using barcode labels" on page 396 provide instructions on how to label silo volumes.

The use of barcodes with matching barcode labels and Sun StorageTek EBS 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 Sun StorageTek EBS 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 Sun StorageTek EBS 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 Sun StorageTek EBS software (or another application) associates with some of the data.

Using a silo with volume import/export capability

Sun StorageTek EBS 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 Sun StorageTek EBS software does not automatically inventory the volume after a deposit.

Either Sun StorageTek EBS 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 Sun StorageTek EBS 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.

Note: 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 Sun StorageTek EBS server is directed to the volumes it can use.

Note: 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 Sun StorageTek EBS 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 Sun StorageTek EBS 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 Sun StorageTek EBS software to query the silo management software to verify that the requested volume exists.

If the volume exists, the volume is allocated to the Sun StorageTek EBS 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.

Note: Silos do not display empty slots.

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*

"Sun StorageTek EBS software interactions with a silo" on page 365 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 Sun StorageTek EBS media database remain intact. If the volume is allocated again, Sun StorageTek EBS 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 Sun StorageTek

EBS server access the data. If the volume is to be removed from the silo for offsite storage, it must be removed with Sun StorageTek EBS 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 399 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.

"Sun StorageTek EBS software interactions with a silo" on page 365 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 Sun StorageTek EBS 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 406 provides information about inventorying a library.

The Sun StorageTek EBS software examines all of the volumes in the silo and compares the new list of volumes to the Sun StorageTek EBS media database. Then the Sun StorageTek EBS software produces a message listing any volumes located in the silo that are not in the media database.

When the Sun StorageTek EBS 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 Sun StorageTek EBS media database, the volume does not need to be loaded. The inventory proceeds rapidly. If, however, the Sun StorageTek EBS 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 48 on page 419 shows the attributes and their descriptions.
5. Click **OK** to close the **Volume Save Sets** window.

Table 48 Volume Save Sets window

Category	Description
Client	Name of the Sun StorageTek EBS 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 33 on page 294 provides a listing of save set values and descriptions.
Size	Save set size, in appropriate units.
Flags	First flag shows which part of the save set is on the volume: • **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. Second flag shows save set status: • **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. • **i**: Still in progress. Optional third flag: • **N**: NDMP save set • **R**: Raw partition backup (such as for a supported module). • **P**: Snapshot Optional fourth flag: • **s**: NDMP save set backed up by **nsrdsa_save** command to a Sun StorageTek EBS 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 420 and "The Save Set List tab" on page 421 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

13

Recovering Data

This chapter covers these topics:

- Choosing a recovery method 424
- Recovering by file selection 425
- Recovering by save set selection 436
- Recovering deduplication data 440
- Recovering with BMR 441
- Recovering encrypted data 441
- Disaster recovery 441
- Directing recoveries to another client 442
- Recovering online indexes for a Sun StorageTek EBS server 452
- Recovering the Console database 456
- Recovering backed-up client files from an old Sun StorageTek EBS server 457
- Recovering the Windows system configuration 458
- Recovering Windows volume mount points 467
- Recovering Windows DHCP and WINS databases 470
- Restoring Windows Content Index Server on Windows 2000 or later 470
- When you restart the service, the CIS re-indexes the entire catalog. 471

Choosing a recovery method

There are two main methods of recovering data:

- Recovery by file selection (index-based recovery)
- Recovery by save set selection

Recovery by file selection is the most flexible and easy to use method to recover data. However, save set recovery can be preferable especially when there is a large amount of data to be recovered or data is not available for quick access because it has passed its browse policy.

When to recover by file selection

Consider recovering data by file selection when:

- The exact name of the file or directory to recover is unknown and you want to browse the file system.
- A directed recovery is required. A directed recovery is a recovery where data that was backed up from one computer is recovered to another computer. "Directing recoveries to another client" on page 442 provides information about directed recoveries.
- Only the files that are marked for recovery are recovered (not extra files).
- It is important to keep the recovery operation as simple as possible. A save set recovery can be a more complex operation.

When to recover by save set selection

Consider recovering data by save set selection when:

- You want to restore a large bulk of lost files, for example, in the event of a total disk failure. In this case, you do not have to mark individual files for recovery.
- The recovery of extra files is acceptable.
- Quick access to the backup data is not available because the data's browse policy has expired. Quick access can be restored by re-creating the browsable client file index with the **scanner**

program. However, this can be a time-consuming process. "Restore a save set entry in the online indexes" on page 454 provides more information about the **scanner** program.

- Memory on the recovery host is an issue. Recovering by save set selection requires less memory than recovering by file selection.

Recovering by file selection

Recovering by file selection enables users to browse and select backed-up files and file systems for recovery. The earliest versions of files and file systems that are available for recovery are determined by the browse policy that was applied when the data was backed up. "About browse and retention policies" on page 206 provides more information about browse policies.

If backed-up data has exceeded its browse policy, you can still recover the data by using save set recovery. "Recovering by save set selection" on page 436 provides more information about save set recovery. You could also recover data by re-creating the client file index with the **scanner** program. Re-creating the client file index enables users to browse and select files for recovery. However, re-creating the client file index can be a time-consuming process. "Restore a save set entry in the online indexes" on page 454 provides more information about re-creating the client file index.

You can also recover files from other client computers by using directed recovery. "Directing recoveries to another client" on page 442 provides details.

Recovering data on UNIX systems

To recover data to a UNIX client:

1. Start the **nwrecover** program by typing this command:

   nwrecover

   - Use the **-s** *servername* option with **nwrecover** to specify a particular Sun StorageTek EBS server on the network to use when recovering data. The /servers file, located in /nsr/res/servers, contains an entry for each available server.

If the **-s** option is not entered and there is only one server detected, that server will be connected to 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.

- Use the **-c** *client* option to specify a particular Sun StorageTek EBS client on the network. If the **-c** option is not used, the current client is assumed.

The **nwrecover** program displays a representation of the client's file system.

2. Mark each directory and/or file to be recovered by selecting the checkbox next to each directory or file listed.

Note: To recover data from a scheduled backup other than the most recent one, you must change the browse time. See "Recovering an earlier version of a file" on page 429.

3. Click **Start** to begin the recovery.
4. In **Recover Options** dialog box, select any appropriate options and click **OK**.

   "Resolving recovered data conflicts" on page 431 and "Relocating recovered data" on page 433 have more information on options in the **Recover Options** dialog box.

   The time it takes to recover the files, depends on file size, network traffic, server load, and tape positioning. During this time, messages appear so that you can monitor the recovery.

   If 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
```

Recovering data on Windows systems

The information in this sections refers to using the NetWorker User program to recover data. If the Sun StorageTek EBS Module for Microsoft Applications Client User program is installed on the client computer, VSS data should be recovered with the NetWorker Module for Microsoft Applications Client User program. The *EMC NetWorker Module for Microsoft Applications Administration Guide provides more*

information about the NetWorker Module for Microsoft Applications Client User program.

Note: If VSS is licensed and enabled, the databases must be taken offline before recovering Micorsoft Exchange Server or SQL Server data.

To recover data to a Windows client using the NetWorker User program:

1. In the NetWorker **User** program, click **Recover** to open the **Source Client** dialog box.
2. Select the source client with the data you want to recover, and click **OK**. The local client is the default selection.
3. Select the destination client for the recovered data, and click **OK**. The local client is the default selection.
4. In the left pane of the **Recover** window, click the appropriate directory folder.
5. Mark each directory and/or file to be recovered by performing one of the following:
   - Select the directory or file, and from the **File** menu, select **Mark**.

     To clear an item, select **Unmark** from the **File** menu.
   - Click the checkbox next to the directory or filename.

     To clear an item, click the checkbox again.

   If you select **My Computer** for recovery, all drives are recovered.

   Note: To recover data from a scheduled backup other than the most recent one, you must change the browse time. "Recovering an earlier version of a file" on page 429 provides details.
6. Click **Start** to begin the recovery.
   - It takes the Sun StorageTek EBS 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.

   Note: If an error occurs while recovering Microsoft Exchange Server or Microsoft SQL Server data by using VSS, you must restart the recovery process.

- If there is a naming conflict during recovery, you are prompted to select a workaround. "Resolving recovered data conflicts" on page 431 provides information about naming conflicts.
- If 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
```

- If 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 Sun StorageTek EBS log file.

Note: The NetWorker User program does not recover any files in parallel. All files are recovered sequentially. To recover files using parallelism, perform recoveries from the command line.

Recovering data from the command prompt

You can also perform a recovery by using the **recover** program. For example, to recover /myfile from the server *jupiter*, type:

```
recover -s jupiter
   Enter directory to browse: /
   recover: Current working directory is /
   recover> add /myfile
   /
   1 file(s) marked for recovery
   recover> recover
   Recovering 1 file into its original location
   Volumes needed (all on-line):
      jupiter.mars.com.001 at //./Tape.0

   Received 1 file(s) from NSR server `jupiter'
   Recover completion time: Tue Jan 21 08:33:04 2009
   recover>
```

- If you do not specify the **-s** option with the **recover** command, the file(s) will be recovered from the Sun StorageTek EBS server defined in the /nsr/res/servers file that comes first in alphabetical order.
- To recover files or directories that begin with a dash (-) such as -Accounting, try one of the following options:

- Run the **recover** command and enter **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 and then enter **add .** to add the current directory and its contents for recovery.
- If current directory is /temp and -Accounting is located in the /temp directory, run the **recover** command and then enter **add /temp/-Accounting** and the file or directory -Accounting and its contents will be added for recovery.

◆ You cannot recover Windows SYSTEM or VSS SYSTEM save sets by using the **recover** command in interactive mode. For information about recovering Windows SYSTEM save sets, see "Recovering the Windows system configuration" on page 458. The *EMC NetWorker Command Reference Guide* or the UNIX man pages provide information about **recover** command.

◆ To avoid using the Windows version of **recover.exe** on Windows operating systems, do one of the following:

- Include <Sun StorageTek EBS_install_path>\bin\recover.exe at the command prompt
- Ensure that <Sun StorageTek EBS_install_path>\bin is listed before %SystemRoot%\System32 in the $PATH environment variable.

Recovering an earlier version of a file

You can browse previous backups to recover an earlier version of a file. This is useful if the required data was deleted from the current file system and is not present in the most recent backup.

The earliest versions of files and file systems that are available for recovery are determined by the browse policy that was applied when the data was backed up. For example, to recover a lost file that was backed up six months ago requires the appropriate browse policy. "About browse and retention policies" on page 206 provides more information about browse policies.

Recovering an earlier version of a file

Before changing the browse time to recover an earlier version of a file, verify that no files are currently marked for recovery. If files are currently marked for recovery and the browse time is changed, the

marked files will still be selected for recovery even though they may have been backed up after the selected browse time.

On Windows

To recover an earlier version on Windows:

1. In the NetWorker **User** program, click **Recover** to open the **Source Client** dialog box.
2. Select the source client whose data you want to recover, and click **OK**.
3. Select the destination client for the recovered data, and click **OK**.
4. From the **View** menu, select **Change Browse Time**.
5. In the **Change Browse Time** dialog box, select a new day within the calendar. Select **Previous Month** or **Next Month** to change from the current month.
6. 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.
7. Click **OK**.
8. Select the files or directories to recover as described in "Recovering by file selection" on page 425.

Note: If there are no entries for files backed up at the time selected in the Time field, the files that were backed up at the closest preceding time are displayed.

On UNIX

To recover an earlier version on UNIX:

1. In the **nwrecover** program, select **Change Browse Time** from the **Options** menu.
2. In the **Choose Date** dialog box, select a new day within the calendar. Click the forward or back arrows to change from the current month.
3. In the optional **Time** attribute, change the time of day to use by changing the displayed time. Use the 24-hour format.
4. Click **OK**.
5. Select the files or directories to recover as described in "Recovering by file selection" on page 425.

Note: If there are no entries for files backed up at the time selected in the **Time** attribute, the files that were backed up at the closest preceding time are displayed.

Resolving recovered data conflicts

By default, the Sun StorageTek EBS server attempts to recover data to its original location. If a name conflict occurs, the Sun StorageTek EBS software can respond either automatically or on an file-by-file basis.

Resolving data conflicts on UNIX

The **Recover Options** dialog box, which appears each time a recovery is started, provides options for resolving naming conflicts during recovery. Table 49 on page 431 lists the available conflict resolution options.

Table 49 Options for resolving naming conflicts on UNIX

Option	Action
Relocate files to	Recovers selected files to the specified location.
Rename recovered file	Renames the recovered files by appending each conflicting name with .R.
Discard recover file	Discards the recovered file and keeps the existing file.
Overwrite existing file	Replaces the file on the file system with the recovered version.
Prompt me at every file conflict	If selected, redisplays the Naming Conflict dialog box for every file that has a naming conflict.

Resolving data conflicts on Windows

By default, the Naming Conflict dialog box appears each time there is a filename conflict during a recovery. Table 50 on page 431 lists the conflict resolution options available during a recovery.

Table 50 Options far resolving naming conflicts on Windows(1 of 2)

Option	Action
Rename Recovered File	Renames the recovered files. "Rules for renaming file on Windows" on page 432 provides more information.

Table 50 Options far resolving naming conflicts on Windows(2 of 2)

Option	Action
Discard Recover File	Discards the recovered file and keeps the existing file.
Overwrite Existing File	Replaces the file on the file system with the recovered version.
Suppress Further Prompting	Automatically performs the selected option if additional conflicts are encountered during this recovery operation.

Rules for renaming file on Windows

When a naming conflict is encountered during recovery, the Sun StorageTek EBS server uses these guidelines to rename the recovered file:

- By default, a tilde (~) is appended to the beginning of the name of the recovered file:

  ~<filenname>

- If a file named ~<filenname> already exists, the recovered file is renamed ~00_<filenname>.
- If a file named ~00_<filenname> already exists, the recovered file is renamed ~01_<filenname>, and so forth to ~99_<filenname>.
- If this fails, the file will not be renamed automatically:
  - If the **recover** command is being used, the Sun StorageTek EBS server fails to rename the file.
  - If the NetWorker **User** program is being used, the user is prompted to type a name for the file.

Resolve name conflicts automatically on Windows

You can configure the Sun StorageTek EBS software to resolve all name conflicts automatically before a recovery is started.

To automatically resolve name conflicts before starting a recovery:

1. In the NetWorker **User** program, click **Recover** to open the **Source Client** dialog box.
2. Select the source client with the data to recover, and click **OK**.
3. Select the destination client for the recovered data, and click **OK**.
4. Select **Recover Options** from the **Options** menu.
5. Select a conflict resolution option, as described in Table 50 on page 431 and then click **OK**.

Another way to avoid name conflicts is to relocate the recovered data prior to recovery. "Relocating recovered data" on page 433 provides more information.

Relocating recovered data

By default the Sun StorageTek EBS server attempts to copy a recovered file to its original location. If an existing file on the local disk has the same name as the file being recovered, the server prompts for instructions on how to recover the file.

Alternatively, you can instruct the Sun StorageTek EBS server to relocate the recovered file prior to the recovery.

Relocate recovered data on Windows

To relocate recovered data:

1. In the NetWorker **User** program, click **Recover** to open the **Source Client** dialog box.
2. Select the source client with the data to be relocated, and click **OK**.
3. Select the destination client for the data, and click **OK**.
4. Select **Recover Options** from the **Options** menu.
5. In the **Relocate Recovered Data To** field, type the directory where the data should be relocated and then click **OK**.

   Note: Provide the full pathname of the relocation directory.

6. Perform a data recovery as described in "Recovering data on Windows systems" on page 426.

Relocate recovered data on UNIX

The **Recover Options** dialog box, which appears each time a recovery is started, allows you to select the location to which files should be relocated during recovery.

To relocate recovered data:

1. In the **nwrecover** program, start a recovery.
2. In the **Relocate file to** field of the **Recover Options** dialog box, click **Browse** and select the directory where the data should be relocated, or type the full pathname.

3. Click **OK**.

Viewing versions of a directory or file

Use the Versions option to view the backup details of a directory or file. Details include the name, size, modified time, volume, location, and so on.

View version details on Windows

To view the version details of backed-up data:

1. In the NetWorker **User** program, select **Relocate** from the **Options** menu.
2. Select the source client with the backup versions to view, and click **OK**.
3. Select the destination client for the recovered data, and click **OK**.
4. Highlight the file or directory to be recovered.
5. From the **View** menu, select **Versions**.

Versions are sorted according to the backup time (the most recent backup first). The selected version matches the one currently selected in the **Recover** window.

View version details on UNIX

To view the version details of backed-up data:

1. In the **nwrecover** program, select an item to view its details.
2. From the **Selected** menu, select **Versions**.

Versions are sorted according to the backup time (the most recent backup first).

Searching for a recoverable file

Users can search for and recover the most recently backed-up version of a file.

On Windows

To search for a recoverable file:

1. In the NetWorker **User** program, click **Recover** to open the **Source Client** dialog box.
2. Select the source client with the data to recover, and click **OK**.

3. Select the destination client for the recovered data, and click **OK**.
4. From the **File** menu, select **Find**.
5. Type the name of the file or directory. Use wildcards to expand the search; without wildcards, partial filenames result in a no match being found.

On UNIX To search for a recoverable file:

1. In the **nwrecover** program, select **Find** from the **File** menu.
2. Type or select the base directory that will be searched.
3. Type the name of the file or directory. Use wildcards to expand the search, or use the **Contains** drop-down button to modify search criteria.
4. Click **More Options** to further modify the search criteria, such as selecting a time of the backup and an owner of the file.
5. Click **Find**.

Viewing the volumes required for data recovery

Users can view a list of the volumes needed to recover files and directories.

On Windows To view a list of the volumes required for data recovery:

1. In the NetWorker **User** program, click **Recover** to open the **Source Client** dialog box.
2. In the **Source Client** dialog box, select the client with the data to be recovered and click **OK**.
3. In the **Destination Client** dialog box, select the client to which the data will be recovered and click **OK**.
4. In the **Recover** window, select the item with the volume requirements to be viewed.
5. Select **Required Volumes** from the **View** menu.

On UNIX To view a list of the volumes required for data recovery:

1. In the **nwrecover** program, select the item with the volume information requirements to be viewed.
2. Select **Required Volumes** from the **Selected** menu.

Recovering by save set selection

Save set recovery enables you to recover data from a save set rather than by browsing and selecting files for recovery. Generally it is preferable to browse and select files for recovery. However, save set recovery may be preferable when there is a large amount of data to recover or when quick access to the backup data is not available because the data's browse policy has expired. "Choosing a recovery method" on page 424 provides more information about deciding whether to perform a save set recovery or a recovery by file selection.

Backed up data that has not exceeded its retention policy is eligible to be recovered by save set recovery. However, if backed up data has exceeded its retention policy, it may still be recoverable by using the **scanner** program. "About browse and retention policies" on page 206 provides information about retention policies. "Restore a save set entry in the online indexes" on page 454 provides information about recovering data that has exceeded its retention policy.

Note: When you perform a save set recover operation, recover the last full backup first, then recover levels 1 to 9 and incremental backups in the chronological order that they were backed up. "Backup levels" on page 190 provides information about the relationship between full backups, backup levels 1 to 9, and incremental backups.

Recover a save set on Windows

Note: Only members of the Windows Administrators group have permission to perform a save set recovery on Windows. In the case of a remote Sun StorageTek EBS client, members of the client's Windows Administrators group must have the Recover Local Data privilege to perform the save set recover. Members of the server's Windows Administrators group automatically have this privilege.

To perform a full save set recovery of a client:

1. In the NetWorker **User** program, select **Save Set Recover** from the **Operation** menu.
2. In the **Source Client** dialog box, select the appropriate client and click **OK**.

3. From the **Save Set Name** list of the **Save Sets** dialog box, select the name of the appropriate save set.
4. Select the particular version of the save set (if there are multiple versions).
5. Click **OK** to begin the recovery. The recovery status appears in the **Recover Status** window.

Recover a save set on UNIX

Note: Only root has the permission to perform a save set recovery on UNIX. In the case of a remote Sun StorageTek EBS client, root@*client* must have the Recover Local Data privilege to perform the save set recover. The root@<ebs_server> user automatically has this privilege.

To perform a full save set recovery of a client:

1. In the **nwrecover** program, select **Recover Save Set** from the **Operation** menu.
2. From the **Name** list, select the name of the appropriate save set.
3. In the **Instances** list, select the particular version of the save set (if there are multiple versions).
4. Click the **Recover** button.
5. Select the appropriate options in the **Recover Options** dialog box and click **OK**.

Recover a save set from the command prompt

To perform a save set recovery from the command prompt, use the **recover** command with this syntax:

```
recover -s servername -S ssid -d directory_path
```

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

Recover individual files or directories from a save set on Windows

To recover specific data from a save set:

1. In the NetWorker **User** program, select **Save Set Recover** from the **Operation** menu.
2. In the **Source Client** dialog box, select the appropriate client and click **OK**.
3. From the **Save Set Name** list of the **Save Sets** dialog box, select the name of the appropriate save set.
4. Select the particular version of the save set (if there are multiple versions).
5. Click **Files** to open the **Files** dialog box.
6. Type the exact names of the files and directories (including the full case-sensitive paths) that you want to recover in the **Files** text box. Each file or directory must appear on a separate line.
7. Click **OK** to begin the recovery. The recovery status appears in the **Recover Status** window.

Recover individual files or directories from a save set on UNIX

To recover specific data from a save set on UNIX, perform a regular recovery as described at "Recovering data on UNIX systems" on page 425.

View volumes required for save set recovery

Use this procedure to display a list of each volume that needs to be available to recover a save set. If no volumes appear during this procedure, then no entries for the volume exist in the media database. Use the **scanner** program to re-create the media database entries. "Restore a save set entry in the online indexes" on page 454 provides information about using the **scanner** program to re-create the media database.

On Windows

To view a list of the volumes required for a save set recovery:

1. In the NetWorker **User** program, select **Save Set Recover** from the **Operation** menu.
2. In the **Source Client** dialog box, select the appropriate client and click **OK**.
3. From the **Save Set Name** list of the **Save Sets** dialog box, select the name of the appropriate save set.

4. Select the particular version of the save set (if there are multiple versions).
5. Click **Required Volumes** to display the volumes required for the recovery.

On UNIX To view a list of the volumes required for a save set recovery:

1. In the **nwrecover** program, select the appropriate save set instance.
2. Select **Required Volumes** from the **Selected** menu.

Relocate recovered save set data

You can relocate recovered save set data from the data's original backup location to another place on the file system.

On Windows To relocate recovered save set data:

1. In the **nwrecover** program, select **Save Set Recover** from the **Options** menu.
2. In the **Source Client** dialog box, select the appropriate client and click **OK**.
3. Click **Options** to open the **Recover Options** dialog box.
4. Type the path in the **Relocate Recovered Data To** text box or click **Browse** to select the destination directory.
5. Click **OK**.

On UNIX The **Recover Options** dialog box, which appears each time a recovery is started, allows you to select the location to which files should be relocated during recovery.

To relocate recovered data:

1. In the **nwrecover** program, start a recovery
2. In the **Relocate file to** field of the **Recover Options** dialog box, click **Browse** and select the directory where the data should be relocated, or type the full pathname.
3. Click **OK**.

Resolve conflicts with recovered save set data

You can configure the Sun StorageTek EBS server to resolve conflicts during the recovery.

On Windows

To resolve data recovery conflicts with save set data:

1. In the NetWorker **User** program, select **Save Set Recover** from the **Operation** menu.
2. In the **Source Client** dialog box, select the appropriate client and click **OK**.
3. Click **Options** to open the **Recover Options** dialog box.
4. Select the appropriate option for the **Duplicate File Resolution** attribute. Table 51 on page 440 lists the options and actions taken for the Duplicate File Resolution.

Table 51 Options far resolving save set recovery naming conflicts on Windows

Option	Action
Always Prompt	Indicates that automatic naming conflict resolution is disabled.
Rename the Recovered File	Renames the recovered files. "Rules for renaming file on Windows" on page 432 provides more information.
Discard the Recover File	Discards the recovered file to prevent overwriting the existing file.
Overwrite the Existing File	Replaces the file with the recovered version.

5. Click **OK**.

On UNIX

The **Recover Options** dialog box, which appears each time a recovery is started, provides options for resolving naming conflicts during recovery. Table 49 on page 431 lists the available conflict resolution options.

Recovering deduplication data

"Recovering deduplicated data" on page 868 describes how to perform a recovery of deduplication data.

Recovering with BMR

Chapter 27, "Support for HomeBase" describes how to perform a BMR recovery or server migration.

The EMC HomeBase product documentation provides details and instructions pertaining to the HomeBase Server and HomeBase Agents.

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 Sun StorageTek EBS server. "Set the Datazone pass phrase for a Sun StorageTek EBS server" on page 97 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_phrase***3**

Note: If an incorrect pass phrase or no pass phrase is entered, encrypted data is not recovered. Instead, the filenames will be created without data. However, if unencrypted data is also selected for recovery, it will be recovered.

Disaster recovery

Recovering data from Sun StorageTek EBS servers that have been damaged or disabled as a result of a disaster or other catastrophic event requires more complex steps than are covered in this chapter. Optimum recovery also requires careful planning. For more information, refer to the *EMC NetWorker Disaster Recovery Guide*. For

the most recent disaster recovery information, refer to the *EMC Sun StorageTek EBS Release Supplement*.

Keep a printed copy of this documentation in a safe, convenient place.

Do not use the **scanner** program to restore the bootstrap after a disaster in which you lose the media database or the server's res files. Instead, use the **mmrecov** command to restore these files.

Note: The **scanner** program can recover only the hash data from deduplication node backups if the save sets have been deleted from the Avamar server that is the deduplication node. If the backup-to-tape method described in "Backup-to-tape for deduplication clients" on page 309 is used regularly, it provides a tape backup of the deduplicated data up to a particular point in time.

Directing recoveries to another client

Directed recovery allows a Sun StorageTek EBS administrator to remotely recover backed-up data from a Sun StorageTek EBS client (the source client) and direct it to another Sun StorageTek EBS client (the destination client).

Note: Directed recovery of the Windows SYSTEM STATE is not supported.

Figure 32 on page 443 shows a user on client saturn performing a directed recovery of data from a remote client (source) to the client mars (destination).

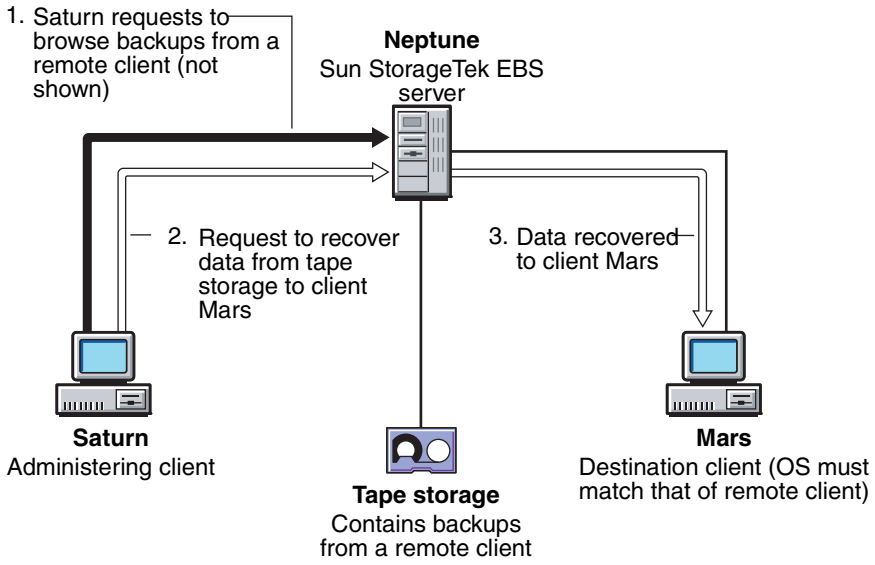


Figure 32 **A directed recovery from a remote client**

Note: For directed recoveries of deduplicated backups, the destination client must also have NetWorker release 7.4 with Service Pack 1 or later installed.

Uses for directed recovery

A directed recovery is particularly useful in these situations:

- When centralized administration is a requirement. All recoveries can be directed from a single administering computer.
- When recovering data to a shared server where users cannot recover the data themselves.
- When a computer is inoperable, a hard disk has crashed, or a network no longer recognizes a computer. Configure the replacement computer by recovering files from another client's backups.

- When you want to transfer files between two clients. For example, if one client's AUTOEXEC.BAT or .profile file is appropriate for a new client, recover it to the new client's hard disk.

Note: The original file ownership and permissions are always retained with directed recoveries.

Requirements for directed recovery

This section contains the requirements for directed recoveries and the procedures to enable directed recovery.

Client restrictions for directed recoveries

The destination client in a directed recovery can be any Sun StorageTek EBS client subject to these guidelines:

General restrictions:

- The source and destination clients must be of the same platform (UNIX-to-UNIX or Microsoft Windows-to-Microsoft Windows).
- The source and destination clients must use the same type of file system (for example, UXFS-to-UXFS, NTFS-to-NTFS).
- The administering host must be a client of the Sun StorageTek EBS server. The administering host is the host from which the graphical user interface is displayed.
- If the administering client is neither the source nor the destination client, it need not have the same platform and type of file system as the other clients.

Windows specific restrictions:

- For Sun StorageTek EBS clients running Microsoft Windows XP Professional, the Sun StorageTek EBS server and clients must be in a Microsoft Windows Domain, regardless of the username and password.
- For a Windows destination client, Windows File and Print Sharing must be enabled. For more information, refer to the Microsoft Windows documentation.
- Directed recovery of SYSTEM or VSS SYSTEM save sets is not supported.

- "Restrictions for Windows command line directed recovery" on page 451 provides additional restrictions.

UNIX specific restrictions

- Restrictions related to relocating non-ASCII directories on UNIX machines:
  - If the remote directory is an existing non-ASCII directory, the locale of the browsing session from which the directed recover is started must match the locale in which the remote directory was created.
  - If the remote directory does not exist, the relocation directory will be created on the remote machine's file system based on the locale of the browsing session.

Enabling directed recoveries

Directed recovery is a restricted Sun StorageTek EBS function available only to users who have all privileges except for the Change security settings privilege, which is not required for directed recovery.

Note: Be selective in granting users the ability to perform directed recoveries. Directed recovery privileges enable users to browse all clients and recover data to any client.

If directed recoveries will be performed on a regular basis, establish the required permissions for each Sun StorageTek EBS server and client in the enterprise prior to accepting data recovery requests. "Managing server access" on page 609 provides general information about access requirements.

Access requirements to perform directed recoveries are as follows:

- Users in the Administrators group on the Sun StorageTek EBS server are automatically granted the necessary privileges. Table 50 on page 431 provides specific access privileges that are required if a user is not a member of the Sun StorageTek EBS Administrators group.
- The recovery must be launched by the root user or Windows Administrator on the administering client.

- The root user or Windows Administrator on the administering client must have the Remote Access All Clients privilege on the Sun StorageTek EBS server.
- *SYSTEM@destination_client* must be listed in the source client's Remote Access list.
- *SYSTEM@destination_client* must have the Recover Local Data privilege.

If the user on the destination client does not have Remote Access All Clients privileges, then the *user@destination_client* must be listed in the source client's Remote Access list. "Editing a client" on page 650 provides information about editing a client.

Table 52 on page 446 shows the access permissions and authorizations required to perform various combinations of directed recoveries.

Table 52 Access permissions used in directed recoveries (1 of 2)

Administering client	Source client	Destination client	Who may perform a directed recovery
UNIX	UNIX	UNIX	• The user must have the Remote Access All Clients privilege or the user@destination_client must be listed in the source client's Remote Access list. • The user@destination_client must have the Recover Local Data privilege.

Table 52 Access permissions used in directed recoveries (2 of 2)

Administering client	Source client	Destination client	Who may perform a directed recovery
Windows	Windows	Windows	• The user must have the Remote Access All Clients privilege or the user@destination_client must be listed in the source client's Remote Access list. • The user@destination_client must have the Recover Local Data privilege. • The user@administering_client must be added to the Windows Administrator Group on the administering client.
Windows	UNIX	UNIX	• The user must have the Remote Access All Clients privilege or the user@destination_client must be listed in the source client's Remote Access list. • The user@destination_client must have the Recover Local Data privilege. • The recovery must be launched by a system administrator on the administering client.
UNIX	Windows	Windows	• The user must have the Remote Access All Clients privilege or the user@destination_client must be listed in the source client's Remote Access list. • The user@destination_client must have the Recover Local Data privilege. • The recovery must be launched by root on the administering client.

Note: For security reasons, if a user with incorrect privileges attempts to perform a directed recovery, no error message will be issued.

Setting the Disable Directed Recover attribute

Directed recovery can be turned on or off for Sun StorageTek EBS clients by using the Disable Directed Recover attribute in the client's nsrla.res database. The default setting for the Disable Directed Recover attribute is No, which means that the client will accept directed recoveries from remote hosts.

To disable a remote host from directing a recovery to a client, set the **Disable Directed Recover** attribute to **Yes**.

"Editing a client NSRLA database" on page 653 provides information about editing the nsrla.res database.

Adding a client computer to the remote client's servers file

"Editing the servers file" on page 674 describes how to add a computer hostname to the remote Sun StorageTek EBS client's servers file.

Perform a directed recovery between Windows clients from the GUI

Note: Directed recovery of save sets on Windows is supported from the command line only. "Perform a directed save set recovery by using the command line" on page 452 provides more information.

To perform a directed recovery:

1. Start the NetWorker User program by using the **-s** option to connect to the Sun StorageTek EBS server to which the source client data is backed up:

   winworkr -s *servername*
   If the **-s** option is not entered and there is only one server detected, that server will be connected to 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.

   Note: Before you can perform a directed recovery of a virtual cluster client, the Client resource must be configured on the Sun StorageTek EBS server. "Client configuration" on page 650 provides instructions on creating a Client resource. Be sure that the Client resource Remote Access attribute contains *both* of the cluster node names.

2. Click **Recover** to open the **Source Client** dialog box.
3. Select the source client, and click **OK**.
4. Select the destination client for the recovered data, and click **OK**.
5. In the **Recover** window, select the files and other items to recover.
6. Click **Start** to begin the directed recovery.

Special data handling features, such as compression, are not available for the items you select.

The selected destination client is the recipient of the recovered data. If files with the same name already exist, they are overwritten unless they have mixed-case names. Recovered files that contain mixed-case

filenames do not replace newer versions of the same file that have been given a single-case filename.

Perform a directed recovery between UNIX clients from the GUI

To perform a directed file recovery:

1. Start the nwrecover program by using the **-s** option to connect to the Sun StorageTek EBS server to which the source client data is backed up:

   nwrecover -s *servername*
   The /servers file, located in /nsr/res/servers, contains entries for the Sun StorageTek EBS servers. If there is more than one server, the **Change Server** dialog box allows you to choose which server to connect to.

   Note: Before you can perform a directed recovery of a virtual cluster client, the Client resource must be configured on the Sun StorageTek EBS server. "Client configuration" on page 650 provides instructions on creating a Client resource. Be sure that the Client resource Remote Access attribute contains *both* of the cluster node names.

2. Select files and other items to recover.
3. Click the **Start a recover operation** button.
4. In the **Recover Options** dialog, click **More Options**.
5. In the **Directed Recover** field, enter or select the name of the client to which the files should be recovered.
6. Click **Recover** to open the **Source Client** dialog box.
7. Click **OK**.

To perform a directed save set recovery:

1. In the **nwrecover** program, select **Recover Save Set** from the **Operation** menu.
2. From the **Name** list, select the name of the appropriate save set.
3. In the **Instances** list, select the particular version of the save set (if there are multiple versions) and click **OK.**
4. In the **Recover Options** dialog, select the relocation path in the window.
5. Click the **Recover** button.

6. Select the appropriate options in the **Recover Options** dialog box and click **OK**.

Perform a directed recovery from the command line

Use the **recover** command to perform a directed recovery from the command line. The **recover** command has a variety of options that you can add to modify the basic browse and recover functionality. Table 53 on page 450 lists the options available with the **recover** command.

Table 53 Recover options

Option	Meaning
-d *destination*	Specifies the name of the directory for relocated recovered files. With this option, include a full path for destination, and use the proper syntax for the destination client's platform. The use of partial paths is not supported.
-s *server*	Specifies the name of the source client's backup server.
-c *client*	Specifies the name of the source client.
-S *SSID*	Specifies a save set to recover.
-R *recover_destination*	Specifies the name of the destination client.
-i (N,Y, R)	Specifies what the Sun StorageTek EBS server should do when there is a naming conflict between a recovered file and an existing file: • **N** — Skip restoring the file. • **Y** — Overwrite the existing file. • **R** — Rename the file. Each conflicted filename is appended with .R.
directory_name	Specifies the initial directory in which to begin browsing.

The *EMC NetWorker Command Reference Guide* or the UNIX man pages provide complete information about the **recover** command.

Example 32 The recover command

In this example, the source client's backup server is *venus*, the source client is *mars*, and the destination client is *jupiter*. All of the files in the specified directory are to be recovered. The recovered files should overwrite any existing files with the same name.

Type this command:

```
recover -s venus -c mars -R jupiter -iY directory_name
```

Restrictions for Windows command line directed recovery

In addition to the restrictions listed in "Client restrictions for directed recoveries" on page 444, be aware of the following considerations when performing Windows directed recoveries from the command line.

- If the Sun StorageTek EBS servers and clients are running Microsoft Windows and the NetWorker server and target recover client are in the same domain:
  - Start the NetWorker server (nsrd) as a domain user that is in the Windows Administrators group on the NetWorker server machine.
- If the Sun StorageTek EBS servers and clients are running Microsoft Windows and the NetWorker server and target recover client are *not* in a domain, or they are *not* in the same domain, ensure that:
  - The NetWorker server (nsrd) is started as a local user that exists on both machines.
  - The local user must have the same password on both machines.
  - The local user must also be in the Windows Administrators group on the NetWorker server machine.
- To avoid using the Windows version of **recover.exe**, do one of the following:
  - Include <Sun StorageTek EBS_install_path>\bin\recover.exe at the command prompt.
  - Ensure that <Sun StorageTek EBS_install_path>\bin is listed before %SystemRoot%\System32 in the $PATH environment variable.

Note: If a command line directed recovery fails, try to perform the directed recovery from the NetWorker User program. "Perform a directed recovery between Windows clients from the GUI" on page 448 provides details.

Perform a directed recovery of specific files and directories by using the command line

To perform a directed recovery of specific files and directories:

1. Type the **recover** command by using this syntax:

```
recover -s servername -c source_client -R
destination_client -i recover_option [directory_name]
```

2. Navigate to the directory with the file to be selected for directed recovery. For example:

   ```
   recover> cd /var/adm
   ```

3. Select a file or directory by typing this command:

   ```
   recover> add file_name
   ```

4. Repeat step 2 and step 3 for each file or directory to be included in the directed recovery.

5. View the selected files or directory by typing this command:

   ```
   recover> list
   ```

6. Start the directed recovery by typing this command:

   ```
   recover> recover
   ```

   - You can run the **recover** program from a Windows client to administer directed recoveries between Sun StorageTek EBS clients on UNIX, but only if both the source and destination clients are UNIX clients.
   - You cannot recover data from UNIX clients to non-UNIX clients, or the reverse.

Perform a directed save set recovery by using the command line

Perform directed save set recover by using the following command:

#**recover -s** *servername* **-R** *destination_client* **-iR -S** *SSID1* **-S** *SSID2*

Note: You can not recover save sets of different clients to a single client by using the same command line.

Recovering online indexes for a Sun StorageTek EBS server

This section includes procedures for recovering the client file index and the media database.

Recover a client file index

Client file index information is recovered by using the **nsrck** program.

To recover a client file index:

1. Log in as root or Windows Administrator and run the **nsrck** command:

   nsrck -L5 *client_name*

   The **-L5** option purges corrupted save set information from the index.

2. If the volume from which the client file index is being recovered is eligible for recycling, set the volume to read-only. This ensures that the volume will not be overwritten during the procedure:

   a. Verify whether the volume is eligible for recycling:

      – To query by date:
      mminfo -mv -c *client_name* **-t** *time*

      – To query by save set ID:
      mminfo -avot -S *ssid*

      mminfo -mv *volume_name*

   b. If the volume is eligible for recycling, type:

      nsrmm -o readonly *volume_name*

3. Run the **nsrck** command with the **-L7** option to recover the client file index:

   nsrck -L7 *client_name*

To recover the client file index from a specific time, replace this step with the instructions included in "Recovering a client file index from a specific time and date" on page 453.

Note: To completely remove a client file index, type **nsrck -R -Y** *client_name*, then type **nsrck -L7** *client_name* to recover the index.

Recovering a client file index from a specific time and date

Use the **-t** *time date* option with **nsrck** to recover client file indexes at a specific time. The time and date formats are in **nsr_getdate** format. For details on the **nsrck** and **nsr_getdate** commands, refer to the *EMC NetWorker Command Reference Guide* or the UNIX man pages.

To recover a client file index for a specific date, type this command:

nsrck -t *"time mm/dd/yyyy"* **-L7** *client_name*

Example 33 Recovering a client file index from a specific time and date

This command recovers a client file index from 4 P.M., March 2, 2009:

nsrck -t "04 pm 03/02/2009" -L7 *client_name*

When you recover a client file index from a time and date in the past, **nsrck** adds the entire contents of the index from that time and date to the current client file index contents. The client file index becomes much larger. Subsequent backups of the client file index will also include the added indexes unless they are explicitly removed by using the **nsrmm** command with the **-d -P** options. For more information about the **nsrmm** command, refer to the *EMC NetWorker Command Reference Guide* or the UNIX man pages.

Restore a save set entry in the online indexes

These sections describe the process of restoring save set entries into either the client file index only, or into both the client file index and the media database.

The procedures in this section use the **scanner** command. When the **scanner** command is used to restore entries to the client file index or the media database, the entries assume the browse and retention policies 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.

Using the scanner command efficiently

In NetWorker 7.5 and higher, when a volume is marked with the scan needed flag, the user can specify the starting file and record numbers that needs to be scanned from the media database. For example:

`scanner -f` *`file`* `-r` *record* `-i` *device*

Use these options when a whole volume scan is not needed. This shortens the time required to run the **scanner** program. For more information, refer to the *EMC NetWorker Command Reference Guide* or the UNIX man pages.

Note: The -i option is not supported for cloud devices.

Restoring a save set entry to the client file index only

If a file is not browsable (which means that the save set's browse policy has expired), but its save set is still tracked by the Sun StorageTek EBS server in the media database (which means that the save set's retention policy has not expired).

To recover the save set's entry back into the client file index:

1. Log in as root or Windows Administrator.
2. Type the **mminfo** command at the command prompt:

   mminfo -a -v *volume_name*

   where *volume_name* is the volume that contains the client file index.
3. From the **mminfo** output, find the *ssid* that contains the required file. Make sure it is not the bootstrap *ssid*.
4. Type the **nsrmm** command with a new browse time:

   nsrmm -S *ssid* **-w** *browse_time*

   where *browse_time* is a time in the future.
5. Type the **scanner** command:

   scanner -i -S *ssid*

   The save set entry is restored to the client file index.

Restoring a save set entry to the client file index and media database

If the save set that contains the file is not browsable and the save set does not appear in the media database, both the browse and retention policies have expired.

To rebuild the save set's entry in the client file index and the media database:

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. Locate *all* the volumes that contain this save set ID.

4. If the save set is found on multiple volumes and the order in which the volumes were written is unknown, complete this step. Otherwise, skip to step 5 .

   a. On each volume that will be reintroduced into the online indexes, run this command:

      scanner -m *device_name*

      The media database is updated with information from each volume.

      Note: If the volume contains data from an earlier version of Sun StorageTek EBS, 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. If the volume already belongs to a pool, the **-b** option will have no effect.

      ---

   b. Query the media database to determine the sequence in which the volumes were written:

      mminfo -a -v

      In step 5 , the volumes must be reintroduced into the online indexes in the same order in which they were written.

5. Starting with the first volume that was written, run the **scanner** program:

   scanner -i *device_name*

   The **scanner** program prompts for a new volume until you terminate it. Scan in the remaining volumes in the order in which they were written.

6. Browse for the file to be recovered. "Recovering by file selection" on page 425 provides information about browsing for a file.

Recovering the Console database

The Console database contains management data such as report information.

The *EMC NetWorker Disaster Recovery Guide* provides information about recovering a corrupted Console database.

Recovering backed-up client files from an old Sun StorageTek EBS server

This section describes how to move a Sun StorageTek EBS client to a new Sun StorageTek EBS server without losing the ability to recover the client files that were backed up on the old Sun StorageTek EBS server. To move a client to a new Sun StorageTek EBS server:

1. Record the **Client ID** attribute of the Sun StorageTek EBS 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**.

      The **Properties** dialog box appears.
   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** dialog box.
2. On the new Sun StorageTek EBS 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. If 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 this name:

      ~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. 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 Sun StorageTek EBS server.

   You can now recover data that was backed up when the Sun StorageTek EBS client was set up on the old Sun StorageTek EBS server.

Note: If **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. If 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 **scanner -i** again after creating the **Client** resource to import save set information into the client file index.

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 entered in the Save Set attribute of the Client resource, these SYSTEM save sets are backed up if no VSS client license exists or VSS is disabled:

- SYSTEM STATE
- SYSTEM FILES
- SYSTEM DB
- SHAREPOINT (only if installed on the client to be backed up)
- ASR

Note: Non-VSS save sets are not supported with Microsoft Windows Vista or Windows Server 2008.

These VSS SYSTEM save sets are backed up if VSS is licensed and enabled. On Windows Vista and Windows Server 2008, VSS save sets are enabled automatically and no VSS license is required:

- VSS SYSTEM BOOT
- VSS SYSTEM FILESET
- VSS SYSTEM SERVICES
- VSS USER DATA
- VSS OTHER
- VSS ASR DISK (Windows Server 2003 only)

If VSS is licensed and enabled, at *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.

Note: 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, reboot the operating system. Additionally, reboot 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 rebooting after each operation. If you attempt to restore the VSS SYSTEM BOOT save set a second time without rebooting after the first restore, this error message will appear:

```
regsow:cannot replace Registry key. Access is denied
```

If 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.

Note: Sun StorageTek EBS software does not determine the Windows operating system version during recovery of the SYSTEM or VSS SYSTEM save sets. If 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.

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 check the size of the SYSTEM or VSS SYSTEM save sets to be restored:

- For NetWorker clients running a 32-bit version of Windows 2003, approximately 500MB of extra disk space is usually sufficient.
- For NetWorker clients running a 64-bit version of Windows 2003, approximately 1GB of extra disk space is usually sufficient.
- For NetWorker clients running a 32-bit or 64-bit version of Windows 2008, approximately 10GB 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 rebooting the client computer.

Note: In Windows 2008 and Windows Vista, the BOOT / BCD (Boot Configuration Data)files that were in %systemdrive%\boot and %systemdrive%\bootmgr file are recovered to a temporary location under %systemdrive%\boot_restored-timestamp. A Windows Administrator can use these files to import BCD settings in a disaster recovery scenario.

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 Sun StorageTek EBS 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 -I 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:

*Sun StorageTek EBS_install_path***bin****recover.exe -iY** [**-s** *Sun StorageTek EBS_server_name*] **-N** "*saveset_name*"

2. Reboot the computer. The original SYSTEM files will be replaced by restored files.

Note: 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
   - VSS OTHER
   - VSS SYSTEM BOOT

   The command that is used to recover each save set should look similar to this:

   *Sun StorageTek EBS_install_path***bin****recover.exe -iY** [**-s** *Sun StorageTek EBS_server_name*] **-N** "*saveset_name*"
2. Reboot the computer. The original VSS SYSTEM files will be replaced by restored files.

Note: If 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 Sun StorageTek EBS 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
- VSS OTHER

Note: 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 the following at the command prompt:

   ```
   Sun StorageTek EBS_install_path\bin\recover.exe -iY
   [-s Sun StorageTek EBS_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:

   ```
   C:\Program Files\Legato\nsr\bin\recover.exe -iY
   -s jupiter -N "VSS SYSTEM SERVICES:\Event Log
   Writer;WMI Writer"
   ```

Note: 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) backup 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 **Sun StorageTek EBS Console Administration** window, click **Media**.
2. In the expanded left pane, select **Volumes**.
3. In the right pane, right-click one of the volumes listed and select **Show Save Sets**.
4. In the **SSID** column, note the appropriate save set ID.
5. Type this command on the Sun StorageTek EBS client host to restore the system state or system database to a particular point in time:

   ```
   Sun StorageTek EBS_install_path\bin\recover.exe -iY
   [-s Sun StorageTek EBS_server_name] -S SSID
   ```

where *SSID* is the save set ID that was noted in step 4 .

Note: If 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 Sun StorageTek EBS log file.

Point-in-time recovery of Microsoft SQL Server or Exchange Server from earlier versions of NetWorker

This version of NetWorker can perform a point-in-time recovery of a Microsoft Exchange Server or Microsoft SQL Server database that was backed up with the NetWorker Client release 7.2.x.

In NetWorker release 7.3 or later, 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.

Performing a point-in-time recovery of an Exchange Server database

If VSS is licensed and enabled, a point-in-time recovery of a Microsoft Exchange Server database can be performed.

Note: These steps are not required if the Sun StorageTek EBS Module for Microsoft Exchange Server is being used.

To perform a point-in-time recovery of an Exchange server database:

1. Take the Microsoft Exchange Server databases offline.
2. Go into the storage group and delete all the *.log and *.chk files from the Exchange data directories.
3. Complete the recovery.
4. Bring the Exchange Server databases back online.

Performing a point-in-time recovery of a SQL Server database

If VSS is licensed and enabled, a point-in-time recovery of a Microsoft SQL Server database can be performed.

To perform a point-in-time recovery of a SQL Server Database:

1. Take the Microsoft SQL databases offline.
2. Run the NetWorker **User** program, and in the left pane of the **Recover** window expand **VSS SYSTEM SERVICES**.
3. Mark the **MSDE Writer** folder.
4. Complete the recovery.
5. Bring the SQL Server databases back online.

Preparing to recover the Windows SYSTEM STATE save set

Table 54 on page 465 describes components of the SYSTEM STATE save set that require special preparation before being recovered.

Table 54 Preparing to recover the SYSTEM STATE save Set (1 of 2)

Component	Recover preparation
Active Directory (if installed)	1. When the computer 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 Sun StorageTek EBS server. "Preparing to recover the system state on a domain controller" on page 999 provides information about these settings.
Certificate Server (if installed)	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 1004 provides more information.
Cluster Server (if installed)	Shut down the Cluster Service on any nodes in the cluster on which the service is running , except for the node on which the recover will be performed. To shut down the cluster service, performing *one* of the following: • 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.

Table 54 Preparing to recover the SYSTEM STATE save Set (2 of 2)

Component	Recover preparation
SYSVOL (if installed)	None
Internet Information Server	None
Performance Counters	None

Preparing to recover the SYSTEM DB save set

Table 55 on page 466 describes components of the SYSTEM DB save set that require special preparation prior to recovery.

Table 55 Preparing to recover the SYSTEM DB save set

Component	Recovery Preparation
Disk Quota Database	• 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 running.
WMI	None

Recovering Windows volume mount points

Volume mount points are an NTFS file system feature of Windows Server 2003, Windows XP Professional, and Windows 2000. 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.

Note: The Sun StorageTek EBS 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 23, "Support for Microsoft 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. Start the NetWorker **User** program.
2. Recover just the mount point, without its data. To recover a mount point on drive D:\ for example:
   a. Click **Recover** to open the **Source Client** dialog box.
   b. Select the source client whose data you want to recover, and click **OK**. The local client is the default selection.
   c. Select the destination client for the recovered data, and click **OK**. The local client is the default selection.
   d. In the **Recover** window, expand the D:\ drive.
   e. Mark *mountpoint_name* for recovery.

Note: When you mark a mount point for recovery, all files and directories beneath it are marked by default. Be sure to clear them before starting the recovery. A mount point may appear as a directory (folder icon) rather than a mount point (volume icon).

f. Click **Start** to run the recovery.

3. Recover the mounted volume's data:

   a. Click **Recover** to open the **Source Client** dialog box.

   b. Select the source client with the data to be recovered, and click **OK**. The local client is the default selection.

   c. Select the destination client for the recovered data, and click **OK**. The local client is the default selection.

   d. In the **Recover** window, expand drive D:\ and *mountpoint_name*.

   e. Mark the data beneath *mountpoint_name* for recovery.

   f. Click **Start** to run the recovery.

"Recovering by file selection" on page 425 provides more information about performing data recoveries.

Recovering nested mount points

To recover nested mount points, first recover the top-level mount point, then work down the hierarchy, performing successive operations to recover the content of each mount volume (files, directories, and nested mount points).

Note: Recover nested mount points and their data in the order shown in this example.

This procedure explains how to recover three nested mount points and their data in the following configuration:

D:\
mountpoint_name1
data1
mountpoint_name2
data2
mountpoint_name3
data3

To recover this mount point tree:

1. Start the NetWorker **User** program.
2. Recover the top-level mount point:
   a. Click **Recover** to open the **Source Client** dialog box.
   b. Select the source client whose data is to be recovered, and click **OK**. The local client is the default selection.
   c. Select the destination client for the recovered data, and click **OK**. The local client is the default selection.
   d. In the **Recover** window, expand the D:\ drive.
   e. Mark **mountpoint_name1** for recovery.

   Note: When you mark a mount point for recovery, all files and directories beneath it are marked by default. Clear these files and directories before starting the recovery.
   At this point, the mount point may appear as a directory rather than a mount point (that is, it might have a folder icon instead of a volume icon).

   f. Click **Start** to begin the recovery.
3. Recover **data1** and **mountpoint_name2**. Clear **data2** and **data3** before starting the recovery.
4. Recover **data2** and **mountpoint_name3**. Clear **data3** before starting the recovery.
5. Recover **data3**.

Recovering Windows DHCP and WINS databases

If VSS is licensed and enabled, the DHCP and WINS databases are automatically included when performing a backup of save set All, Consequently, these procedures are not required.

Recover a DHCP database

To recover the DHCP database:

1. Using the NetWorker **User** program, recover the %SystemRoot% \System32\dhcp directory.
2. Using Microsoft DHCP administrative tools, 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 a WINS database" on page 106 provides more information.
2. Use Microsoft WINS administrative tools to restore the **WINS** database.

Note: For detailed instructions about Microsoft WINS administrative tools, refer to the Microsoft documentation.

Restoring Windows Content Index Server on Windows 2000 or later

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

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.

Note: If 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:

- If 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.

If 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.

14

Reporting

This chapter covers these topics:

- Gathering report data 474
- Data retention and expiration policies 475
- Setting expiration policies for data retention 477
- Report categories 477
- Report types 478
- Configuring reports 481
- Viewing reports 484
- Preconfigured reports 495
- Customizing and saving reports 531
- Sharing reports 533
- Exporting reports 534
- Command line reporting 535
- Printing reports 537

Gathering report data

To facilitate trend analysis, capacity planning, and problem detection, Sun StorageTek EBS software automatically collects data on a continual basis from the Sun StorageTek EBS enterprise. Sun StorageTek EBS Console stores the collected information in the Console database for a specified number of days, as described in "Data retention and expiration policies" on page 475.

Sun StorageTek EBS 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 477 provides detailed information about the various types of reports. The EMC Dantz® Retrospect documentation provides information on Retrospect events.

The following options are available through the Sun StorageTek EBS Console reporting feature:

- Data collection for the entire enterprise or for specific Sun StorageTek EBS or Retrospect 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 Sun StorageTek EBS 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.

Enable the gathering of report data

The Gather Reporting Data feature is set by default when a host is added to the enterprise. If the Sun StorageTek EBS or Retrospect server to be monitored is not yet in the enterprise, you can enable the Gather Reporting Data feature when adding the Sun StorageTek EBS or Retrospect server to the enterprise. "Adding a Sun StorageTek EBS, Retrospect, or Avamar host" on page 591 provides details.

To enable the gathering of report data:

1. From the **Console** window, click **Enterprise**.
2. Select the **Sun StorageTek EBS** or **Retrospect** server for which the collection of report information is to be enabled.
3. Right-click the Sun StorageTek EBS managed application, then select **Properties**. The **Managed Application Properties** dialog box appears.
4. Under **Features**, select **Gather Reporting Data**, then click **OK**.

Disable the gathering of report data

To disable the gathering of report data:

1. From the **Console** window, click **Enterprise**.
2. Select the **Sun StorageTek EBS** or **Retrospect** server for which the collection of report information is to be enabled.
3. Right-click the Sun StorageTek EBS managed application, then select **Properties**. The **Managed Application Properties** dialog box appears.
4. Under **Features**, clear the **Gather Reporting Data** checkbox, then click **OK**.

Data retention and expiration policies

NetWorker Console provides separate expiration polices 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 476. 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, *except* in the save set output). Retention policy for completion data can affect multiple reports.	Savegroup and save set completion data and drive data.	One month
Completion Message (in Backup Status reports, *only* 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
Save Set Data (in all Backup Statistics reports). Retention policy for save set data can affect multiple reports.	Save set records.	One year
Audit Data (in User Audit reports). Retention policy for audit data affects only Audit reports.	Reports on all Sun StorageTek EBS tasks (except License Manager tasks) performed by specified users (but only when the Sun StorageTek EBS User Auditing system option is activated).	One year

Users 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 477.

Note: Reports not mentioned in Table 56 on page 476 have no retention policies.

These different policies give administrators the flexibility to retain certain types of information for less time than others, as demonstrated in Example 34 on page 476.

Example 34 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 Sun StorageTek EBS 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 Sun StorageTek EBS 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 *Sun StorageTek EBS Installation Guide* provides information about estimating the size of the Console database .

Report categories

Table 57 on page 478 describes the various report categories included in Sun StorageTek EBS software.

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
Sun StorageTek EBS 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. "Sun StorageTek EBS backup statistics reports" on page 495 provides detailed information about reports.
Sun StorageTek EBS Backup Status	Provide status information about savegroup completion and save set backups. "Sun StorageTek EBS backup status reports" on page 502 provides detailed information about the reports in this category.
Events	Provide summary and detailed information about Sun StorageTek EBS events. "Event reports" on page 517 provides detailed information about the reports in this category.
Hosts	Provide a listing of Sun StorageTek EBS servers in the Enterprise, including information about event and reporting features. "Host reports" on page 521 provides detailed information about the reports in this category.
Users	Provide lists of defined Sun StorageTek EBS Console users, logout and login reports, audit reports, and users with restricted views. "User reports" on page 522 provides detailed information about the reports in this category.
Devices	Provides information about the way devices are being used. "Device reports" on page 524 provides detailed information about the reports in this category.
Inactive Files Report	Manages inactive files on a client or group and sets the Sun StorageTek EBS software to automatically generate a list of inactive files in an environment. "Inactive files report" on page 529 provides detailed information about the reports in this category.

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. Basic reports can be selected from within the Sun StorageTek EBS Backup Statistics, Sun StorageTek EBS or Retrospect Backup Status, Events, Hosts, and

Users report categories. "Basic reports" on page 479 provides a detailed description.

Drill-down reports are available within the Sun StorageTek EBS Backup Statistics, Sun StorageTek EBS Backup Status, Events, and Devices report categories. "Drill-down reports" on page 480 provides a detailed description.

For information on:

- Retrospect events and error messages, refer to the Dantz Retrospect documentation.
- Managing Retrospect servers. "Managing various servers in the enterprise" on page 590 provides details.

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 Sun StorageTek EBS 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

The reports that come with Sun StorageTek EBS software are known as "canned report"s, and include 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 531 and "Sharing reports" on page 533 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 Sun StorageTek EBS 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 482 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.

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 Sun StorageTek EBS software vary. Some acceptable input formats for a collection of common locales are shown in Table 59 on page 483.

Table 59 Date and time input formats for common locales

Language	Date formats	Time formats
US English	• EEEE, MMMM D, YYYY (Monday, March 8, 2009) • MMMM D, YYYY (March 8, 2009) • MMM D, YYYY (Mar 8, 2009) • M/D/YY (3/8/07)	• h:mm:ss a z (11:27:30 P.M. PST) • h:mm:ss a (11:27:30 P.M.) • h:mm a (11:27 A.M.)
UK English	• DD MMMM YYYY 08 March 2009) • DD-MMM-YYYY (08-Mar-2009) • DD/MM/YY (08/03/07)	• HH:mm:ss z (23:27:30 PST) • HH:mm:ss (23:27:30) • HH:mm (23:27)
French	• EEEE D MMMM YYYY (lundi 8 mars 2009) • D MMMM YYYY (8 mars 2009) • D MMM YYYY (8 mar. 2009) • DD/MM/YY (08/03/07)	• HH:mm:ss z (23:27:30 PST) • HH:mm:ss (23:27:30) • HH:mm (23:27)
German	• EEEE, D. MMMM YYYY (Montag, 8. März 2009) • D. MMMM YYYY (8. März 2009) • DD.MM.YYYY (08.03.2009DD) • MM.YY (08.03.07)	• HH:mm:ss z (23:27:30 PST) • HH:mm:ss (23:27:30) • HH:mm (23:27)
Japanese	• YYYY/MM/DD (2009/03/08) • YY/MM/DD (07/03/08)	• HH:mm:ss z (23:27:30 JST) • HH:mm:ss (23:27:30) • HH:mm (23:27)
Simplified Chinese	• YYYY-M-D (2009-3-8) • YY-M-D (07-03-8)	• HH:mm:ss (23:27:30)

In Table 59 on page 483, 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 include:

- Hour
- Day
- Week
- Month
- 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

Reports can be displayed in either interactive or document mode. Depending on the report, you may also elect to display the content of the report as a table or as a chart.

In addition, reports can be printed and exported to various file formats. "Printing reports" on page 537 and "Exporting reports" on page 534 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.

Note: 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.

Sun StorageTek EBS supports these report modes:

- "Interactive mode" on page 485
- "Document mode" on page 487
- "Interactive and document mode chart types" on page 489

In addition, Sun StorageTek EBS includes these restrictions and processing considerations:

- "Restricting report views" on page 494
- "Background processing of reports" on page 495

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 and tables

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 and charts

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 490.

  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 to display a tool tip that includes this information:
  - Save Set Name
  - Start Time
  - End Time
  - Client Name
  - Throughput (B/Sec)

    In Drive View, the tool tip includes information:

    – Drive
    – Start Time
    – End Time
    – 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. Sun StorageTek EBS 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 Sun StorageTek EBS 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

Report data displayed as a table in document mode

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.
- One or more *data* columns represent report information that can be summed as quantitative data.

The control columns show information such as:

- Server name
- Group name
- Client name
- Save set name
- Save time
- Save set ID
- Backup type
- Backup level

In charts, these topics are generally shown as *X*-axis data.

The data columns show information such as:

- Amount of data (in MB)

- Number of files
- Number of save sets
- Duration (in seconds)

In charts, these topics are generally shown as *Y*-axis data.

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

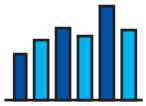
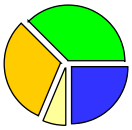
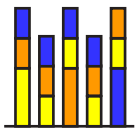
These chart types are available in both interactive and document mode:

- Bar chart
- Pie chart
- Plot chart
- Stacking bar chart
- Gantt chart (for Drive Utilization reports only)

Note: When viewing a Drive Utilization report as a chart, it automatically displays as a Gantt chart, and the chart type cannot be changed. "Types of Device reports and configuration" on page 524 provides more information. The Gantt chart type is not an option for other reports.

Table 60 on page 490 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 Sun StorageTek EBS 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 35 on page 492 describes appropriate use of the stacking bar chart type.

Example 35 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 Sun StorageTek EBS 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 Sun StorageTek EBS 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 Sun StorageTek EBS 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.

View reports

Before displaying a report, select the scope of the report by specifying Configuration parameters. "How to configure a report" on page 482 provides information on configuring 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.

Restricting report views

When a Sun StorageTek EBS Console user is added or reconfigured, the user's views of Sun StorageTek EBS 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 Sun StorageTek EBS 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 Sun StorageTek EBS 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.

Note: 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, Sun StorageTek EBS software includes a variety of reports that can be viewed and output in a variety of ways.

Sun StorageTek EBS backup statistics reports

The different types of reports included within the Sun StorageTek EBS Backup Statistics report category provide backup statistics for each selected Sun StorageTek EBS server within the enterprise.

Sun StorageTek EBS 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 Sun StorageTek EBS backup statistics reports and configuration

The Sun StorageTek EBS 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 Sun StorageTek EBS Backup Statistics report category are described in Table 61 on page 496. The specific parameters available depend on which Sun StorageTek EBS Backup Statistics report is selected.

Table 61 Sun StorageTek EBS 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.	• BOOTSTRAP • INDEX • FILES • DB2 • INFORMIX • MSEXCHANGE • MSSQL • NOTES • NMEO • SAP • SYBASE
Level	Select one or more backup levels.	• Full • Incremental • Skip • Level 1-9 (Partial list of options)
Save Time	Limits the report to a specified time range. **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 Sun StorageTek EBS Backup Statistics report category are listed in Table 62 on page 498 and Table 63 on page 504.

Save set data retention policy and configuration

Settings for the save set retention policy impact the data that is available to the Sun StorageTek EBS Backup Statistics reports.

Impact of save set data retention policy

If a save set retention policy of six months is specified, Sun StorageTek EBS software cannot query the database for a time range that extends back more than six months. The report cannot display data that has expired, because that data has been removed from the database. Thus, even if a save time parameter of one year is specified, the report can display only six months of data if the limit of the save set retention policy is six months.

Backup statistics basic reports

Within the Sun StorageTek EBS Backup Statistics report category, choose any of the basic reports described in Table 62 on page 498. 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 482 provides information on selecting and removing parameters .

Note: Except for the Deduplication Backup Summary report, these basic reports do not distinguish between regular and deduplication clients.

Note: If a time range for Save Time is not specified, all available data for that report is displayed.

Table 62 Sun StorageTek EBS Backup Statistics Basic Reports (1 of 3)

Report name	Purpose	Configuration parameters
Server Summary	For all or specified servers, provides statistics on the amount of data, number of files, and number of save sets backed up. The Server Summary report focuses on the highest level of the Sun StorageTek EBS datazone.	• Server Name • Save Time • Backup Type • Backup Level
Group Summary	For all or specified groups, provides statistics on the amount of data, number of files, and number of save sets backed up. The Group Summary report focuses on the savegroups. It provides a perspective on the differences between various savegroups in the enterprise or on specified servers.	• Server Name • Group Name • Save Time • Backup Type • Backup Level
Client Summary	For all or specified clients, provides statistics on the amount of data, number of files, and number of save sets backed up. The Client Summary report focuses on the clients of a particular server and savegroup or of a number of servers and savegroups. It provides a perspective on the differences between clients in the enterprise or on a specified server.	• Server Name • Group Name • Client Name • Save Time • Backup Type • Backup Level
Save Set Summary	For all or specified save sets, it provides statistics on the amount of data, number of files, and number of save sets backed up. The Save Set Summary report allows exploration of save set backup dynamics by viewing save set data for a particular client, server, and group, or for a number of clients, servers, and groups simultaneously.	• Server Name • Group Name • Client Name • Save Time • Save Set Name • Backup Type • Backup Level
Level Summary	Indicates how many full, incremental, skip, or level 1-9 backups have been completed for the scope defined by the configuration parameters. The Level Summary report can help determine whether the backup schedules are well balanced.	• Server Name • Group Name • Client Name • Save Time • Save Set Name • Backup Level • Backup Type

Table 62 Sun StorageTek EBS Backup Statistics Basic Reports (2 of 3)

Report name	Purpose	Configuration parameters
Backup Type Summary	Indicates how many of the backup types listed in Table 61 on page 496 have been completed for the scope defined by the configuration parameters.	• Server Name • Group Name • Client Name • Save Time • Save Set Name • Backup Type • Backup Level
Monthly Summary	Displays statistics on a month-to-month basis.	• Server Name • Group Name • Client Name • Save Time • Save Set Name • Backup Type • Backup Level
Daily Summary	Displays statistics on a day-to-day basis.	• Server Name • Group Name • Client Name • Save Time • Save Set Name • Backup Type • Backup Level
Group Summary by Server	Similar to a Server Summary report, it shows backup statistics in a server-centric manner. The Group Summary by Server report also shows the breakdown of statistics by savegroup within each server.	• Server Name • Group Name • Save Time • Backup Type • Backup Level

Table 62 Sun StorageTek EBS Backup Statistics Basic Reports (3 of 3)

Report name	Purpose	Configuration parameters
Group Summary by Client and Server	Similar to a Client Summary report, it shows backup statistics in a client-centric manner. The Group Summary by Client and Server report also identifies the group and server to which each client belongs.	• Server Name • Group Name • Client Name • Save Time • Backup Type • Backup Level
Save Set Details	Displays details about each save set, including backup duration. The Save Set Details and Save Set Details by Client reports: • Provide backup duration statistics. • Include the Save Set Name configuration parameter.	• Server Name • Group Name • Client Name • Save Time • Save Set Name • Backup Type • Backup Level
Save Set Details by Client	Similar to the Save Set Details report, but also identifies the client, group, and server to which each save set belongs. Only the Save Set Details and Save Set Details by Client reports: • Provide backup duration statistics. • Include the Save Set Name configuration parameter.	• Server Name • Group Name • Client Name • Save Time • Save Set Name • Backup Type • Backup Level

Backup statistics drill-down reports

The drill-down reports consist of multiple Sun StorageTek EBS Backup Statistics basic reports, connected in a predetermined sequence. Table 63 on page 504 lists the available Backup Statistics drill-down reports.

"Drill-down reports" on page 480 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 482 provides information on selecting and removing parameters.

Sun StorageTek EBS backup status reports

The Sun StorageTek EBS Backup Status reports consolidate information about the success of scheduled group backups. As with the Sun StorageTek EBS Backup Statistics reports, these reports can provide either an enterprise-wide or a more focused summary of activity over a specified time range.

The Sun StorageTek EBS Backup Status reports provide the same basic function as selecting Show Details for a group in the Monitoring window of the Administration window. The Sun StorageTek EBS Backup Status reports, however, allow you to select the scope and level of detail.

Note: 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.

Sun StorageTek EBS 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 Sun StorageTek EBS Backup Statistics and Sun StorageTek EBS 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 Sun StorageTek EBS modules (such the Module for Oracle, Module for EMC, 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 Sun StorageTek EBS software to identify the backup type and include it in the reports. The Sun StorageTek EBS Module for Oracle (NMO), however, does not include an identifying prefix and therefore NMO is not included in the list of backup types. As a result, the NMO files are handled under the normal file type.

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)

A backup command is always used for database backups (it is sometimes used also for regular file backups). A system administrator who has created scripts for backing up the database files should be able to scan the backup commands and recognize the script name used for backing up Oracle files. This method can be used as a workaround for distinguishing the NMO backup type.

Types of Sun StorageTek EBS backup status reports and configuration

The Sun StorageTek EBS 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 Sun StorageTek EBS Backup Status Report category are described in Table 63 on page 504.

Table 63 Sun StorageTek EBS backup status parameters

Parameter	Description	Options
Server Name	Selects one or more Sun StorageTek EBS 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.	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.	• BOOTSTRAP • DB2 • FILE • INDEX • INFORMIX • Lotus Notes • MSEXCHANGE • MSSQL • NMEO • ORACLE - RMAN • SAP • SYBASE
Level	Selects one or more backup levels.	• Full • Incremental • Skip • Level 1-9 (Partial list of options)
Status	Selects status.	• Successful • Failed • Interrupted

The parameters available for each report type are listed in Table 64 on page 505 and Table 65 on page 508.

Completion data retention and Sun StorageTek EBS backup status

The settings for the completion data policy impact the data that is available to the Sun StorageTek EBS 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 Sun StorageTek EBS Backup Status report category, choose any of the basic reports listed in Table 64 on page 505. 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 482 provides information on selecting and removing parameters.

Note: If no time range is specified for Group Start Time, all available data for that report is displayed.

Table 64 Sun StorageTek EBS backup status basic reports (1 of 3)

Report name	Purpose	Configuration parameters
Server Summary	Provides a report on the status of savegroup backups for all or specified servers. The Server Summary report focuses on the highest level of the Sun StorageTek EBS datazone. Use this report to: • Compare server workloads. • Troubleshoot problems associated with backups.	• Server Name • Status • Group Start Time
Group Summary	Provides a report on the status of backups for all or specified groups, including backup duration and success ratio. The Group Summary report focuses on the savegroups. Use this report to: • Compare backup status of the different savegroups on one server. (The Group Summary by Server report also serves this purpose.) • Check the backup status of the Default Savegroup on multiple servers in order to look for differences between them.	• Server Name • Group Name • Status • Group Start Time
Monthly Summary	Shows status on a month-to-month basis.	• Server Name • Group Name • Status • Group Start Time

Table 64 Sun StorageTek EBS backup status basic reports (2 of 3)

Report name	Purpose	Configuration parameters
Daily Summary	Shows status on a day-to-day basis.	• Server Name • Group Name • Group Start Time • Status
Group Summary by Server	Similar to a Server Summary report, it provides status in a server-centric manner. The Group Summary by Server report also shows status for each group on a server. Use this report to: • See how successful different savegroups are. • Troubleshoot problems associated with backups.	• Server Name • Group Name • Group Start Time • Status
Group Details	Lists the group start time, backup duration, and status for each savegroup backup.	• Group Name • Server Name • Group Start Time • Status
Group Summary by Client and Server	Provides backup status information for all or specified clients. The Group Summary by Client and Server report focuses on either of the following: • The clients of a particular server and savegroup. • A number of servers and savegroups. The report provides a perspective on the differences between clients in the enterprise or on a specified server.	• Server Name • Group Name • Client Name • Group Start Time • Status • Backup Type • Backup Level
Daily Save Set Summary	Combines a focus on the time period with a focus on save sets. Use this report to get a perspective on the success of different save sets.	• Server Name • Group Name • Client Name • Save Set Name • Status • Backup Type • Backup Level • Group Start Time

Table 64 Sun StorageTek EBS backup status basic reports (3 of 3)

Report name	Purpose	Configuration parameters
Save Set Details	Provides additional information about individual save set backups, including: • Save set ID • Save type • Save set start time • Backup level	• Server Name • Group Name • Client Name • Group Start Time • Save Set Name • Status • Backup Type • Backup Level
Save Set Details by Client	Similar to the Save Set Details report, but also identifies the client, group, and server to which each save set belongs. Use this report to get detailed information about each save set and its status for a particular client.	• Server Name • Group Name • Client Name • Group Start Time • Save Set Name • Status • Backup Type • Backup Level
Save Set Output	Shows messages associated with a particular save set backup. **Note:** A Save Set Output report can be generated without specifying any parameters. Doing so could, however, make it difficult to determine which messages come from which save set.	• Server Name • Group Name • Client Name • Save Set Name • Group Start Time

Backup status drill-down reports

The drill-down reports are comprised of multiple Sun StorageTek EBS Backup Status basic reports, connected in a predetermined sequence. For a list of available Backup Status drill-down reports, see Table 65 on page 508.

"Drill-down reports" on page 480 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 482 provides information on selecting and removing parameters.

Table 65 Sun StorageTek EBS backup status drill-down reports

Report name	Purpose	Sequence	Configuration parameters
Daily Backup Status	Shows daily backup status of different datazone components for comparison.	1. Daily Summary 2. Server Summary 3. Group Details 4. Save Set Details by Client 5. Save Set Output	• Server Name • Group Name • Group Start Time • Status
Group Status over Time	Combines a focus on the time period with a focus on savegroups. You can: • Compare status of servers and groups during the specified month. • Identify trends over time for groups and servers.	1. Monthly Summary 2. Group Summary by Server	• Server Name • Group Name • Group Start Time • Status
Group Status by Server	Allows you to view datazone status reports and compare them to date-centric views.	1. Server Summary 2. Group Summary 3. Monthly Summary 4. Daily Summary	• Server Name • Status • Group Start Time
Client Status over Time	Allows you to view client backup status and then navigate to greater levels of detail.	1. Group Summary by Client and Server 2. Daily Save Set Summary 3. Save Set Details 4. Save Set Output	• Server Name • Group Name • Client Name • Group Start Time • Status • Backup Type • Backup Level

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. There are four different types of reports that can be generated from the recover statistics:

- Server Summary
- Client Summary
- Recover Details
- Recover Summary Over Time

Note: 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.

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.

Types of Sun StorageTek EBS recovery reports and configuration

The Sun StorageTek EBS recovery report category includes basic and drill-down reports. The different types of reports included within the Sun StorageTek EBS Recover Statistics report category provide recover statistics for each selected Sun StorageTek EBS server within the enterprise.

The Configuration tab allows you to limit the scope of the report that was selected.

The parameters available within the Sun StorageTek EBS Recovery report category are described in Table 66 on page 510. The specific

parameters available depend on which Sun StorageTek EBS Recovery Statistics report is selected.

Table 66 Sun StorageTek EBS 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	• Successful • Failed

The parameters available for each report type in the Sun StorageTek EBS Recovery Statistics report category are listed in Table 67 on page 511.

Recovery Statistics basic reports

Within the Sun StorageTek EBS Recovery Statistics report category, choose any of the basic reports described in Table 67 on page 511. 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 482 provides information on selecting and removing parameters.

Note: If a time range for Save Time is not specified, all available data for that report is displayed.

Table 67 Sun StorageTek EBS Recovery Statistics basic reports

Report name	Purpose	Configuration parameters
Server Summary	For all or specified servers, provides statistics on the recoveries that were performed by NetWorker servers, including the amount of data, number of files, and recovery status (successful or failed).	• Server Name • Recover Status • Start time
Client Summary	For all or specified clients, provides statistics on the recoveries that were performed for specific clients, including the amount of data, number of files, and recovery status (successful or failed).	• Server Name • Client Name • Recover Status • Start time
Recover Details	For all or specified servers and clients, allows you to obtain more details regarding recovery operations.	• Server Name • Source Client • Recover Status • Start time

Server Summary Report

The Server Summary Report lets you obtain information about recoveries that were performed by NetWorker servers. To generate this report, specify the following information:

- The name of one or more NetWorker servers
- The desired time range for the recovery operations
- The Completion status of the recovery (Success, Failure, or both)

When the report is generated, the following data is displayed:

- Server name
- Size (amount of data)

- Number of files
- Total Requests
- Status (“Successful” or “Failed”)

Client Summary Report

The Client Summary Report lets you obtain information about recoveries that were performed for specific clients. To generate this report, specify the following information:

- The name of one or more NetWorker servers
- The name of one or more source clients
- The desired time range for the recovery operations
- The Completion status of the recovery (Success, Failure, or both)

When the report is generated, the following data is displayed:

- Source client name
- Size (amount of data)
- Number of files
- Total Requests
- Status (“Successful” or “Failed”)

Recover Details Report

The Recover Details Report lets you obtain more detailed information about recoveries that were performed by a NetWorker server. To generate this report, specify the following information:

- The name of one NetWorker server
- The name of one source client
- The desired time range for the recovery operations
- The Completion status of the recovery (Success, Failure, or both)

When the report is generated, the following data is displayed:

- Source client
- Target Client
- Initiating client
- User
- Size
- Number of files

- Start time
- Duration
- Status

Note: This report is available in Table mode only.

Recovery Statistics drill-down report

This drill-down report consists of multiple Sun StorageTek EBS Recovery Statistics basic reports, connected in a predetermined sequence. "Drill-down reports" on page 480 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 482 provides information on selecting and removing parameters.

Table 68 Sun StorageTek EBS Recovery Statistics drill-down reports

Table 5.

Report name	Purpose	Sequence	Configuration parameters
Summary over Time	Allows you to view server recovery status and then navigate to greater levels of detail.	1. Server Summary 2. Client Summary 3. Recover Summary	• Server Name • Start time • Recover Status

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. The following drill-down sequence is used for the reports:

- Server Summary
- Client Summary

- Recover details

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.

Deduplication Statistics reports

The Sun StorageTek EBS Deduplication Statistics reports, available from the **Reports** task pane in the Console window, provide deduplication backup statistics for each selected Sun StorageTek EBS 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
- Deduplication Backup Summary

Types of Sun StorageTek EBS Deduplication Statistics reports and configuration

The Sun StorageTek EBS Deduplication 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 Sun StorageTek EBS Deduplication Statistics report category are described in Table 69 on page 515. The specific parameters available depend on which Sun StorageTek EBS Deduplication Statistics report is selected.

Table 69 Sun StorageTek EBS Deduplication 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. **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.	• Full • Incremental • Skip • Level 1-9 (Partial list of options)

The parameters available for each report type in the Sun StorageTek EBS Backup Statistics report category are listed in Table 70 on page 516.

Deduplication Statistics basic reports

Within the Sun StorageTek EBS Deduplication Statistics report category, choose any of the basic reports described in Table 70 on page 516. 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 482 provides information on selecting and removing parameters.

Note: If a time range for Save Time is not specified, all available data for that report is displayed.

Table 70 Sun StorageTek EBS Deduplication Statistics basic reports

Report name	Purpose	Configuration parameters
Client Summary	For all or specified clients, provides deduplication statistics on the following: • Number of Save Sets: The number of save sets included in the backup. • Number of Files: The number of files included in the backup. • NetWorker size: The amount of the data that would have been moved using traditional NetWorker backup. • Protected size: The size of the last full Avamar backup. • Data moved: The size of the actual data moved during the backup. • NetWorker metadata size: The size of the NetWorker metadata	• Server Name • Client Name • Group Name • Save Set Name • Save Time • Backup Level
Save Set Summary	For all or specified save sets, provides deduplication statistics on the following: • Number of Save Sets: The number of save sets included in the backup. • Number of Files: The number of files included in the backup. • NetWorker size: The amount of the data that would have been moved using traditional NetWorker backup. • Protected size: The size of the last full Avamar backup. • Data moved: The Size of the actual data moved during the backup. • NetWorker metadata size: The size of the NetWorker metadata	• Server Name • Client Name • Group Name • Save Set Name • Save Time • Backup Level
Save Set Details	Displays details about each save set, including backup duration. The Save Set Details report provides statistics on the following: • Save Set ID • Save time • Backup level • NetWorker size: The amount of the data that would have been moved using traditional NetWorker backup. • Protected size: The size of the last full Avamar backup. • Data moved: The Size of the actual data moved during the backup. • Deduplication Ratio: Percentage of savings by using deduplication. • NetWorker metadata size: The size of the NetWorker metadata • Number of Files: The number of files included in the backup.	• Server Name • Client Name • Group Name • Save Set Name • Save Time • Backup Level

Deduplication Statistics drill-down report

The drill-down report, Deduplication Backup Summary, consists of multiple Sun StorageTek EBS Deduplication Statistics basic reports, connected in a predetermined sequence. "Drill-down reports" on page 480 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 482 provides information on selecting and removing parameters.

Event reports

These reports provide summary information about current events on Sun StorageTek EBS, Retrospect, 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 Sun StorageTek EBS Console displays the current events of the Sun StorageTek EBS and Retrospect 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 EMC Retrospect documentation provides information on Retrospect events. The EMC Avamar documentation, Chapter 26, "Support for Avamar Deduplication" and "Dismissing an event" on page 545 provide 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 71 on page 518. 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 71 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 Sun StorageTek EBS Retrospect Avamar
Priority	Selects only priority events. Priority represents the relative severity of the event. Table 81 on page 543 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 Table 72 on page 519. 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 482 provides information about selecting and removing parameters.

Table 72 Event basic reports

Report name	Purpose	Configuration parameters
Server Summary	For all or specified Sun StorageTek EBS or Retrospect servers, it provides a summary report on the active events. This report does *not* specify priority or category. It lists only the number of events per server.	• Server Name • Server Type
Priority Summary	Provides a summary of currently active events according to their priority. The report can be limited to show only selected priorities. Table 81 on page 543 provides a list of event priorities.	• Priority • Server Type
Category Summary	Provides a summary of currently active events according to their category. The report can be limited to show only selected categories.	• Category • Server Type
Event Details	Provides additional information about each currently active event, including: • Error message • Note (if one exists) • Check mark (indicating whether there is an annotation) For information on Retrospect events and error messages, refer to the Dantz Retrospect documentation.	• Server Name • Server Type • Priority • Category
Annotation Details	Provides the actual text of the Sun StorageTek EBS Console annotation. "Working with annotations" on page 544 provides information about annotations. **Note:** An Annotation Detail report can be generated without specifying any parameters. This could, however, make it difficult to determine which annotation is associated with which event.	• Server Name • Server Type • Event Time • Priority • Category

Additional information

These resources provide more information about the contents of Event reports:

- "Working with notes" on page 543 provides information on Sun StorageTek EBS Console notes.

- The *EMC NetWorker Error Message Guide* provides descriptions of Sun StorageTek EBS software error messages and troubleshooting procedures.
- Chapter 29, "Troubleshooting" provides information on troubleshooting Sun StorageTek EBS software issues.
- The Dantz Retrospect documentation provides information on Retrospect events and error messages.
- "Managing various servers in the enterprise" on page 590 provides information on administering Retrospect through Console.

Event drill-down reports

The drill-down reports consist of multiple Event basic reports, connected in a predetermined sequence. Table 73 on page 521 provides a list of available Event drill-down reports. "Drill-down reports" on page 480 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 482 provides information on selecting and removing parameters.

Table 73 Event drill-down reports

Report name	Purpose	Sequence	Configuration parameters
Events by Server	Provides a summary report on the currently active events that are running on the server, and then allows you to drill down to increasing levels of detail.	1. Server Summary 2. Event Details 3. Annotation Details	• Server Name • Server Type
Events by Priority	Provides a summary of currently active events according to their priority, and then allows you to drill down to increasing levels of detail.	1. Priority Summary 2. Event Details 3. Annotation Details	• Priority • Server Type
Events by Category	Provides a summary of currently active events according to their category, and then allows you to drill down to increasing levels of detail.	1. Category Summary 2. Event Details 3. Annotation Details	• Category • Server Type

Host reports

The Hosts report category includes only basic reports, no drill-down reports. There are two basic reports, as described in Table 74 on page 521.

Table 74 Host reports

Report name	Purpose	Configuration parameters	Default
Host List	Provides an overview of servers in the enterprise, including: • Whether it is a Sun StorageTek EBS or Retrospect server. • Whether the Capture Events feature is enabled for the server. • Whether the Gather Report Data feature is enabled for the server. • Where the server is located in the enterprise path.	None	All servers
Enterprise Inventory	Allows movement through the Enterprise. Limit the report's scope by first viewing one of the lower-level folders within the Enterprise: • Start from Enterprise folder. • Start from selected folder.	Enterprise Path	Start from Enterprise folder

"Enterprise" on page 588 provides a description of the Enterprise and its folders.

User reports

The Users report category provides information on Sun StorageTek EBS Console user activity. Chapter 16, "Console Server Management" provides information about Sun StorageTek EBS Console users and creating user accounts.

The Users report category includes only basic reports, no drill-down reports. Select any of the basic reports listed in Table 75 on page 523.

Note: The **Full Name** and **Description** information appears in the User reports only if it was included when the user was created.

Table 75 User reports(1 of 2)

Report name	Purpose	Configuration parameters
User List	Lists the user IDs for all users that have been given access to the Sun StorageTek EBS Console. Report includes this information for each user: • Username • Full name (if available) • Description (if available) • Last login time • Last logout time	None
Users by Last Login Time	Lists the last login time for all Sun StorageTek EBS Console users, starting with the most recent login. Report includes this information for each user: • Last login time • Last logout time • Username • Full name (if available) • Description (if available) Report also includes a listing of all logins available in the records.	Last Login Time (range)
Users by Last	Lists the last logout time for all Sun StorageTek EBS Console users, starting with the most recent logout. Report includes this information for each user: • Last logout time • Last login time • Username • Full name (if available) • Description (if available) Report also includes a listing of all logouts available in the records.	Last Logout Time (range)

Table 75 User reports(2 of 2)

Report name	Purpose	Configuration parameters
Users Currently Logged In	Lists the usernames and login/logout times of all Sun StorageTek EBS Console users that are currently logged in. Report includes this information for each user: • Username • Full name (if available) • Description (if available) • Last login time • Last logout time	None
User Audit	Allows the administrator to audit (track) a user's Sun StorageTek EBS server activities that were performed by using the Sun StorageTek EBS Console. For example, the audit can identify a user who has been labeling (overwriting) tapes that should have been archived. The information tracked comes from the Console database. **Note:** Ensure that User Audit is enabled in the System Options dialog box in the Administration window.	• Event Time • Server Type • Server Name • User Name • Category • Operation • Object Type • Object Name
User Host Permissions	Lists the access permission for selected users on selected computers. Report includes this information for each user: • Username • Hostname	• User Name • Hostname

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 Sun StorageTek EBS 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. Table 76 on page 526 describes the Drive Utilization report.

The Drive Utilization Report, which is a drill-down report, supports Sun StorageTek EBS servers running Sun StorageTek EBS 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.

Table 76 Device reports

Report name	Purpose	Sequence	Configuration parameters
Drive Utilization	Starts with a summary report on the backup activity level of one or more devices for a particular time period, and then allows you to drill down to see the save set activity level on a particular device at a particular time. **Note:** You cannot specify a time range that exceeds eight days. That is, the date entered in the To field cannot exceed eight days from the date entered in the From field. If typing a relative time in the To field, the value cannot exceed 8 days. If a save set starts or ends outside the specified range, the report will display this data.	1. Drive View 2. Save Set View	• Server Name • Drive • Time (range)

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 480 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 77 on page 527. The specific

parameters available depends on which Cloud backup and recover report is selected.

Table 77 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. **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 described in Table 78 on page 528. 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 482 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 489 provides general information on charts.

Table 78 Cloud backup and recover reports(1 of 2)

Report name	Purpose	Configuration parameters
Device Backup Summary	The Device Backup Summary report displays the scheduled backup history for each Cloud storage device and the login account.	• Server name • Device name • Start time
Backup Details	The Backup Details report displays historic information of all scheduled backups to the Cloud.	• Server name • Device name • Start time
Backup Summary over Time	The Backup Summary over Time drill-down report displays the history of backups performed by the NetWorker Server to the Cloud over a period of time.	• Server name • Device name • Start time

Table 78 Cloud backup and recover reports(2 of 2)

Report name	Purpose	Configuration parameters
Device Recover Summary	The Device Recover Summary report displays recovery information that is performed by the NetWorker Server for each Cloud storage device. **Note:** 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.	• Server name • Device name • Start time
Recover Details	The Recover Details report displays recovery information that is performed by the NetWorker Server for each Cloud storage device.	• Server name • Device name • Start time
Recover Summary over Time	The Recover Summary over Time report is a drill-down report that displays the history of recoveries from the Cloud performed by the NetWorker Server over a period of time.	• Server name • Device name • Start time

Inactive files report

A Sun StorageTek EBS administrator can manage inactive files on a client or group and set the Sun StorageTek EBS 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 Sun StorageTek EBS 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 Sun StorageTek EBS 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 Sun StorageTek EBS 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 Sun StorageTek EBS servers. Only servers that have the Gather Reporting Data attribute turned on will appear in the selection list.
- One or more Sun StorageTek EBS groups for the selected Sun StorageTek EBS servers.

Client File Details

The Client File Details report provides information about inactive files backed up for selected Sun StorageTek EBS clients. Data will be provided for every requested Sun StorageTek EBS 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 Sun StorageTek EBS servers. Only servers that have the Gather Reporting Data attribute turned on will appear in the selection list.
- One or more Sun StorageTek EBS groups for the selected Sun StorageTek EBS servers.
- One or more Sun StorageTek EBS clients for the selected Sun StorageTek EBS servers.

Customizing and saving reports

Canned reports can be changed and then saved under different names. Users can preserve the report configuration parameters that are most useful for the enterprise.

A customized report is a changed copy of a canned report. 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 options:

- Delete
- Rename
- Save
- Save As...
- Share
- Print
- Export

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 Sun StorageTek EBS reporting purposes, the terms *customized* report and *saved* report are basically 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 Sun StorageTek EBS Console. This also makes them accessible by the command line reporting feature. "Command line reporting" on page 535 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 494 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 formats and shared with others. Table 79 on page 534 lists the formats into which reports can be exported.

Table 79 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. **Note:** For raw data only.

Exporting a report

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 (HTML, PDF, CSV, or PostScript).

   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.

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 Sun StorageTek EBS 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 1.5 or later to run the **gstclreport** command. The JRE must be installed before installing Sun StorageTek EBS software.

Additionally, you must add an environment variable named JAVA_HOME to your NetWorker server host. Set the JAVA_HOME environment variable to the location of your JRE. Your operating system documentation provides information about setting up environment variables.

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.

15 Monitoring

This chapter covers these topics:

- Overview of monitoring 540
- Events..... 540
- Monitoring window 545

Overview of monitoring

Sun StorageTek EBS Console provides the ability to view details of current Sun StorageTek EBS and Retrospect server activities, operations related to devices and libraries, and events that require user intervention. It also provides the ability to view system-health events from Avamar servers. It makes administration of servers more efficient by providing a centralized means of monitoring all activity throughout an enterprise. "Managing various servers in the enterprise" on page 590 provides details on adding hosts to be monitored.

Events

An event signals that user intervention is required. For example, if a Sun StorageTek EBS server needs a new tape, the server alerts users to the situation by posting an event to the Console window.

Sun StorageTek EBS software generates an event based on various factors, including:

- The software or hardware encounters an error that requires user intervention to resolve.
- A Sun StorageTek EBS savegroup has failed.
- Capacity monitoring (for example, reaching the space threshold on the deduplication node).
- Sun StorageTek EBS 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 Sun StorageTek EBS Console (instead of by the License Manager) approaches its expiration date. In this situation, a message is recorded in the Sun StorageTek EBS 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)
- Sun StorageTek EBS activities (in seconds)
- Sun StorageTek EBS libraries (in hours)
- Avamar events (in hours)

Note: Event polling for Sun StorageTek EBS libraries and deduplication events can only be done to a maximum of once per hour. "Console system options" on page 576 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 Sun StorageTek EBS software can monitor that server for events. This option is selected by default when a host is added.

To disable or reenable 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**, then click **OK**.
   - To disable captured events, select **Features,** clear the **Capture Events** checkbox, then click **OK**.

For Avamar servers, the **Capture Events** option monitors only system-level events. The Avamar documentation provides other event information.

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 80 on page 542 describes the various columns and the information they provide for Sun StorageTek EBS events.

The Dantz Retrospect documentation provides details on Retrospect events and error messages.

Table 80 Events columns

Column	Description
Priority	Represents the relative severity of the problem by displaying one of seven icons. Table 81 on page 543 describes each priority.
Server Name	Identifies the host that caused the event to be generated.
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 482 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 544 provides more information.
Note	Provides an editable field for making brief notes associated with an event. "Working with notes" on page 543 provides more information.

Event priorities

Each event is designated one of seven possible priorities. Table 81 on page 543 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 81 Event priorities

Icon	Priority	Description
	Alert	Severe error condition that demands immediate attention.
	Critical	Error condition detected by the Sun StorageTek EBS and Retrospect server that should be fixed by a qualified operator.
	Emergency	Condition exists that may cause Sun StorageTek EBS and Retrospect software to fail unless corrected immediately. This icon represents the high*est* priority.
	Information	Information about the current state of the server. This icon represents the *lowest* priority.
	Notification	Important information.
	Waiting	Indication that the Sun StorageTek EBS and Retrospect 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 Sun StorageTek EBS 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 Sun StorageTek EBS 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 Sun StorageTek EBS or Retrospect software are automatically dismissed in the Console window when the problem that triggered the event is resolved.
- 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.

Monitoring window

Open the Monitoring window from the Administration window taskbar to view details of current Sun StorageTek EBS server activities, such as:

- Running of automatic and manual savegroups
- Archiving, cloning, recovering, and browsing of client file indexes

The Monitoring window also displays alerts and log messages, as well as operations related to devices and jukeboxes. While the Monitoring window is used primarily to monitor Sun StorageTek EBS server activities, it can be used to perform certain operations, such as starting, stopping, or restarting a group backup.

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.

Viewing the Monitoring window

To view the full Monitoring window:

1. From the **Administration** window, click **Monitoring**. The Monitoring window includes a docking panel that displays specific types of information.
2. Select the types of information you want to view from the docking panel.

A smaller window appears within the Monitoring window for each type you select. This 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 82 describes the various types of information available in the docking panel and the details each one provides.

Table 82 Monitoring window panel (1 of 2)

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.
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.
Alerts	Lists the priority, category, time, and message of any alerts.
Devices	Lists devices, device status, storage nodes, libraries, volumes, pools, and related messages.

Table 82 Monitoring window panel (2 of 2)

Window	Information provided
Operations	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. The Operations window will display library data only for servers running Sun StorageTek EBS software release 7.3 or later. **Note:** 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.
Log	Lists messages generated by the Sun StorageTek EBS server, including the priority of each message, the time the message was generated, the source of the message, and the category.
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.

Viewing backup group information

The Groups window displays groups that are in the process of completing, or have completed, their backup. Use this window to identify which groups backed up successfully and which groups failed.

The backup of a client group may fail for one of these reasons:

- The Sun StorageTek EBS server failed.
- The Sun StorageTek EBS client failed.
- The network connection failed.

To find out more about a backup failure, check Group Backup details. "Viewing group backup details" on page 549 provides more information.

Groups window

The backup status is represented by an icon. Table 83 on page 548 lists and describes each of the icons.

Table 83 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.

It is reported when the group is in a probing state. A message is sent when a group starts and finishes probing, and the results of the probing (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.



IMPORTANT

Be sure to check the results of the probe in the log window to ensure that the probes indicated the backup be taken.

Group 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.

Note: When a group backup is started manually, the Sun StorageTek EBS 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 Sun StorageTek EBS server immediately backs up the clients in the group, overriding the scheduled backup start time. The group icon changes to the clock icon until the backup has completed or is interrupted.

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 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.

Viewing session information

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

- Save
- Recover
- Clone
- Browse
- All

The default setting for the Sessions window is to display Save sessions. "Changing displayed session types" on page 550 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 Sun StorageTek EBS 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.

Note: The ability to change the type of session being displayed is available only for Sun StorageTek EBS release 7.3 and higher.

Other panel information

This section describes the other panel information available in the Monitoring window. These include:

- "Alerts window" on page 551
- "Devices window" on page 552
- "Operations window" on page 554
- "View operation details" on page 554 .

Viewing alerts

The Alerts window displays alerts generated by a particular Sun StorageTek EBS server. It includes priority, category, time, and message information.

Alerts window

The alert priority is represented by an icon. Table 84 on page 551 lists and describes each of the icons.

Table 84 Alerts window icons

Icon	Label	Description
	Alert	Severe error condition that demands immediate attention.
	Critical	Error condition detected by the Sun StorageTek EBS server that should be fixed by a qualified operator.
	Emergency	Condition exists that could cause Sun StorageTek EBS software to fail unless corrected immediately. This icon represents the *highest* priority.
	Information	Information about the current state of the server. This icon represents the *lowest* priority.
	Notification	Important information.
	Waiting	The Sun StorageTek EBS 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.

Monitoring the status of devices

The Devices window allows you to monitor the status of all devices, including NDMP devices. If the Sun StorageTek EBS 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.

Devices window

The device status is represented by an icon. Table 85 on page 552 lists and describes each of the icons.

Table 85 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.

Monitoring device operations

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 555 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 Sun StorageTek EBS or Retrospect server.
- Time the operation started.
- Type of operation.
- Duration of the operation.
- Status messages from the operation.
- Any error messages.

Note: 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.

Operations window

The operation status is represented by an icon. Table 86 on page 554 lists and describes each of the icons.

Table 86 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.

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

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.

Viewing log files from the console

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

Note: If a particular log file is no longer available, check the log file on the Sun StorageTek EBS server. The log files are located in this directory: <Sun StorageTek EBS_install_path>\logs. "Viewing log files" on page 904 provides information about viewing log files.

Log window

The log priority is represented by an icon. Table 87 on page 556 lists and describes each of the icons.

Table 87 Log window icons

Icon	Label	Description
	Alert	Severe error condition that demands immediate attention.
	Critical	Error condition detected by the Sun StorageTek EBS server that should be fixed by a qualified operator.
	Emergency	Condition exists that could cause Sun StorageTek EBS software to fail unless corrected immediately. This icon represents the *highest* priority.
	Information	Information about the current state of the server. This icon represents the *lowest* priority.
	Notification	Important information.
	Waiting	The Sun StorageTek EBS 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.

Viewing archive information

The Archive Requests window displays the current status of all archive requests that are scheduled on the Sun StorageTek EBS 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.

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 278.

Archive requests window

The archive status is represented by an icon. Table 88 on page 558 lists and describes each of the icons.

Table 88 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.

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.

16

Console Server Management

This chapter covers these topics:

- Managing Sun StorageTek EBS Console server users..... 562
- Moving the Console server..... 572
- Setting system options..... 575
- Setting environment variables 580
- Accessing the Console Configuration Wizard 582
- Sun StorageTek EBS Console server maintenance tasks 582
- Displaying international fonts in non-US locale environments 585

Managing Sun StorageTek EBS Console server users

The NetWorker Console (Console) server restricts user privileges based on the three roles described in Table 89 on page 562. These roles cannot be deleted and their privileges cannot be changed.

Table 89 Console roles

User Role	Privileges
Console Security Administrator	• Add, delete, and edit users • Configure login 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.

When Console is first opened, the default administrator user is assigned to all three Console user roles.

Login authentication for Console users

Console server supports two types of login authentication:

- Native Console authentication
- External authentication to an LDAP v3 compliant server

Native Console authentication is enabled by default and requires that user names and passwords be maintained on the Console server. No additional set up is required to enable native Console login authentication.

LDAP authentication enables you to log in to the Console server with user names and passwords that are maintained on a centralized LDAP v3 compliant server such as a Microsoft Active Directory

server. Console user privileges are controlled by mapping LDAP user roles or LDAP user names to Console user roles. There is no need to add user names and passwords on the Console server. However, one must still add LDAP user names to user groups on the NetWorker server to enable privileges on the NetWorker server.

Note: You must use either native Console authentication or LDAP authentication. To switch from one mode to the other, use the Configure Logon Authentication wizard.

Enabling LDAP login authentication

To enable LDAP authentication:

1. Log in to the Console server as a user, such as the default administrator, who belongs to the Console Security Administrator role.
2. On each NetWorker server, add an external LDAP user to the NetWorker server Administrator's user group. This step ensures that once LDAP is enabled, at least one user will be able manage the NetWorker server and to add additional NetWorker users as required.

   The LDAP user that you add should also belong to the LDAP user roles or LDAP user names that you specify later in <ln Link>step 7.

   a. Click the **Enterprise** button on the taskbar.

   b. Highlight a host in the navigation tree, right-click **NetWorker**, and select **Launch Application**. The Administration window appears.

   c. Click the **Configuration** button on the taskbar.

   d. In the navigation tree, select **User Groups**.

   e. In the **User Groups** list, right-click **Administrators** and select **Properties**.

   f. Add the LDAP user to the **User** attribute. Use the following format to add the user:

      ```
      User=LDAP_username, host=console_host
      ```

      where *console_host* is the name of the Console server host.

   g. Click **OK**.

3. From the **Console Setup** menu, select **Configure Login Authentication** to launch the **Configure Login Authentication** wizard.

   The **Select Authentication Method** panel appears.

4. Select the **External Repository** radio button and click **Next**. The **Manage Authentication Authorities** panel appears.

5. Click **Add** and then provide information about your authentication authority in the remaining fields. For help on a field, click the question mark (**?**) at the bottom left side of the wizard panel.

6. Click **Next** when finished. The **Setup Console Security Administrator** panel appears.

7. Enter the LDAP user roles or LDAP user names that will be mapped to the Console Security Administrator role and click **Finish**.

   Note: The LDAP user that was added to the NetWorker Server Administrator's list in <ln Link>step 2 must also be added to the Console Security Administrator role.

   You can now log in to the Console server using an LDAP user name and password that was mapped in this step.

Note: In LDAP mode, the user name *Administrator* is not allowed to log in even if it is defined in LDAP.

Post LDAP set up requirements

After LDAP authentication is enabled, perform the following tasks:

- Map additional LDAP user roles to Console user roles. "Adding Console users and assigning to Console roles" on page 565 provides more information.
- Optionally, restrict access from the Console server to managed servers. "Restricting a Console user's view of managed servers" on page 566 provides more information.
- Add LDAP user names to NetWorker server user groups. This task enables users to perform operations on the NetWorker server. "Granting Console users privileges on a NetWorker server" on page 568 provides more information.

Adding Console users and assigning to Console roles

This section describes how to add Console users and assign those users to Console roles. These procedures differ depending on whether native NetWorker login authentication or external LDAP login authentication is being used.

By default, Console users can view all managed NetWorker servers. However, Console users must still be granted explicit privileges on a NetWorker server. "Granting Console users privileges on a NetWorker server" on page 568 describes how to manage access to NetWorker servers.

Adding Console users when using native authentication

To add Console users when using native login authentication:

1. Log in to the Console server as 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 **General** tab displayed.
4. Enter a username.
   - Do not exceed 20 characters.
   - Do not use spaces, or any of these characters:: < > /
   - Do not use characters with an ASCII value less than or equal to 32.
   - Do not 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 to be given to the user. Table 89 on page 562 provides more information about Console user roles.
7. Enter the user password.

   Note: 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 edited users. For customers upgrading from a previous release, this requirement will be enforced when user passwords are changed.

8. Reenter the password in the **Confirm Password** attribute and the click **OK**.

Adding Console users when using LDAP authentication

When an LDAP user logs in for the first time, a user object is automatically created on the Console server. You need only 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.

To add Console users when using LDAP authentication:

1. (Optional) Create the Console user manually:
   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**, then select **New**. The Create User dialog box appears, with the Identity tab displayed.
   d. In the User Name attribute, enter the LDAP user name.
   e. Optionally, enter the full name of the LDAP user and a general description in the remaining attributes.
   f. Click **OK**.
2. Map LDAP users to Console roles.
   a. In the left pane, select **Roles**.
   b. In the right pane, right-click a Console user role and select **Properties**.
   c. In the External Roles attribute, add each LDAP user role or LDAP user name that is to be mapped. Add each entry on a separate line.
3. Click **OK**.

Restricting a Console user's view of managed servers

You can restrict Console users from viewing Sun StorageTek EBS servers by setting the permissions on the user object. "Implications of restricting a Console user's view of managed servers" on page 567 provides information about the implications of restricting user views, see . "Restricting report views" on page 494 provides information on how restricting views affects reporting.

To set permissions for a user to view servers:

1. Log in as Console Security Administrator.
2. From the **Console** window, click **Setup**.
3. In the left pane, click **Users**.
4. Right-click a username, then select **Permissions**. The **Edit User** dialog box appears, with the **Permissions** tab displayed.
5. To grant permission for the user to view various hosts, use the arrow keys to select the allowed hosts.
6. Click **OK**.

Note: Console users can manage data only on the servers they are allowed to view. However, to perform operations on the NetWorker server, Console users must also be granted explicit privileges on the NetWorker server. "Granting Console users privileges on a NetWorker server" on page 568 provides more information.

Implications of restricting a Console user's view of managed servers

Users can manage data, such as report data, only on servers they are allowed to view. However, even when they have permission to view a given server, the management operations they can perform on it are determined by their privileges on that server. "Granting Console users privileges on a NetWorker server" on page 568 provides information on how these privileges are set.

The effects of restricting user views for various functions are as follows:

- In the **Events** window: The user sees only events from allowed Sun StorageTek EBS servers. The Retrospect events will be displayed in the **Events** window if the user adds the Retrospect server host in the **Enterprise** window of the Console.
- In the **Enterprise** window: The user sees all of the hierarchy folders, but only the allowed Sun StorageTek EBS and Retrospect servers appear in those folders.
- In the **Libraries** window: The user sees only the devices controlled by allowed Sun StorageTek EBS and Retrospect servers.

- In the **Reports** window: The user sees report data only from allowed Sun StorageTek EBS servers, groups, clients and Retrospect servers. "Managing various servers in the enterprise" on page 590 provides details on managing Retrospect servers.
- In the **Setup** window:
  - The user sees properties for all users, and his or her own properties and permissions.
  - The user can edit his or her own properties, but not permissions. Only the Console Security Administrator can edit user permissions, and see and edit permissions.

Note: Because each user can view and manage different sets of Sun StorageTek EBS and Retrospect servers, the contents of the reports can vary among users. For example, a shared backup summary report entitled "Building C Backups" will show different data for different users (even if the report is run at the same time) if the users' permissions include different Sun StorageTek EBS servers. This applies to all report types, whether canned or customized, private or shared.
If no data is available for a given server, that server will not appear in any lists or reports, regardless of the user permissions.

"Restricting a Console user's view of managed servers" on page 566 provides information on how to restrict user views.

Granting Console users privileges on a NetWorker server

By default, each Console user must be granted explicit privileges on the NetWorker servers that they will manage.

Note: If you do not need to control NetWorker server privileges by individual user, you can disable this feature and give all users the same privileges that are granted to the Console server by default. To do so, clear the User Authentication for Sun StorageTek EBS system attribute, which is described in "Setting system options" on page 575.

To grant Console users privileges on a NetWorker server:

1. Log in to the Console server as a user who belongs to the Administrator's group on the NetWorker server to which you are adding users.

- If you are using native Console authentication, the default Console administrator user has administrator privileges on the NetWorker server.
- If you are using external LDAP authentication, the LDAP user that was added in <ln Link>step 2 of the section titled, "Enabling LDAP login authentication" on page 563, has administrator privileges on the NetWorker server.

2. Click the **Enterprise** button on the taskbar.
3. Highlight a host in the navigation tree, right-click NetWorker, and select **Launch Application**. The Administration window appears.
4. Click the **Configuration** button on the taskbar.
5. In the navigation tree, select **User Groups**.
6. In the User Groups list, right-click a user group and select **Properties**.
7. Add users to the User attribute. Use the following format to add a user:

   ```
   User=username, host=console_host
   ```

   where *console_host* is the name of the Console server host.

   Note: If LDAP authentication is enabled, the username entered in this step must be the LDAP user name and *not* the LDAP user role.

8. Click **OK**.

   Chapter 17, "Sun StorageTek EBS Server Management" provides more information about NetWorker server use groups and their privileges.

Resetting user preferences

An option is available in the Users task of the Console to reset a user's preferences to the settings as they were when the user first logged in.

To reset a user's preferences:

1. From the **Console** window, click **Setup**.
2. In the left pane, click **Users**.
3. Right-click a username, then select **Reset User Preferences**.

4. A dialog displays asking if you really want to remove all of their preferences before performing the reset. Click **Yes** to reset the preferences.

Deleting a user

Once a user has been deleted, that user can no longer log in to the Console server. Only the Sun StorageTek EBS administrator can delete other users.

To delete a user:

1. Log in as a Console Security Administrator.
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.
   - 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.
   - If you are using LDAP authentication, you must also remove the user from the LDAP user role, which is maintained on the LDAP server.

Editing user properties

To change a password or other descriptive information about a user, edit the user properties. All users can edit their own properties. Only a Sun StorageTek EBS administrator, however, can edit properties of other users.

To edit user properties:

1. From the **Console** window, click **Setup**.
2. In the left pane, click **Users**.
3. Right-click a username, then select **Properties**. The **Edit User** dialog box appears, with the **General** tab displayed.
4. Change the user information as necessary. Other user information is displayed, but is not editable, including last-login information.

Note: If you are using native Console authentication, you can also change the administrator password by using the Console Configuration Wizard. "Accessing the Console Configuration Wizard" on page 582 provides more information.

5. Click **OK**.

Resetting the administrator password (native authentication only)

If the default administrator password is lost or forgotten, it can be reset by using the GST_RESET_PW environment variable.

Note: These steps apply only when native Console authentication is being used.

Microsoft Windows

To reset the native Console security administrator password on Windows:

1. From the **Start** menu, go to **Control Panel** > **System** and set the **GST_RESET_PW** variable in the **Environment Variables** dialog box to a value of **1**.

   Note: How you navigate to the **Environment Variables** dialog box differs depending on which Windows operating system you use. Consult the documentation that accompanies your operating system for details.

2. Restart the **GST Service**. When the **GST Service** starts, the password resets.
3. Log in to the Console server and type **administrator** as the user name and type **administrator** as the password.
4. Return to the Environment Variables dialog box and delete the **GST_RESET_PW** variable. This step prevents the password from being reset at each Console log in attempt.

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 your shell, for example, in ksh:

   ```
   export GST_RESET_PW "non null value"
   ```

2. Stop and restart the Console server. "Stopping and starting the Console server" on page 69 provides details. This should be done in the same shell where the password was set.
3. Log in to the Console server and type **administrator** as the user name and type **administrator** as the password.
4. Set **GST_RESET_PW** back to **null** using the appropriate command for your shell, for example, in ksh:

```
export GST_RESET_PW=
```

The next time the Console server is restarted, the password will not be reset again.

Moving the Console server

The Console server of Sun StorageTek EBS can be moved from one computer to another only if both computers use the same operating system. Reasons to move the Console server component include:

- The current computer has insufficient processing capabilities. For example, if there is a need for more memory or a faster processor.
- The current computer has insufficient space for the Console database.
- The current computer was damaged beyond repair.

To move the Console server component to another computer:

1. Back up the existing Console database with the **savepsm** command. "Performing a manual backup of the Console database" on page 100 provides more details.
2. Set up the new computer with the same operating system as the one that is running the current software and is connected to the network.
3. On the new computer, install the Sun StorageTek EBS Client software and the Sun StorageTek EBS Console server component.
   - This step includes meeting the prerequisites for installing Sun StorageTek EBS software. For example, the Sun StorageTek EBS client must already be installed. If using NetWorker License Manager, the License Manager software must either already be installed or be available on another host. For information, consult the *Sun StorageTek EBS Installation Guide*.

- If you are using the License Manager and it is moved from one computer to another, enter the new License Manager hostname in the Console window. For information, refer to the *EMC License Manager Installation and Administration Guide*.

4. On the remote Sun StorageTek EBS server, set up the new target computer as a Client resource. "Task 6: Create a backup Client resource" on page 81 provides information about creating a Client resource.
5. For the Client resource of the source Console server, add the appropriate users to the **Remote Access** attribute on the **Globals (2 of 2)** tab.
6. Stop the Console server service on the source Console server. "Stopping and starting the Console server" on page 69 provides more information.
7. Stop the Console server service on the target Console server.
8. Ensure the GST service is not running on the original backup host or on the host being used for the directed recover. Use the appropriate command or tools for your operating system to ensure the GST service is not running.
9. For UNIX systems, set the appropriate library path environment variable to the following:

   Console_install_dir`/sybasa/lib64`
   Console_install_dir`/sybasa/lib (Linux)`

   The environment variable to set varies by platform:

   - Solaris/Linux: LD_LIBRARY_PATH
   - AIX: LIBPATH
   - HP-UX: SHLIB_PATH

   Note: If the Console is not installed in the default /opt/LGTOnmc directory on Solaris, add *Console_install_dir*/bin to the LD_LIBRARY_PATH environment variable.

10. On the Sun StorageTek EBS server, you must have:
    - Client resources for the original backup host and the directed recover host.
    - 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 new host using the **-c** option to identify the original host. Also use the **-O** option to recover only the database and not the gstd_db.conf credential file.

    recoverpsm [**-f**][**-d** *recover_directory*] **-s** *Sun StorageTek EBS_server* **-c** *original_Console_server* **-S** **gst_on**_*original_Console_server* **-0**

    Table 90 on page 575 provides a description of **recoverpsm** options. The *EMC NetWorker Command Reference Guide* or the UNIX man pages provide a complete description of the command and its options.

12. Use the **recover** program or the NetWorker **User** or **nwrecover** program to recover the gstd_db.conf configuration file to the new location. This will maintain your existing database login credentials.

    For NetWorker 7.5 and higher, the gstd_db.conf file is located in the save set named CONSOLE_BACKUP_FILES.

    Note: If you were using LDAP authentication, you must configure your LDAP authorities again on the new Console server using the Configure Login Authentication wizard.

13. If *recover_directory* is not the same as the Console database directory, copy the database file and the gstd_db.conf file, to the directory where the Console database resides.
14. Start the Console server service on the target Console server.
15. If you are using the License Manager and the License Manager host has changed, obtain a Host Transfer Affidavit from EMC support.

The host ID that is used for the authorization code is that of the License Manager computer. If the License Manager computer has not changed, then new authorization codes are unnecessary

Table 90 The recoverpsm options

Option	Meaning
-f	Instructs the software to overwrite existing Console database files.
-d *recover_directory*	Specifies the destination directory for the recovered Console database. Include a full path for the directory. Partial paths are not supported.
-S gst_on_*source_Console_server*	Specifies the existing Console database on the source Console server. You must prepend **gst_on**_ to the name of the source Console server.
-O	Omit the recovery of the database credential file.
-s StorageTek_EBS_*server*	Specifies the name of the Sun StorageTek EBS server where the Console database backup resides.
-c *source_Console_server*	Specifies the short name of the source Console server, where the existing Console database resides.

Setting system options

Sun StorageTek EBS Console includes several options that affect performance. These options enable users to fine-tune the performance of the Console server. Only Console application 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**. The **System Options** dialog box appears**.**
3. Set a value, or enable or disable the appropriate system option.
4. Click **OK**.



IMPORTANT

Do not adjust these system options without careful consideration. A mistake in setting system options or environment variables can seriously degrade performance.

Console gets two configuration options from the environment (GST_MAXLOGSIZE and GST_MAXLOGVERS). Normally the default values are used; these environment variables can be changed by advanced users. "Setting environment variables" on page 580 for more information provides more information.

Table 91 Console system options (1 of 3)

System option	Purpose
Log on Banner	Specifies the text to display in the user log on screen.
Debug level **Note:** Alternatively, GST_DEBUG can be set as an environment variable. This allows users to troubleshoot the Console if they cannot access the Console window. The environment variable overrides the system option version (Debug level). To avoid having it override the system option indefinitely, unset the environment variable after use, and set the system option for general use.	Set the level of debugging information written to the **gstd** log file. (Increasing this value increases the number of operation and status messages that Console records.) • To increase the number of program messages written to **gstd** log file, choose a larger value. • To use less disk space, choose a smaller value **Range:** 1 to 10.
Polling interval for events and reporting (secs)	Set the minimum number of seconds that must pass before Console checks a Sun StorageTek EBS server for events. • To make Console more responsive to new events, decrease the value. • To decrease the load on the CPU, increase the value. **Range:** >= 2 **Default:** 20
Polling interval for NetWorker activities (secs)	Set the minimum number of seconds that must pass before Console checks a Sun StorageTek EBS server for updates to the Administration windows. • To make the Administration windows more responsive to changes on the Sun StorageTek EBS server, decrease the value. • To decrease the load on the CPU, increase the value. **Range:** >= 2 **Default:** 10

Table 91 Console system options (2 of 3)

System option	Purpose
Polling thread factor	Determine the number of programming threads that are used when polling Sun StorageTek EBS servers. • To make Console more responsive, increase the value. • To decrease the load on the CPU, decrease the value. **Range:** 1 – 10 **Default:** 5
Polling interval for NetWorker libraries (hours)	Set the minimum number of hours that must elapse before the Console checks a Sun StorageTek EBS server for library information. • To poll the Sun StorageTek EBS server for library information more frequently, decrease the number of hours. • To poll the Sun StorageTek EBS server for library information less frequently, increase the number of hours. **Range:** Greater than or equal to 1. **Default:** 12
Polling interval for Avamar events (hours) **Note:** In order for Avamar events to appear in NMC, they must be set to **Acknowledge** in the Events Management tab of the Avamar Administrator program. Chapter 26, page 26-855 provides more information.	Set the minimum number of hours that must elapse before the Console checks for Avamar events. Avamar events can include: • **hfscheck** — validation of the data stored on the Avamar server. • Disk health/hardware — status of the disks on the Avamar server (limited to “bad” or “failed” disk status). • Replication events that occur on the Avamar server. Avamar events can be polled at hourly intervals: • To poll for Avamar events more frequently, decrease the number of hours. • To poll for Avamar events less frequently, increase the number of hours. **Range:** Greater than or equal to 1. **Default:** 12

Table 91 Console system options (3 of 3)

System option	Purpose
NetWorker user auditing	Enable auditing of user Console requests. • When enabled, this function logs all Console user requests to Sun StorageTek EBS servers. • When this function is disabled, Console user requests to Sun StorageTek EBS servers are not logged. **Range:** Not applicable **Default:** Disabled Audit records are written to the Console database, and the results are viewable in the Reports task of Console's main window. Chapter 14, page 14-473 provides for information about viewing reports.
User authentication for NetWorker	Set individual user authentication for functions a user can perform on a Sun StorageTek EBS server. **Note:** Requires restarting the Console server to take effect. This option restricts user access to a given Sun StorageTek EBS server *only* if the user groups on that server have been set up to do so. User groups can be created, modified, or deleted from the Administration window. • This option does not set restrictions on user views. • When this function is disabled, all Console users who are allowed to access a Sun StorageTek EBS server have the same privileges on that server. • When this function is enabled, access privileges on Console servers can vary among Console users. **Range:** Not applicable **Default:** Enabled
RPC ping via UDP when connecting to a NetWorker server	Enable the use of user datagram protocol (UDP) when Sun StorageTek EBS Console first checks whether a Sun StorageTek EBS server is up and running. Set this option if the firewall allows for UDP connections. Note that there is a significant trade-off in performance. **Range:** Applicable **Default:** Enabled

Individual User Authentication

Console security administrators can restrict or grant Console user access to Sun StorageTek EBS servers based on the Console username if the User Authentication for Sun StorageTek EBS system option has been enabled (with a subsequent restart of the Console server). This system option is enabled by default.

Requests to Sun StorageTek EBS servers through the Administration window always come from the Console server, regardless of any system option settings.

When the User Authentication for Sun StorageTek EBS system option is enabled, the requests to Sun StorageTek EBS servers appear (from the Sun StorageTek EBS side) to be coming from users on the Console server, rather than from the **gstd** process owner on the Console server.

When the system option is enabled, a Sun StorageTek EBS server allows requests only from users who belong to the Administrators group of the Sun StorageTek EBS server. The username of the Console daemon process owner must be included in the Sun StorageTek EBS Administrators group on Sun StorageTek EBS servers to which Console users have access. Sun StorageTek EBS administrators must also ensure that the usernames of Console users who have permission to view and manage those servers are included in the Sun StorageTek EBS server's Administrators group. Table 92 provides examples. "Administrator privileges" on page 610 provides detailed information about the Administrators group.

Note: The username of the root or system user on the Console server *must* be specified, regardless of whether individual user authentication is used.

Table 92 Usernames for individual user authentication

On Console server	Operating system	Example of username to add to the Sun StorageTek EBS Administrators group[1]
Root or system user on the Console server–the Console daemon process owner (mandatory)	Microsoft Windows	system@winbox.petropep.com
	UNIX	root@unixbox.petropep.com
Sun StorageTek EBS Administrator	Microsoft Windows	administrator@winbox.petropep.com
	UNIX	administrator@unixbox.petropep.com
User	Microsoft Windows	user1@winbox.petropep.com
	UNIX	user1@unixbox.petropep.com

[1] Where winbox.petropep.com and unixbox.petropep.com are Console servers.

Chapter 17, page 17-587 describes how to include the user names in the Sun StorageTek EBS Administrators group

Impact on network connections

More network connections might need to be available when individual user authentication is enabled. This could have an impact on firewall considerations. Appendix B, "Firewall Support" provides information about firewalls.

When the User Authentication for Sun StorageTek EBS system option is set, a separate network connection is created from the Console server to a Sun StorageTek EBS server for each Console user that has an Administration window open to that server. This option is set by default.

When the user authentication for Sun StorageTek EBS system option is *not* set, there is only one network connection from the Console server to a Sun StorageTek EBS server that is being managed through Sun StorageTek EBS.

Setting environment variables

Table 93 on page 581 describes the three environment variables available in this Sun StorageTek EBS release. Variables are checked only when the Console server service starts, which is similar to how the **nsrd** daemon handles its corresponding variables.

Note: Environment variable names are case-sensitive.

Table 93 Console environment variables

Purpose	Environment variable
Set the maximum size of the **gstd** log file before a new log file is created at Console server startup time. If this variable is undefined, the maximum size defaults to 1 MB. • To preserve more log information, increase the value. • To use less disk space, decrease the value. "The gstd log" on page 959 provides more information about the **gstd** log file. **Range:** >= 1 **Default:** 1 MB	GST_MAXLOGSIZE
Set the maximum number of historical versions of the **gstd** log file that Console retains. Older versions are deleted when the number of versions exceeds this value. • To preserve more log information, choose a larger value. • To use less disk space, choose a smaller value. "The gstd log" on page 959 provides more information about the **gstd** log file. **Range:** 1–999 **Default:** 4	GST_MAXLOGVERS
Set the level of debugging information written to the **gstd** log file. Setting this value higher increases the number of operation and status messages that Console records. GST_DEBUG can also be set as a system option (Debug level). • To increase the number of program messages written to gstd.log, choose a larger value. • To use less disk space, choose a smaller value. **Range:** 1–20 **Default:** 0	GST_DEBUG **Note:** This environment variable can be used to troubleshoot the product in cases where the user interface is not accessible. The environment variable overrides the system option version (Debug level). To avoid having it override the system option indefinitely, unset the environment variable after use, and set the system option for general use.

Setting environment variables

To set environment variables:

Note: Open the /etc/init.d/**gst** file in a text editor. (This file is a Bourne shell script.) By default, this is a read-only file. Check the file permissions and change them, if necessary, before editing it.

5. 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 to be set.
- *variable_value* is the value to be assigned to the environment variable.

6. Save the changes.
7. Stop and restart the Console server. "Stopping and starting the Console server" on page 69 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.

Accessing the Console Configuration Wizard

Note: Only Console Application Administrators can use the Console Configuration Wizard.

The following tasks can be performed from the Console menu options or from the Console Configuration Wizard.

- Set the Sun StorageTek EBS Administrator password.
- Specify the Sun StorageTek EBS server that will back up the Console database.
- Add Sun StorageTek EBS 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**.

Sun StorageTek EBS Console server maintenance tasks

This section lists maintenance tasks associated with the Sun StorageTek EBS Console server.

Changing the service port used by the Sun StorageTek EBS Console database

By default, the Sun StorageTek EBS Console database uses port 2638 for TCP/IP communications. Other applications may also use this port or, as is the case with EMC AlphaStor or EMC DPA installed with default options, may have their own instance of the iAnywhere database installed, which can create conflicts if both applications are installed on the same host.

To resolve this problem, change the service port as follows:

1. Stop the GST Service. "Stopping and starting the Console server" on page 69 has information about stopping the GST service.
2. Stop any other products that are using the iAnywhere database, for example, EMC AlphaStor or EMC DPA.
3. Ensure that no **dbsrv9** processes are running. If the **dbsrv9** processes are running, stop them.
4. Open a terminal or command prompt window.
5. Depending on the shell you are using, use the appropriate command (for example. **setenv** for csh, **export** for sh) to update the library path environment variable in the terminal window to the following location:
   - Solaris: **/opt/LGTOnmc/bin:/opt/LGTOnmc/sybasa/lib**
   - Linux/AIX/HPUX: **/opt/lgtonmc/bin:/opt/lgtonmc/sybasa/lib**
   - Microsoft Windows (assumes default installation location): **C:\Program Files\Legato\Management\GST\sybase\bin**

   The environment variable to set varies by platform, for example:
   - Solaris/Linux: LD_LIBRARY_PATH
   - AIX: LIBPATH
   - HP-UX: SHLIB_PATH
6. Edit the gstd.conf file to add or change the following line:

   db_svc_port=*port_number*

   For example:

   db_svc_port=2639

   The gstd.conf is located in the following locations:

- Solaris: /opt/LGTOnmc/etc
- Linux/AIX/HP-UX: /opt/lgtonmc/etc
- Microsoft Windows: C:\Program Files\Legato\Management\GST\etc

7. Run the gstconfig command to update the port value in the Sun StorageTek EBS Console. The gstconfig command is located in the following locations:
   - Solaris: /opt/LGTOnmc/bin
   - Linux/AIX/HP-UX: /opt/lgtonmc/bin
   - Windows: C:\Program Files\Legato\Management\GST\bin
8. Close the terminal or command prompt window.
9. Restart the GST Service.

Changing database connection credentials

When the Sun StorageTek EBS Console server is started for the first time, it automatically generates login credentials that it will use to log into the Sun StorageTek EBS Console database. This information is stored internally by the Sun StorageTek EBS Console server and the user does not need to know the required credentials. However, it may be necessary to force the Console server to change the database connection credentials.

To force the server to change the credentials:

1. Stop the GST Service. "Stopping and starting the Console server" on page 69 has information about stopping the GST service.
2. Set the environment variable GST_RESET_DBPWD to any value. On Microsoft Windows system, this should be set as a System Variable and the system should be rebooted after the variable is set.
3. Restart the GST Service.
4. Delete the GST_RESET_DBPWD environment variable. On Microsoft Windows system, the system should be rebooted after the variable is set.

NMC server IP address/hostname updates

If the IP address/hostname of the NMC server is modified or protocols such as IPv6 are added or removed, perform the following:

1. Stop the GST Service. "Stopping and starting the Console server" on page 69 has information about stopping the GST service.
2. Navigate to the NetWorker bin directory and run the platform-specific command:
   - On Windows, go to C:\Program Files\Legato\Management\GST\bin, and run **gstconfig**.
   - On Solaris, as root:
     LD_LIBRARY_PATH=/opt/LGTOnmc/bin:/opt/LGTOnmc/sybasa/lib export LD_LIBRARY_PATH
     ./gstconfig
   - On Linux, as root:
     LD_LIBRARY_PATH=/opt/lgtonmc/bin:/opt/lgtonmc/sybasa/lib
     export LD_LIBRARY_PATH
     ./gstconfig
   - On AIX:
     LIBPATH=/opt/lgtonmc/bin:/opt/lgtonmc/sybasa/lib
     export LIBPATH
     ./gstconfig
   - On HP-UX:
     SHLIB_PATH=/opt/lgtonmc/bin:/opt/lgtonmc/sybasa/lib
     export SHLIB_PATH
     ./gstconfig
3. Restart the GST service.

Displaying international fonts in non-US locale environments

To use or view data from a localized Sun StorageTek EBS server, ensure that the appropriate font is available to the Console server. The Sun StorEdge Enterprise Backup Software Installation Guide provides more information about displaying international fonts on a Console server that is operating in English mode.

17

Sun StorageTek EBS Server Management

This chapter covers these topics:

- Enterprise 588
- Configuring a Sun StorageTek EBS server 599
- Configuring ConnectEMC 600
- Enabling BMR server support 604
- Parallelism and multiplexing 604
- Managing server access 609
- Sun StorageTek EBS User Groups 613
- Server communication issues within Microsoft Windows 618
- Notifications 619
- Indexes 629
- Updating Sun StorageTek EBS clients by using the software distribution feature 641
- Message log management 641
- Internationalization 644

Enterprise

The Enterprise is a visual representation of the Sun StorageTek EBS Console control zone. Sun StorageTek EBS and Retrospect servers in the Enterprise can be monitored for events, and various reports can be generated on events, backups, and user activity. Chapter 14, "Reporting" provides detailed information about Console reports.

Enterprise components

The Enterprise includes these components:

- Hosts

  A host, also known as a *managed node*, is the Sun StorageTek EBS, Retrospect, or Avamar server being monitored. A host terminates a branch in the Enterprise.

  Note: EMC Dantz Retrospect hosts can be monitored 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 Sun StorageTek EBS servers

Use the Enterprise to organize the Sun StorageTek EBS 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 36 on page 589 provides information.
- By function of the computers backed up by the Sun StorageTek EBS 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 36 An enterprise arranged by geographic location

Figure 33 shows an Enterprise arranged by geographic location. There are three folders, one for each country where the Sun StorageTek EBS servers are located: USA, France, and Australia. Each folder contains a number of hosts that correspond to Sun StorageTek EBS servers named for the city where they are located. The Australia folder, for instance, contains three host computers labeled *perth1*, *perth2*, and *sydney*.

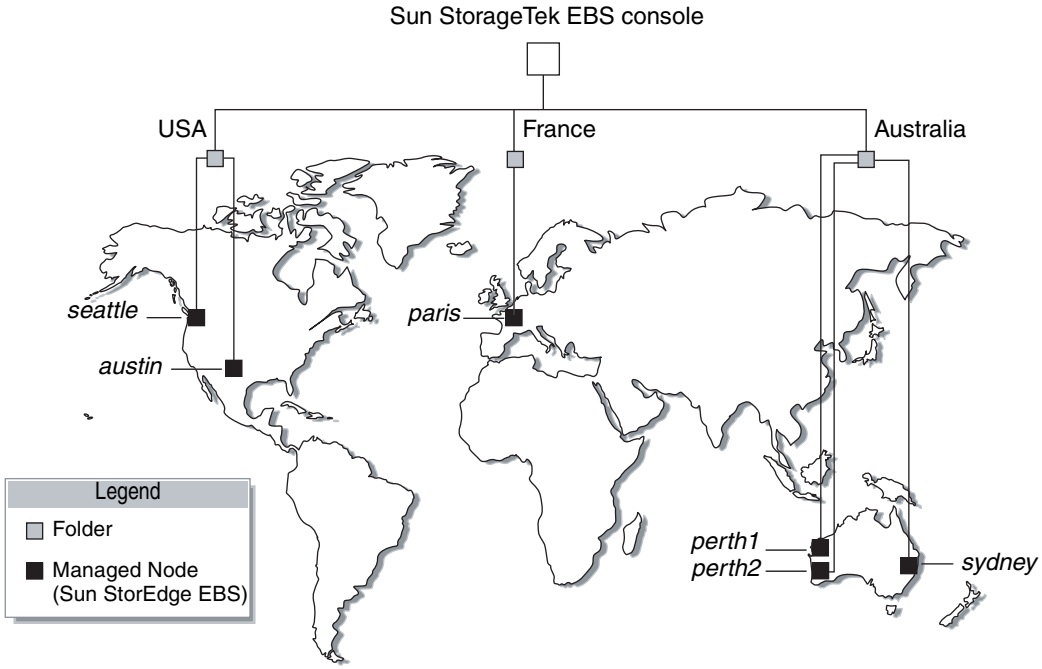


Figure 33 Sun StorageTek EBS servers worldwide

Viewing the enterprise

In the Console window, the organization of the Sun StorageTek EBS and Retrospect 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 Sun StorageTek EBS servers.
   - The right pane displays the contents of the selected folder or host.
2. Select a view option as described in Table 94.

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

Sun StorageTek EBS software enables management of Sun StorageTek EBS, Retrospect 7.5, and Avamar 3.7 and later servers within the Enterprise. You can add, delete, move, and copy Sun StorageTek EBS, Retrospect, and Avamar servers as needed. All of these functions can be performed through the Console window. When using Sun StorageTek EBS software to manage a large number of Sun StorageTek EBS 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 595 provides further information.

The addition of Sun StorageTek EBS, Retrospect, and Avamar servers provides a centralized means of managing, monitoring and reporting activity throughout the Enterprise. The server activities include, but

are not limited to, operations related to devices and libraries, and events that require user intervention.

For details on:

- Gathering Sun StorageTek EBS report data. Chapter 14, "Reporting" provides details.
- Monitoring Sun StorageTek EBS, Retrospect, and Avamar events "Viewing deduplication events in NMC" on page 869 and Chapter 15, "Monitoring" provide details.
- Retrospect events and error messages. The Dantz Retrospect documentation provides more information.
- Avamar events and error messages. The Avamar documentation provides more information.

Requirements for adding a Retrospect server as a host

The requirements for administering a Retrospect server by using the Console window include the following:

- The Retrospect server must be a Windows computer with Microsoft Remote Desktop Connection, or Microsoft Terminal services software installed.
- The Retrospect server must be configured to allow the host to connect to it from the client with the remote access software installed.
- Port 3000 must be available for Console and Retrospect to communicate.
- The Retrospect client must be a Windows computer with Microsoft Remote Desktop Connection, or Microsoft Terminal services software installed. Ensure that the mstsc.exe file is installed and runs in the host computer's PATH.

Adding a Sun StorageTek EBS, Retrospect, or Avamar host

The Console window can display Sun StorageTek EBS, Retrospect, or Avamar server events, which allows Sun StorageTek EBS, Retrospect, and Avamar server activity-reports to be generated.

Sun StorageTek EBS Console supports hosts that use Sun StorageTek EBS server software release 7.2.*x* or later, Retrospect server software release 7.5, 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**.

   Note: Hostnames and aliases cannot exceed 80 characters.

4. Select **Sun StorageTek EBS**, **Retrospect**, or **Avamar**, then 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 582 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 Sun StorageTek EBS, Retrospect, 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**.

You can also use the "drag-and-drop" feature 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**.

You can also use the "drag-and-drop" feature to move hosts.

Configuring remote access for a Retrospect host

By configuring remote access, you can use Sun StorageTek EBS Console to initiate a remote desktop session with the Retrospect host. From the Retrospect host, you can then manage and administer Retrospect clients.

To initiate a connection to a Retrospect backup server, use Microsoft Remote Desktop Connection or Microsoft Terminal software. Ensure that port 3000 is available for Console and Retrospect to communicate. The appropriate Microsoft documentation provides details on how to configure the Microsoft Remote Desktop Connection or the Microsoft Terminal software.

Managing folders in the enterprise

Sun StorageTek EBS 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

Note: 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.

   Note: You can also use the "drag-and-drop" feature while holding down the **Ctrl** key to copy folders.

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.

You can also use the "drag-and-drop" feature 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 Sun StorageTek EBS 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 Sun

StorageTek EBS managed nodes with the features Capture Events and Gather Reporting Data enabled. "Using the gstmodconf command" on page 597 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 593 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 37:

Example 37 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 38 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 38 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 39 shows the error that is generated if **gstmodconf** is used to add a host that already exists in the Enterprise:

Example 39 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 40 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 40 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: clnt_create: Remote system
   error-Connection refused.
```

Configuring a Sun StorageTek EBS server

To use the Administration window to modify a Sun StorageTek EBS server, the server must meet the following requirements:

- The server must be included in the Enterprise.

  Note: Dantz Retrospect servers can also be included in the Enterprise.

  "Adding a Sun StorageTek EBS, Retrospect, or Avamar host" on page 591 provides information about adding hosts to the Enterprise.

- The server must not restrict users from viewing the Sun StorageTek EBS server.

The Administration window can be used to modify an existing Sun StorageTek EBS server by specifying attributes in the Properties dialog box of the Sun StorageTek EBS 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

Note: When setting up the Sun StorageTek EBS server, be sure to enter the Sun StorageTek EBS product serial number. The product serial number is located on the Enabler Certificate that was sent when the Sun StorageTek EBS product was ordered.

To set up the Sun StorageTek EBS 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 Sun StorageTek EBS server

Chapter 28, "Licensing" provides information on how to license the Sun StorageTek EBS 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.

Note: 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**.

Configuring ConnectEMC

If ConnectEMC was selected during the installation of Sun StorageTek EBS 7.6, you will need to configure the program using NMC and the ConnectEMC Console on the 32-bit Windows system you specified as the location to install the ConnectEMC software. The ConnectEMC Console icon appears on the desktop of this system.



IMPORTANT

Ensure that the Sun StorageTek EBS server and ConnectEMC host name are resolved with short name and fully qualified domain name (FQDN) prior to running ConnectEMC. You must be able to

ping all hosts by IP address, short name and FQDN name. If not, installing ConnectEMC fails and the Sun StorageTek EBS/ConnectEMC send home feature fails.

The ConnectEMC for Windows User Guide, available on Powerlink from **Support > Technical Documentation and Advisories > Software ~ C ~ Documentation > ConnectEMC**, provides more instructions for configuring ConnectEMC.

Configure ConnectEMC in the ConnectEMC Console

To configure ConnectEMC in the ConnectEMC Console:

1. Double-click the **ConnectEMC Console** desktop icon. If the message "Invalid callhome config file" appears, click **OK** to continue.
2. When the Console appears, the left side of the button menu contains an octagonal **STOP** icon. If the icon is red, click the **STOP** icon.
3. In the left pane, click the **[+]** box next to **Connection** and select **General**.
4. From the **Connection** menu, select **Add Connection > Email Connection** from the list box.
5. In the **Email** window, enter the following values:

   In the Connection settings section:

   - **Recipient** — **emailalert@emc.com**
   - **EMail Server Hostname/IP** — The name or IP address of the local SMTP server that supports SMTP relaying.
   - **Sender** — The administrator's email address

   In the General Settings section:

   - **Alert Type** — Select **Use this connection first**
   - **Email Type** — Select **ACSII**

   All other values — Leave as default values
6. Click **Test** to send a test email to the ConnectHome server. Verify that the email was sent successfully.
7. Click **Save**.

8. The left side of the menu contains a green octagonal **Run** icon. Click the **Run** icon.
9. From the menu, select **File > Exit**.

(Optional) Resetting the EMC Polling Interval

The Polling Interval by default is set to 5 seconds. If you want to reset the ConnectEMC Polling Interval, while still in the ConnectEMC Console and with the ConnectEMC service set to Stop, perform the following:

1. In the far left pane, click the **Settings** folder.
2. Click **General**. **POLL_INTERVAL** should display in the right pane under **Variable Name**.
3. In the **Setting Value** column for the **POLL_INTERVAL**, type an interval value in seconds (for example, type **60** for a 1 minute interval).
4. Select the floppy disk icon on the left side of the menu to save changes.
5. If necessary, restart the ConnectEMC service (see step 8 above).

Configure ConnectEMC in NMC

After configuring the settings for ConnectEMC using the ConnectEMC console, some tasks need to be performed in the Sun StorageTek EBS Management Console. These tasks are:

- Specifying the name of the ConnectEMC host
- Setting Autostart to run nsrconnect

Name of the ConnectEMC host

When ConnectEMC is installed, the host is set to the Sun StorageTek EBS server by default, if no host was specified during installation.

To set or verify the name of the ConnectEMC host:

1. Launch NMC.
2. Click the **Configuration** button. An entry for ConnectEMC appears in the tree of the left navigation pane.
3. Click **ConnectEMC**. In the ConnectEMC window pane, right-click **DefaultConnectEMC** and select **Properties**.

4. In the ConnectEMC host field, type the name of the host that will be used to run the ConnectEMC daemon if the host specified is not correct.

Configure ConnectEMC to send

To configure ConnectEMC to send:

1. Connect to the Sun StorageTek EBS server using NMC.
2. Click the **Configuration** button. An entry for ConnectEMC appears in the tree of the left navigation pane.
3. Click **ConnectEMC**. In the ConnectEMC window pane, right-click **DefaultConnectEMC** and select **Properties**.
4. Confirm that the correct Sun StorageTek EBS server is specified in the ConnectEMC Host field.
5. Select **Enabled** from the Autostart drop-down menu.
6. Click **Ok**.

Setting Autostart to run nsrconnect

nsrconnect is the Sun StorageTek EBS interface to ConnectEMC. The nsrconnect task is run to create the .xml file that contains the system configuration information stored in the RAP database. When Autostart is set to Enabled, nsrconnect will be launched one day after the Sun StorageTek EBS server is started (or restarted), and once per month thereafter. For example, if the Sun StorageTek EBS server (nsrd) is started on Monday, nsrconnect launches on Tuesday, and then again a month from Tuesday.

To start the nsrconnect task immediately, set AutoStart to Start now.



IMPORTANT

Start Now can only be initiated on the server once every 15 minutes due to the 15 minute throttle limit.

Configure Sun StorageTek EBS to pass contact information to Connect Home server

In order to quickly and accurately identify the source of server information, provide the company contact information:

1. Right-click the Sun StorageTek EBS server in the navigation pane, and select **Properties**.
2. Click the **Customer Information** tab.

3. Complete the following fields:
   - Contact Name
   - Company
   - Street Address
   - City/Town
   - State/province
   - Zip/postal code
   - Country
   - Phone Number
4. Click **Ok**.

Enabling BMR server support

Connection with a HomeBase Server is enabled during the set up of the Sun StorageTek EBS server. This connection enables the delivery of profile data from the Sun StorageTek EBS client to the HomeBase Server.

"Configuring a client for BMR support" on page 877 provides information for configuring and enabling a BMR client.

Parallelism and multiplexing

Parallelism is a general term within the Sun StorageTek EBS 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 Sun StorageTek EBS environment.

Parallelism

Several attributes in various Sun StorageTek EBS 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 Sun StorageTek EBS 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 Sun StorageTek EBS server:

- For the Client resource for a Sun StorageTek EBS 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.
- For all other clients (that is, clients that are not also Sun StorageTek EBS servers), the default value is 4. Many features within Sun StorageTek EBS, such as deduplication, NDMP backups, and VSS, will not work correctly with client parallelism set to a value that is higher than 4.

For all clients other than the Client resource of the Sun StorageTek EBS server, to avoid disk contention, you should generally 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), higher client parallelism will result in multiple save streams attempting to access the disk at the same time.

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 Sun StorageTek EBS software. Each enabled storage node connected to the Sun StorageTek EBS server can increase the parallelism maximum. The maximum parallelism for any Sun StorageTek EBS server and storage node combination can vary. The *EMC NetWorker Release Notes* provides more information.

Optimally, your Sun StorageTek EBS 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 Sun StorageTek EBS 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

The maximum number of active devices for a storage node is defined by using the Max active devices attribute of the Storage Node resource, found in the Sun StorageTek EBS 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 Sun StorageTek EBS software may use from the storage node. 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 Sun StorageTek EBS 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 Sun StorageTek EBS 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 Sun StorageTek EBS environment.

The Sun StorageTek EBS 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
- Max sessions
- Pool parallelism

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 Sun StorageTek EBS 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, described below.

Instead of a hard limit, the Target sessions attribute aids in load balancing for devices by determining when the Sun StorageTek EBS software should begin writing save streams to another device:

- If a device is already receiving the number of backup sessions determined by the target sessions value, the Sun StorageTek EBS 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 Sun StorageTek EBS 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 Sun StorageTek EBS server to multiplex multiple save sets to the same device, rather than write

each save set to a separate device, the Sun StorageTek EBS 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.

Note: When the Sun StorageTek EBS 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. It does not make a difference if the device with the lowest target session value is associated with a different pool than the incoming savestreams.

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 Sun StorageTek EBS 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 Sun StorageTek EBS 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 Sun StorageTek EBS 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.

Managing server access

Sun StorageTek EBS server access is managed by granting access privileges to Sun StorageTek EBS Console users.

Granting access to manage Sun StorageTek EBS servers from Console server

If the Console server and the Sun StorageTek EBS server are installed on separate hosts, the Console administrator must be granted access to administer and monitor the target Sun StorageTek EBS server. To grant access, use the **nsraddadmin** command from the command prompt on the Sun StorageTek EBS server. Grant access on each Sun StorageTek EBS server that is to be managed from the Console server. The *Sun StorageTek EBS Installation Guide* provides information about granting access with the **nsraddadmin** command. The *EMC NetWorker Command Reference Guide* or the UNIX man pages provides information about the **nsraddadmin** command.

Administrator privileges

Use the Administration window to grant Sun StorageTek EBS administrator privileges to additional Console users. This is done by adding users to the Administrators group for the Sun StorageTek EBS server. "Sun StorageTek EBS User Groups" on page 613 provides more information.

When granting privileges to users, note the following requirements:

- To change security settings, the user must be a member of the Sun StorageTek EBS Administrators group.
- To perform a scheduled backup of SYSTEM and VSS SYSTEM save sets, the user must be a member of the Sun StorageTek EBS Administrators group.
- To add to or change the configuration of the server, media devices, and media libraries, the user must have the Configure Sun StorageTek EBS privilege.

Sun StorageTek EBS software ships with default administrator settings that give root users on UNIX and Linux, and members of the Windows Administrators group on Windows, permission to change its configurations. The Administrators group in the Server resource should include the following entries, depending on the platform of the client trying to administer the Sun StorageTek EBS server:

- Microsoft Windows:

  ```
  user=administrator, host=server_name
  ```

- UNIX and Linux:

```
user=root, host=server_name
```

Note: If a hostname is used, the Sun StorageTek EBS server must be able to convert the hostname to a resolved TCP/IP address.

Additional Console users or user groups can be granted administrative privileges by adding them to the Administrators group.

Usernames should be listed in the form of name/value pairs by using the 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 grant administrative privileges to a user that is named *jdoe* on the host named *jupiter*, enter this line in the Users attribute:

user=jdoe,host=jupiter

Note: The formats *user@host*, *host/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 preferred.

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 Sun StorageTek EBS from any host (this is the same as an entry of user=*username*). 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 41 Using the Sun StorageTek EBS Administrators Group

This example shows what to enter to provide Sun StorageTek EBS 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 and Sun StorageTek EBS administrative privileges

The Sun StorageTek EBS 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 individual 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 administrative rights to the server. The syntax for this user description is as follows:

- Single user:

  `user=`*user_name*`, domain=`*domain_name*

  For example:

  ```
  user=joe,domain=NT-ENG
  ```

- Group:

  `group=`*group_name*`, domainsid=`*domain_id*

  For example:

```
group=Administrators,domainsid=S-1-5-32-323121-123
```

Sun StorageTek EBS User Groups

Sun StorageTek EBS includes an access control feature that is configured through the User Group resource. This feature allows Sun StorageTek EBS administrators to create user groups, define the privileges associated with those user groups, and assign users to various groups. Privileges include such things as permission to monitor Sun StorageTek EBS software, as well as to back up and recover local data.

To create user groups, you must have an Access Control enabler installed on the Sun StorageTek EBS server. One enabler per datazone is required. This feature is available only on servers with Sun StorageTek EBS release 7.0 and later.

Preconfigured User Groups

Sun StorageTek EBS comes with two preconfigured user groups:

- Administrators
- Users

Administrators

Members of the Administrators group have permission to perform all Sun StorageTek EBS functions, except with regard to some utilities that require root or Microsoft Windows Administrator access. The UNIX root user and members of the Microsoft Windows Administrators group are always members of this group and cannot be removed from the group.The *EMC NetWorker Command Reference Guide* or the UNIX man pages provide information on the permission requirements for a specific utility.

Privileges associated with the Administrators group cannot be changed.

Users

By default, members of the Users group are granted permission to back up and recover local data and to monitor Sun StorageTek EBS operations. They cannot view or edit configurations.

Customizing privileges

The privileges associated with the Users group can be customized to fit your requirements. Privileges associated with the Administrators group cannot be edited. Table 95 lists these preconfigured privileges and their associated permissions.

Table 95 User Group privileges (1 of 3)

Privilege	Permissions
Change security settings	Allows group members to edit these resources and attributes: • User Group resources • Remote Access attribute in the Client resource • Archive Users attribute in the Client resource • Administrators attribute in the Server resource • Notification resource • Owner Notification attribute in the Client resource **Note:** If you change the users listed in the Administrator attribute of the Server resource, the changes are automatically reflected in the Administrators group as well. Similarly, if you change the members of the Administrators group, the changes are automatically reflected in the Administrator attribute. Members of groups with this privilege can grant access to restricted functions to any user. User groups with this privilege must also have these privileges: Configure Sun StorageTek EBS, Operate Sun StorageTek EBS, Monitor Sun StorageTek EBS, Operate Devices and Jukeboxes, Backup Local Data, and Recover Local Data.
Remote access all clients	Allows group members to remotely browse and recover data associated with any client, as well as view configurations for all Client resources. This privilege is required to perform Directed Recovers. This privilege supersedes the Remote Access attribute in the Client resource. User groups with this privilege must also have these privileges: Operate Sun StorageTek EBS, Monitor Sun StorageTek EBS, Operate Devices and Jukeboxes, Backup Local Data, and Recover Local Data.
Configure Sun StorageTek EBS	Allows group members to configure resources associated with the Sun StorageTek EBS server, its storage nodes, and clients. This includes creating, editing, and deleting resources. Users with this privilege cannot configure User Group resources or the Remote Access or Archive Users attributes in the Client resource. Permission to change these settings is granted only through the Change Security Settings privilege. User groups with this privilege must also have these privileges: Operate Sun StorageTek EBS, Monitor Sun StorageTek EBS, Operate Devices and Jukeboxes, Backup Local Data, and Recover Local Data.

Table 95 User Group privileges (2 of 3)

Privilege	Permissions
Operate Sun StorageTek EBS	Allows group members to perform Sun StorageTek EBS server operations. For example, members can: • 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. User groups with this privilege must also have these privileges: Monitor Sun StorageTek EBS, Operate Devices and Jukeboxes, Backup Local Data, and Recover Local Data.
Monitor Sun StorageTek EBS or Retrospect	Allows group members to: • Monitor Sun StorageTek EBS or Retrospect operations, including device status, save group status, and messages. • View media databases information. • View Sun StorageTek EBS configuration information (except the security settings described in the Change Security Settings privilege). 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.

Table 95 User Group privileges (3 of 3)

Privilege	Permissions
Operate devices and jukeboxes	Allows group members 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. User groups with this privilege must also have Monitor Sun StorageTek EBS privileges.
Recover local data	Allows group members to recover data from the Sun StorageTek EBS 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. **Note:** This privilege does not provide permission to view information about other clients. This privilege does not override file-based permissions. Users can only recover files with the appropriate user permissions 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.
Backup local data	Allows group members to manually back up data from their local client to the Sun StorageTek EBS server, 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. This privilege does not provide permission to view information about other clients. **Note:** This privilege does not override file-based permissions. Users can only back up files with the appropriate user permissions 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.

Working with User Group resources

This section provides procedures for creating, editing, copying, and deleting a User Group resource.

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, enter the names of the users that should be members of the user group by using the format described in "Administrator privileges" on page 610.
5. Select the privileges to assign to the user group.
6. Click **OK**.

Editing a User Group resource

To edit 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. Make any necessary changes, then click **OK**.

Note: You cannot change the privileges associated with the Administrators group. Also, you cannot remove the UNIX root user or the Windows Administrator for the server being configured.

Copying a User Group 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 Group resource

Note: You cannot delete a preconfigured user group.

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.

Server communication issues within Microsoft Windows

This sections addresses various client/server communication issues that occur when running Sun StorageTek EBS 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 Sun StorageTek EBS 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 Sun StorageTek EBS 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 permissions 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 Sun StorageTek EBS server but not another, check Backup Operators group on the problematic computer to ensure that you are properly assigned.

Dynamic Host Configuration Protocol

A Sun StorageTek EBS 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 Sun StorageTek EBS server becomes invalid. If the Sun StorageTek EBS server is a DHCP client, a static TCP/IP address for the server must be preserved. Sun StorageTek EBS clients can still use dynamic addressing.

Backup and Recover Server service

In Microsoft Windows operating systems, the Sun StorageTek EBS 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.

Notifications

A notification provides information about events that occur in a Sun StorageTek EBS environment. You can configure the events to be reported, as well as how the Sun StorageTek EBS server reports them to you. Specific programs can be executed when an event occurs, including third-party programs. By default, the Sun StorageTek EBS server sends notifications to log files located in the <Sun StorageTek EBS_install_dir>\logs directory.

Note: Retrospect server events can be displayed in the Console window.

Preconfigured notifications

Sun StorageTek EBS is preconfigured to provide most of the notifications required to monitor Sun StorageTek EBS and Retrospect

events. Table 96 lists these preconfigured notifications and the associated actions performed by the Sun StorageTek EBS server.

Table 96 Preconfigured notifications (1 of 3)

Notification	Action performed by the Sun StorageTek EBS server
Bootstrap	Sends the bootstrap information to the default printer of the Sun StorageTek EBS server.
Bootstrap backup failure	Sends this email to root/Administrator: Bootstrap backup failed.
Cleaning cartridge expired	Windows: Reports expired cleaning cartridge to the media log file. UNIX: Sends this email to root: replace the expired cleaning cartridge.
Cleaning cartridge required	Windows: Reports the need for cleaning cartridge to the media log file. UNIX: Sends this email to root: insert a cleaning cartridge.
Client install	Windows: Reports hostname and Sun StorageTek EBS client software version information to the media log file. UNIX: Sends this email to root: host *host_name* installed *product_version*. Where *host_name* is the name of the Sun StorageTek EBS host, and *product_version* is the Sun StorageTek EBS client software release and build number.
Device cleaned	Windows: Reports that a device has been cleaned to the media log file. UNIX: Sends this email to root: device cleaning is complete.
Device cleaning required	Windows: Reports that a device requires cleaning to the media log file. UNIX: Sends this email to root: clean the device.
Device disabled	Windows: Reports that a device has been automatically disabled to the media log file. UNIX: Sends this email to root: a device has been automatically disabled.
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 media log file. UNIX: Sends this email to root: advanced file volume is full.
Index size	Windows: Sends a message about the status of the index to the index log file. UNIX: Sends this email to root: check the size of the client file index because it will soon exceed the space available.

Table 96 Preconfigured notifications (2 of 3)

Notification	Action performed by the Sun StorageTek EBS server
Log default	Windows: Sends data about Sun StorageTek EBS events to the messages log file. UNIX: Directs the UNIX syslog facility (/usr/bin/logger) to log information from the /nsr/logs/messages file into the appropriate system log file. SuSE Linux: System log notifications might fail on SuSE Linux Enterprise Server 8 and 9. The logger binary appears in the /bin directory instead of /usr/bin. This causes the Log default notification to fail. To resolve this issue: • Create a link in the /usr/bin directory to /bin/logger. **Note:** The default Tape Mount Request 1 and Tape Mount Request 2 notifications must be updated if a link is not created. • Modify the Log Default notification and change the /usr/bin/logger filepath to /bin/logger.
Sun StorageTek EBS Daemons Not Running	Sends this email to root/Administrator: Sun StorageTek EBS daemons not running.
New Virtual Machine	Windows: Sends message that new virtual machines have been detected to the messages log file. UNIX: Sends email that new virtual machines have been detected.
Registration	Windows: Sends messages about the registration status of your Sun StorageTek EBS products to the messages log file. UNIX: Sends this email to root: check the registration status.
Resource File Corruption	Sends this email to root/Administrator: Resource file corruption.
Savegroup completion	Windows: Reports the success of client save sets to the savegrp log file. UNIX: Sends this email to root: degree of success in completing scheduled backups, cloning, and archive operations.
Save set marked suspect	Sends this email to root/Administrator: Save set marked suspect.
SNMP notification request	Sends event notifications to a network management console (this notification only appears if the NetWorker SNMP module has been purchased and enabled). Chapter 22, "SNMP Module" provides details on SNMP notifications
Storage Node Registration	Registers a client as a storage node.
Tape mount request 1	Windows: Requests media be mounted in a device and displays a pending message. UNIX: Sends a request to the messages file to mount a backup volume, using a priority of Waiting.
Tape mount request 2	Windows: Requests media be mounted in a device and displays a pending message. UNIX: Sends a request to the messages file to mount a backup volume, using a priority of Critical.

Table 96 Preconfigured notifications (3 of 3)

Notification	Action performed by the Sun StorageTek EBS server
Tape mount request 3	Windows: Sends a request for mounting a backup volume, using a priority of Alert, to the media.log file. UNIX: Sends an email to root to request that the tape be mounted.
Tape mount request 4	Sends an email to root/Administrator to request that the tape be mounted.
Verify Label failed on unload	Sends an email to root/Administrator: Unable to read the tape label.
Volume Marked full	Sends an email to root/Administrator: Volume marked full.

Customizing notifications

Notifications require the following three elements:

- "Events" on page 622
- "Actions" on page 623
- "Priorities" on page 626

Events

Table 97 lists the events that trigger a notification.

Table 97 Events (1 of 2)

Event	Description
adv_file	The file system is full and is waiting for additional space.
Bootstrap	The bootstrap backup failed.
Cleaning cartridge expired	A cleaning cartridge has expired and needs replacing.
Cleaning cartridge required	Mount the cleaning cartridge.
Client	Sun StorageTek EBS 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.
Hypervisor	New virtual clients have been auto-discovered, or auto-discovery failed.

Table 97 Events (2 of 2)

Event	Description
Index attention	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.
Registration	Product registration needs attention.
Savegroup	A backup group has completed the backup.
Server	Other server events (for example, restarting the Sun StorageTek EBS server).
Storage node	A storage node has been installed.
Write completion	A write operation is complete.

Note: The Dantz Retrospect documentation provides details on Retrospect events displayed in the Console window.

Actions

Table 98 lists the actions that the server takes after an event notification occurs.

Table 98 Actions (1 of 2)

Action	Description
eventlog	(Windows only) Logs the notification message to the event log.
nsrlog	(Windows only) Sends a message about an event to a file.
nsrlpr	(Windows only) Prints information to a printer. "Using nsrlpr to print notifications" on page 624 provides more information about using **nsrlpr**.
logger	(UNIX only) Uses the UNIX **syslog** facility (/usr/bin/logger) to log information or send messages.

Table 98 Actions (2 of 2)

Action	Description
lp	(UNIX only) Prints the notification.
mail	(UNIX only) Sends email to the specified user.
nsrtrap	Sends notifications to an SNMP management console.

Third-party programs can also be used for the action, as long as the programs support reading from *standard input*. For example:

- With UNIX systems, you can use a third-party email program rather than the **mail** command.
- With Windows systems, you can use a third-party email 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 group, or users with the Change Security Settings privilege, can change the Action attribute of an existing notification.

Note: The Dantz Retrospect documentation provides details on Retrospect actions.

Using nsrlpr to print notifications

Sun StorageTek EBS 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** command.

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, enter this at the command prompt:

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*

Note: Print jobs sent by the Sun StorageTek EBS 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 website provide more information.

Priorities

Each Sun StorageTek EBS 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 Sun StorageTek EBS server sends a mount backup volume request, the priority assigned to the message is Waiting. The priority of the second request is Critical. The priority of the third request is Alert.

Table 99 lists the priorities upon which notifications are based.

Table 99 Priorities

Priority	Description
Information	Information about the current state of the server.
Notice	Important information.
Warning	A non-fatal error has occurred.
Waiting	The Sun StorageTek EBS server is waiting for an operator to perform a routine task, such as mounting a backup volume.
Critical	The server detected an error that should be fixed.
Alert	A severe condition exists that requires immediate attention.
Emergency	A condition exists that may cause Sun StorageTek EBS to fail unless corrected immediately.

Note: Event priorities are sorted alphabetically, rather than by severity.

Logging event notifications

Sun StorageTek EBS keeps two general notification log files. By default, these files are located in <Sun StorageTek EBS_install_dir>\logs.

- messages log file

  The data in the messages log file is generated by nsrlog, which is part of the Sun StorageTek EBS event notification mechanism. The nsrlog program is triggered by a notification, and it prints the file to the messages log file.

- daemon log file

  The **nsrd**, **nsrexecd**, and their subordinate processes redirect their output to the daemon log file.

"Viewing log files" on page 904 provides information about viewing log files.

To better access and use these event logs on Windows systems, an Event Logging mechanism enables applications to log messages 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 100.

Table 100 Event Viewer messages

Event Viewer category	Displayed information
Source	Events from Sun StorageTek EBS software always designate Sun StorageTek EBS as the source.
Category	Mapped from Sun StorageTek EBS notification event type (savegroup, server, registration, and so on).
Severity	Mapped from Sun StorageTek EBS notification priority: • Alert and Emergency are mapped to Error. • Priorities between Critical and Warning are mapped to Warning. • Notification and Information are mapped to Information.
Event ID	Events from Sun StorageTek EBS software always designate the numeral 1 for the ID.

Creating a custom notification

Sun StorageTek EBS also provides preconfigured notifications. "Preconfigured notifications" on page 619 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 101 on page 631 provides command options.
7. Click **OK**.

Editing a notification

Note: You cannot change the name of 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**.

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.

Note: You cannot delete any of the preconfigured notifications.

Indexes

The Sun StorageTek EBS 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 636 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 Sun StorageTek EBS 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 permissions
- 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.

Note: If the file system that contains the indexes gets full, the Sun StorageTek EBS 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.

Sun StorageTek EBS 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 206 provides information on policies.

Automated index activities

The Sun StorageTek EBS 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 Sun StorageTek EBS 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 638 provides more information.

Checking online indexes

Each time the Sun StorageTek EBS 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 online indexes for a Sun StorageTek EBS server" on page 452.

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 101 identifies the index information displayed for each client.

Table 101 Indexes window information

Column	Description
Client Name	Names of the Sun StorageTek EBS 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 102 identifies the information in the **Save Sets** dialog box for each save set.

Table 102 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 636 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.

   Note: 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 Sun StorageTek EBS 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 Sun StorageTek EBS 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 Sun StorageTek EBS server displays this prompt:

```
Cross-checking may take considerable time. Would you
like to cross-check
client_name?
```

4. Click **Yes** to continue. The Sun StorageTek EBS 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 Sun StorageTek EBS software normally designates a default location for the client file index on the Sun StorageTek EBS server. This default location is:

- For UNIX: /nsr/index/client_name
- For Windows: <*Sun StorageTek EBS_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 75 provides instructions.
6. Click **OK**.

Changing the client file index location for an existing client

Note: 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 636 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: /nsr/indexes/jupiter contains 54 records
   occupying 7 KB
   nsrck: Completed checking 1 client(s)
   ```

   Note: 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 Sun StorageTek EBS server.
2. Copy the client file index from its current location to the new location. For example:
   - For Windows, copy <*Sun StorageTek EBS_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 635 provides instructions.

Managing the size of the online indexes

Over time, the size of the online indexes on the Sun StorageTek EBS 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 Sun StorageTek EBS 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 638 provides details.
- Delete volume-based entries from the client file index. "Deleting volume-based online index entries" on page 639 provides details.
- Adjust the Browse Policy and Retention Policy attributes of clients backing up to the Sun StorageTek EBS 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 220 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 636 provides details.

Reducing media database size

Reduce the size of media database on the Sun StorageTek EBS server by using one or more of these methods:

- Remove volumes that contain recyclable save sets from the Sun StorageTek EBS inventory. "Removing volume-based entries from the online indexes" on page 638 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.

  Note: 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 392 provides details.

  When a volume is recycled, the Sun StorageTek EBS 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 220 provides details.

- Compress the media database."Compressing the media database" on page 640 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 Sun StorageTek EBS 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 and recover these files by using the **NetWorker User** program on Windows or the **nwrecover** program on UNIX. The save set recover procedure must be used to recover data once the client file index

entries are removed. "Recovering by save set selection" on page 436 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 Sun StorageTek EBS 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 online indexes for a Sun StorageTek EBS server" on page 452 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.

Note: The presence of a clone of the particular volume prevents the deletion of the volume entry in the media database. This is because the Sun StorageTek EBS 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 638 provides details about this option.
   - **File Index Entries Only**. "Removing client file index and media database entries" on page 639 provides details about this option.
5. Click **OK**.

The Sun StorageTek EBS 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
```

Deleting deduplication save sets

"Deleting deduplication save sets" on page 871 provides information on deleting deduplication save sets.

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:

     <*Sun StorageTek EBS_install_dir*>\mm\.cmprssd
   - On UNIX:

     /nsr/mm/.cmprssd
2. Type at the command prompt:

   ```
   nsrim
   ```

Querying deduplication save sets using mminfo

"Querying deduplication save sets using mminfo" on page 871 provides information on querying deduplication save sets.

Updating Sun StorageTek EBS clients by using the software distribution feature

Use the software distribution feature to remotely distribute and update the Sun StorageTek EBS software from a centralized Sun StorageTek EBS server to one or more Sun StorageTek EBS clients.

You can update these Sun StorageTek EBS software packages on computers that have the Sun StorageTek EBS release 7.3 or later client software already installed:

- Client
- Storage node
- Man pages
- NMO, NMSQL and NMExch Sun StorageTek EBS Application Modules

Note: The software distribution feature is not supported on HP Tru64 UNIX, IRIX, Mac OS X, NetWare, Open VMS platforms, and cluster environments.

The software distribution feature can be used to:

- Manage the software repository.
- Inventory Sun StorageTek EBS software installed on Sun StorageTek EBS clients.
- Update Sun StorageTek EBS software packages on existing Sun StorageTek EBS clients.
- Monitor software distribution inventory and upgrade operations (only available from the Software Administration Wizard).

Software distribution can be performed on the Sun StorageTek EBS client software using either the Software Administration Wizard or the **nsrpush** command. The *Sun StorageTek EBS Installation Guide* or the **nsrpush** man page provides information on how to update the Sun StorageTek EBS client software.

Message log management

The Sun StorageTek EBS server stores the messages generated by the Sun StorageTek EBS server daemons in a message:

- On Windows: <Sun StorageTek EBS_install_dir>\logs
- On UNIX: /nsr/logs

You can use environment variables and/or UNIX startup scripts to automatically control the size of the log.

Setting the UNIX startup script to trim log files

To modify the way that Sun StorageTek EBS services manage the Sun StorageTek EBS log files, the NSR_MAXLOGSIZE environment variable can be set in the Sun StorageTek EBS startup script before the line:

```
(echo 'starting Sun StorageTek EBS daemons:') >
/dev/console
```

Table 103 lists the location of Sun StorageTek EBS startup scripts:

Table 103 Sun StorageTek EBS startup script location

Operating system	Startup script location
Solaris, Linux, and IRIX	/etc/init.d/networker
HP-UX	/sbin/init.d/networker
HP Tru64	/sbin/init.d/NSRstartstop
AIX	/etc/rc.nsr

Enter the environment variables by using the format (bourne shell example):

```
ENV_VAR_NAME=value
export ENV_VAR_NAME
```

These environment variables can be used to trim log files at startup:

- To change the maximum size of log files, change the NSR_MA XLOGSIZE value. The default value for NSR_MAXLOGSIZE is 1048576 bytes.
- To change the maximum number of log files that are saved, change the NSR_MAXLOGVERS value. The default value is 4.

  For example:

  NSR_MAXLOGVERS=3

  export NSR_MAXLOGVERS

```
NSR_MAXLOGSIZE=524288
export NSR_MAXLOGSIZE
```

Restart the Sun StorageTek EBS server for the environment variables to take effect.

Every time the Sun StorageTek EBS server starts, it checks the size of the daemon log file. When the daemon log file reaches the size defined in NSR_MAXLOGSIZE or the default size of 2 MB, the following occur:

1. The existing daemon.raw log file is renamed daemon_<date_time>.raw and a new empty daemon.raw log file is created.
2. This process is repeated until the value in NSR_MAXLOGVERS is reached, at which point the oldest log is removed.

Note: The trimming mechanism functions only when you start **nsrd** on UNIX or the Sun StorageTek EBS Backup and Recover Server service on Windows. The daemon and service do *not* check periodically to determine if the log file has exceeded NSR_MAXLOGSIZE. If the **nsrd** daemon or Sun StorageTek EBS Backup and Recover Server service runs for a long time, the log file can still grow very large.

To activate the trimming mechanism:

- On UNIX, enter **nsr_shutdown** to stop the Sun StorageTek EBS daemons, and then use the Sun StorageTek EBS startup script to restart the daemons.
- On Windows, stop and restart the Sun StorageTek EBS Backup and Recover Server service.

The *EMC NetWorker Command Reference Guide* or the UNIX man pages provide information about the **nsr_shutdown** command and its options.

Savegroup job information

By default, when a group is backed up, details of the corresponding savegroup's child jobs are temporarily saved to files in the following location, NetWorker_*install_dir*\nsr\tmp\sg*savegroup_name*. There is no easy way to associate the output of a savegroup's child jobs with the parent savegroup job. Additionally, these output files were overwritten when the group is next backed up.

Beginning with NetWorker 7.5, you can specify that the details of a savegroup's child jobs be saved in files that are named according to the child job ID. These files are saved to, NetWorker_*install_dir*\nsr\logs\sg*savegroup_name**Job_ID* where *Job_ID* is a file whose name corresponds to a savegroup child job ID.

To help locate and retrieve information about a savegroup's child jobs, use the **jobsquery** command to search the NetWorker jobs database. The man pages or the *EMC NetWorker Command Reference Guide provide* more information.

To have savegroup child logs saved by job ID:

1. In the **Administration** window, click **Configuration**.
2. In the left pane of the **Configuration** window, select the Sun StorageTek EBS server.
3. From the **File** menu, select **Properties**.
4. Select the **Configuration** tab.
5. Select the **Savegroup log by job id** attribute.
6. Click **OK**.

Note: When the Savegroup log by job id attribute is enabled, the job files saved to NetWorker_*install_dir*\nsr\logs\sg*savegroup_name* will be purged based on the settings in the NetWorker server attributes named, Jobsdb retention in days and Jobsdb maximum size (KB).

Internationalization

Sun StorageTek EBS releases 7.4 and later have been internationalized. As a result, the Sun StorageTek EBS software now supports language packs, which can be installed as part of the Sun StorageTek EBS installation, or can be installed separately after the Sun StorageTek EBS software has been installed. The *Sun StorageTek EBS Installation Guide provides more information*.

Internationalization support in the Sun StorageTek EBS software is dependent on the underlying operating system's internationalization support. If you are planning on using non-English data in the Sun StorageTek EBS 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 Sun StorageTek EBS software in a multi-language environment, which are addressed in the following sections.

Log file viewer

Sun StorageTek EBS log files must be viewed using the **nsr_render_log** program. "Viewing log files" on page 904 provide more information.

Interoperability with previous releases of Sun StorageTek EBS

Multiple locales within the same datazone are fully supported only if all Sun StorageTek EBS installations in the datazone are at release level 7.4. The following limitations apply:

- In datazones with a Sun StorageTek EBS release 7.3 server and Sun StorageTek EBS 7.4 clients, support for scheduled backups of path or filenames containing non-ASCII characters is limited to the support provided by Sun StorageTek EBS 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 Sun StorageTek EBS 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 and Linux:

- 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 Sun StorageTek EBS 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 and Linux, 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.

Group completion messages

Group completion messages displayed in the Sun StorageTek EBS console, displayed either through the Group details, Group status, or monitoring, will be displayed in the language used by the Sun StorageTek EBS server that generated the message.

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 and Linux 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 671 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.

18

Sun StorageTek EBS Client Management

This chapter covers these topics :

- Sun StorageTek EBS client overview 650
- Client configuration 650
- Creating a client probe 655
- Associating a probe with a Client resource 656
- Deduplication client 656
- Configuring a client for BMR support 657
- Lockbox pass phrase management 657
- Sun StorageTek EBS authentication 658
- Multiple clients from the same computer 671
- Scheduled backups of non-ASCII files or directories 672
- Adding or changing the Sun StorageTek EBS servers for a client ... 673
- Editing the servers file 674
- Client priority 674

Sun StorageTek EBS client overview

A Sun StorageTek EBS client is both a physical computer with Sun StorageTek EBS client software installed on it, and a Sun StorageTek EBS server resource that specifies a set of files and directories to be included in a scheduled backup. As such, a single Sun StorageTek EBS client computer can have several Client resources specified that can back up to the same or even different Sun StorageTek EBS servers.

"Defining a client and save set combination" on page 672 further explains concept of a client computer with multiple Sun StorageTek EBS Client resources.

Sun StorageTek EBS client software is available for a variety of platforms. No matter which platform the client resides on, it can be backed up to any Sun StorageTek EBS server. For example, you can back up a Sun StorageTek EBS client on a Microsoft Windows computer to a Sun StorageTek EBS server on a Solaris computer.

Client configuration

Before a client can be backed up by a Sun StorageTek EBS server, the client computer must have the appropriate Sun StorageTek EBS client software installed. The Sun StorageTek EBS Installation Guide provides more information.

Creating a client

Sun StorageTek EBS client creation is part of the process of creating a scheduled backup. "Setting up a scheduled backup" on page 75 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 651 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 652.
7. Stop all Sun StorageTek EBS services.
8. On the Sun StorageTek EBS 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:

  <*StorageTek_EBS_install_path*>`\index\`*client_name.domain.com*

9. Restart the Sun StorageTek EBS 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 75 provides instructions on how to create a client.

Deleting a client

When a client is deleted, the Sun StorageTek EBS 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 Sun StorageTek EBS 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 81 provides instructions on how to create a client.

The Sun StorageTek EBS 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 Sun StorageTek EBS 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: "Enabling directed recoveries" on page 445 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 Sun StorageTek EBS 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
      ```
4. 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
   ```

5. Type **Yes** when prompted to confirm the change.

Note: 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:

1. Log in as root or as Windows Administrator on the Sun StorageTek EBS client.
2. Type the following at the command prompt:

   ```
   nsradmin -p nsrexec
   ```

   The **nsradmin** prompt appears.
3. 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";
   ```
4. 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

Note: 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 656.

Note: The **Command Options:** field is applicable 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 655 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.

Deduplication client

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 Sun StorageTek EBS side.

"Creating a deduplication node and a replication node" on page 860 provides information on how to create and configure aSun StorageTek EBS deduplication client.

Configuring a client for BMR support

The HomeBase Agent is automatically installed with the Sun StorageTek EBS client. After enabling communication with a HomeBase Server, configure a Sun StorageTek EBS client or start by editing an existing client. "Configuring a client for BMR support" on page 877 provides detailed instructions.

Note: A HomeBase Agent license batch code is required to enable the client to generate profiles and send them to the HomeBase Server. Contact your EMC HomeBase representative for licensing information.

Lockbox pass phrase management

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.

To use the lockbox pass phrase management feature for a NetWorker module:

1. Log in as root, or as Windows administrator on the Sun StorageTek EBS 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 Sun StorageTek EBS server

"Set the Datazone pass phrase for a Sun StorageTek EBS server" on page 97 provides information on how to set the Datazone pass phrase on a NetWorker server.

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.

Sun StorageTek EBS authentication

This section describes how to configure authentication between Sun StorageTek EBS hosts. This section also includes special considerations for authentication.

Two types of authentication are supported for Sun StorageTek EBS 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. Sun StorageTek EBS hosts and Sun StorageTek EBS user permissions are authenticated by using **nsrauth**. The **nsrauth** authentication

mechanism is available for hosts that run Sun StorageTek EBS software release 7.3 or later.

Each Sun StorageTek EBS 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 Sun StorageTek EBS permissions

Authentication for backwards compatibility (oldauth)

For compatibility with earlier Sun StorageTek EBS 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 660 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 Sun StorageTek EBS 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 Sun StorageTek EBS 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 Sun StorageTek EBS release 7.3 or later server be allowed to authenticate with legacy Sun StorageTek EBS clients by using **oldauth.**
- That all other Sun StorageTek EBS 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 659 provides more information.

To specify the authentication strength between Sun StorageTek EBS hosts:

1. From the **Administration** window, click **Configuration**.

2. Open the **Local Host**s 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 Sun StorageTek EBS 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 Sun StorageTek EBS client name.
     - *10.102.0.0/255.255.0.0* - A subnet IP address representing all Sun StorageTek EBS 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



IMPORTANT

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 662 provides details.

6. Click **OK**.
7. Restart the Sun StorageTek EBS services for the host with the authentication relationships to be configured. "Stopping and starting a Sun StorageTek EBS server, client, or storage node" on page 70 provides information about restarting Sun StorageTek EBS services.

Example 42 Enforcing strong authentication on selected hosts

In this example, a Sun StorageTek EBS 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 Sun StorageTek EBS 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 Sun StorageTek EBS services for the Sun StorageTek EBS server. "Stopping and starting a Sun StorageTek EBS server, client, or storage node" on page 70 provides information about restarting Sun StorageTek EBS services.

Maintaining Sun StorageTek EBS local host authentication credentials

Each Sun StorageTek EBS host that uses **nsrauth** strong authentication has unique credentials that are used to identify itself to other Sun StorageTek EBS hosts during the **nsrauth** strong authentication process. These credentials are known as local host authentication credentials.

Additionally, each Sun StorageTek EBS host maintains a copy of the local host credentials that belong to every Sun StorageTek EBS 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 659 provides details.

Example 43 Securing the initial peer host authentication process

When two Sun StorageTek EBS 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 Sun StorageTek EBS server *jupiter* will authenticate with the Sun StorageTek EBS 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 666 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 669 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 44 Local host credentials are exported

In most cases, local host credentials are retained. For example, when the Sun StorageTek EBS 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 Sun StorageTek EBS 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. If necessary, re-import 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 666 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 667 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 45 Local host credentials are not exported

This scenario is similar to "Local host credentials are exported" on page 665, 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 Sun StorageTek EBS host. "Creating new local host certificate keys" on page 667 provides details.
2. Delete all peer certificates that are set up for the Sun StorageTek EBS host whose credentials were created in step 1 . "Deleting Sun StorageTek EBS local host Peer resources" on page 668 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 Sun StorageTek EBS 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 Sun StorageTek EBS 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 Sun StorageTek EBS 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 Sun StorageTek EBS host.
7. On all other local host resources, delete the Peer resource that was set up for the Sun StorageTek EBS host whose certificate was changed in this procedure. "Deleting Sun StorageTek EBS local host Peer resources" on page 668 provides details.

Note: Consider exporting the new certificate information to a credentials file. "Local host credentials are exported" on page 665 provides information about the benefits of doing so.

Maintaining local host Peer resources

Each Sun StorageTek EBS host has unique authentication credentials that are used during the **nsrauth** strong authentication process. These credentials include certificate information.

Each Sun StorageTek EBS host retains a record of the certificate of all Sun StorageTek EBS hosts to which it has successfully authenticated. These records are maintained as part of each Sun StorageTek EBS host's Peer resource database.

Deleting Sun StorageTek EBS local host Peer resources

When the credentials for a Sun StorageTek EBS host change, so does its certificate. All hosts that have previously authenticated with the changed Sun StorageTek EBS 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 Sun StorageTek EBS host whose Peer resource must be deleted.
4. Right-click the Peer resource that corresponds to the Sun StorageTek EBS 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 Sun StorageTek EBS host whose credentials have changed, if this has not already been done. "Exporting local host credentials" on page 666 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 Sun StorageTek EBS host whose credentials have changed. "Securing the initial peer host authentication process" on page 663 provides details.

Creating a custom certificate and private key for a host

Sun StorageTek EBS automatically creates certificate and private keys for each Sun StorageTek EBS 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 this directory:

   Sun StorageTek EBS_install_path/nsr/res

   If the Sun StorageTek EBS client service, **nsrexecd**, has already been started on the host:

   a. Start the **nsradmin** program:

      nsradmin -p nsrexec

      The **nsradmin** prompt appears.

   b. Type this command:

      . type: nsrla

   c. Import the custom certificate and private key file:

      update nw instance info operations: import;

      nw instance info file: *Sun StorageTek EBS_install_path/nsr/res/certificate_file*

   d. 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 Sun StorageTek EBS 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 Sun StorageTek EBS 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 180 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 Sun StorageTek EBS 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 Sun StorageTek EBS 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.

Adding or changing the Sun StorageTek EBS servers for a client

Sun StorageTek EBS uses the contents of the /nsr/res/servers (UNIX), or the *Sun StorageTek EBS_install_path*\res\servers (Windows) file on each Sun StorageTek EBS client to control who has client-tasking rights (the right to request a program to be executed on another client). This client-tasking might be any of the following:

◆ Server that performs an archive request

◆ Scheduled backup

◆ Another client that requests a directed recover

During a Sun StorageTek EBS software installation, you can add the names of Sun StorageTek EBS servers to this file. To add additional hosts at a later date, a text editor must be used to add the hostname to the file.

◆ So that the client with the tasking rights can back up to other Sun StorageTek EBS servers, the names of the additional Sun StorageTek EBS servers must be added to this file.

◆ So that other clients can perform directed recovers to the client with the tasking rights, their names must be added to the servers file.



IMPORTANT

If the servers file is empty, then any Sun StorageTek EBS host has tasking rights. This is a potential security concern.

Editing the servers file

To give clients client-tasking rights to a Sun StorageTek EBS client:

1. Stop the Sun StorageTek EBS host whose tasking rights are being updated. "Stopping and starting a Sun StorageTek EBS server, client, or storage node" on page 70 provides more information.
2. Open the servers file in a text editor.

   The default installation location for this file is:

   - /nsr/res/servers (UNIX)
   - *Sun StorageTek EBS_install_path*\res\servers (Windows)
3. Enter one hostname per line.
4. Save the changes and exit the text editor.
5. Restart the Sun StorageTek EBS host.

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 Sun StorageTek EBS 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 78 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.

19 Cluster Support

This chapter covers these topics:

- Cluster terminology 678
- Cluster-aware and cluster-unaware Sun StorageTek EBS servers ... 678
- Cluster licensing 680
- Backing up data on UNIX cluster environments 680
- Configuring a virtual client to back up to a local storage node 684
- Configuring a virtual client to be recovered from a local storage node 685
- Configuring an external client to a virtual server 686
- Recovering data in a UNIX cluster environment 687
- Tracking scheduled saves 688
- Using a physical storage node and virtual server on UNIX platforms 691
- Using autochangers and stand-alone tape devices 692
- Client Backup Configuration Wizard support requirements 704
- NetWorker cluster performance issues 706
- Setting NetWorker environment variables in a Sun cluster..... 708

Cluster terminology

This section defines terms specific to Sun StorageTek EBS cluster support that are used in this document:

- Node — A physical computer that is a member of a cluster.
- Private disk — A local disk on a cluster node. A private disk is not available to other nodes within the cluster.
- Shared disk — A disk that is shared by the cluster nodes.
- Failover — A means of ensuring application availability by relocating resources in the event of a hardware or software failure. Two-node failover capability allows operations to switch from one cluster node to the other. Failover capability can also be used as a resource management tool.
- Highly available application — An application installed in a cluster environment and configured for failover capability.
- Cluster-aware Sun StorageTek EBS server — A Sun StorageTek EBS server that can recognize cluster clients.
- Cluster client — A Sun StorageTek EBS client within a cluster. This can be either a virtual client, or a Sun StorageTek EBS Client resource that backs up the private data that belongs to a one of the physical nodes.
- Virtual client — A Sun StorageTek EBS Client resource that backs up data that belongs to a highly available service or application within a cluster. Virtual clients can fail over from one cluster node to another.
- Stand-alone server — A Sun StorageTek EBS server that runs within a cluster, but not configured as a highly available application. A stand-alone server does not have failover capability.
- Virtual server — A Sun StorageTek EBS server configured to run as a highly available application. A virtual server can fail over from one cluster node to another.

Cluster-aware and cluster-unaware Sun StorageTek EBS servers

To perform scheduled backups within a cluster, you must have a cluster-aware Sun StorageTek EBS server and Sun StorageTek EBS

client. Not all Sun StorageTek EBS servers and clients are cluster-aware.The *EMC NetWorker Software Compatibility Guides* provides information about which versions are cluster aware.

These issues apply if either the server or client are not cluster-aware:

- If a Sun StorageTek EBS *server* does not have cluster-client support, it cannot recognize cluster clients. When a scheduled backup is performed, the Sun StorageTek EBS server backs up only the private data of the physical node. The data is backed up to the same index representing the physical node.
- If a Sun StorageTek EBS *client* does not have cluster-client support, it backs up all data that belongs to the virtual servers and physical nodes into the same index representing the physical node. It does not recognize the difference between virtual and physical clients.

Cluster licensing

Chapter 28, "Licensing" provides information on Sun StorageTek EBS licensing support for clusters.

Backing up data on UNIX cluster environments

This section describes how to configure scheduled saves or data from a physical cluster client and a virtual cluster client, as well as how to perform a manual backup of data from a virtual cluster client.

These topics provide more information:

- "Configuring a scheduled save from a physical cluster client" on page 680
- "Configuring a scheduled save from a virtual cluster client" on page 680
- "Performing a manual backup from a virtual cluster client" on page 682

Configuring a scheduled save from a physical cluster client

Backing up data from a physical client within a cluster is similar to backing up data on a computer that is not part of a cluster. Chapter 2, "Backing Up Data" provides detailed instructions on how to back up data.

To prepare for regular backups, install the Sun StorageTek EBS client software on all physical clients within the cluster. The *Sun StorageTek EBS Installation Guide* provides detailed installation and configuration instructions.

Configuring a scheduled save from a physical cluster client on HACMP

Configuring a scheduled save from a virtual cluster client

Instructions for how to configure a scheduled save from a virtual cluster client are the same whether the Sun StorageTek EBS server is inside the cluster or outside the cluster.

To configure a scheduled save from a virtual cluster client under the Sun StorageTek EBS server:

1. (Optional) Define save groups as required.

   Note: Save groups must have the **Autorestart** attribute enabled and the **Manual Restart** option disabled in order for their save sets to restart after a virtual client or Sun StorageTek EBS server failover.

2. Make each physical client within the cluster a Sun StorageTek EBS client of the Sun StorageTek EBS server.

   For each physical client in the cluster:

   a. Create a new Sun StorageTek EBS Client resource.

   b. For the **Name** attribute, enter the name of the physical client.

3. Make each virtual client within the cluster a Sun StorageTek EBS client of the Sun StorageTek EBS server.

   For each virtual client in the cluster:

   a. Create a new Sun StorageTek EBS client.

   b. For the **Name** attribute, enter the name of the virtual client.

   c. For the **Remote Access** attribute, add entries for each physical client within the cluster. For example:

      ```
      root@clus_phys1
      ```

   d. For the **Group** attribute, select a group.

      If you are using virtual Sun StorageTek EBS server that resides in a cluster, the first time the Sun StorageTek EBS server runs, it creates the Client resource for the virtual Sun StorageTek EBS server.

4. Run a test probe to verify that the Client and Group resources are properly configured.

   On the cluster node or stand-alone server on which the Sun StorageTek EBS server resides, enter this command:

   savegrp -pv -c *virtual_client group_name*

Note: "Tracking scheduled saves" on page 688 provides information on if the test probe does not display the appropriate scheduled backups and index.

Performing a manual backup from a virtual cluster client

Backing up data from a virtual cluster client is similar to backing up data from noncluster clients. For example:

```
save -c client saveset
```

Performing a manual backup by using nonroot administrative privileges

For all cluster platforms except HP MC/ServiceGuard, AutoStart, and VERITAS Cluster Server, performing a manual backup by using nonroot privileges is similar to backing up data from noncluster clients by using root privileges. For example:

```
save -c client saveset
```

HP MC/ServiceGuard

To perform a manual backup of data from a virtual cluster client with nonroot privileges on HP MC/ServiceGuard:

- On each node on which you want to run the **save** command, edit or create the .rhosts file in your home directory to include the names of all the nodes within the cluster. For example:

  nodeA

  nodeB

- As root on all nodes in the cluster:

  a. Edit or create the /etc/cmcluster/cmclnodelist file and add this information to the file:

  nodeA *user_name*

  nodeB *user_name*

  b. Copy the file to each node in the cluster.

Note: If the cmclnodelist file exists, the cluster software ignores any .rhosts file.

AutoStart for AIX, HP-UX, Linux, and Solaris

To perform a manual backup of data from a virtual cluster client with nonroot privileges on AutoStart for AIX, HP-UX, Linux, and Solaris, use the **ftconsole** program to add nonroot users to the security access list for the FT domain.

VERITAS Cluster Server

You can perform a manual backup of cluster data that belongs to a physical or a virtual cluster client with nonroot privileges by using

VERITAS Cluster Server (VCS) version 4.0 or lower. To do so, set the server's AllowNativeCliUsers attribute to **1**. Setting this attribute ensures that the nonroot user is not prompted for a password when executing the server's commands. This attribute instructs the server to authenticate by using the user's operating system username and password, rather than the server username and password. The server default Cluster Guest permissions are sufficient to use the Sun StorageTek EBS client software.

In VCS version 4.1 or later, the AllowNativeCliUsers attribute has been deprecated. To perform a manual backup of cluster data that belongs to a physical or a virtual cluster client with nonroot privileges by using VCS version 4.1 or later, use the VCS **halogin** command. For more information about the **halogin** command, see the VCS documentation.

Note: When a nonroot user is running a NetWorker program such as **save** or **nsradmin** on a VCS cluster, the user must have permissions to query VCS. Otherwise, the NetWorker program will not be cluster-aware because the user is not authenticated by the VCS cluster. For information on how to set up VCS authentication, see the VCS documentation.

HACMP for AIX

To back up a physical client on an HACMP 4.5 node, each node requires persistent IPs or an extra NIC that is configured outside of the control of the HACMP 4.5 environment. For a physical client, the Sun StorageTek EBS software requires an address that uniquely connects to a physical client.

The service and boot addresses of HACMP 4.5 do not meet this requirement because clusters configured with IP address takeover (IPAT), replace the boot address with the service address when a resource group is attached.

Service addresses are associated with a resource group, not physical nodes. The output of the hostname command on a computer must correspond to a pingable IP address. The computer hostname must also be set to the name equivalent of the address used by the physical client's persistent IP or dedicated NIC. With either method, persistent IP or dedicated NIC, the primary network adapter must be used (for example, en0).

Configuring a virtual client to back up to a local storage node

Typically, the Sun StorageTek EBS server backs up the data on a virtual client within a cluster to the first storage node listed in the virtual client's Storage Nodes attribute.

You can configure a virtual client within a cluster to direct its backups to the storage node on the physical host on which it resides. To do this, type the following in the Storage Nodes attribute of the virtual client:

```
curphyhost
```

Note: curphyhost is only applicable to virtual clients and should never be inserted into the clone storage nodes list even for the NetWorker Server's NSR resource.

This is done by typing curphyhost in the Storage Nodes attribute of the virtual client.

For example, consider a two-node cluster where:

- Nodes A and B are the two physical nodes in the cluster.
- The virtual client is saturn, which can reside on Node A or fail over to Node B.

In a normal Sun StorageTek EBS backup of saturn—without "curphyhost" listed in the virtual client's Storage Nodes attribute—the save data is directed to the remote device (rd=) on Node A. When saturn fails over to Node B and a backup for saturn is initiated, the save data is still directed to the remote device (rd=) on Node A.

Listing "curphyhost" first in saturn's Storage Nodes attribute modifies the Sun StorageTek EBS operation so that if saturn fails over to Node B and a backup of saturn is initiated, the save data is directed to the remote device (rd=) on Node B. This action takes place because, after the failover, saturn resides on Node B—the current physical host.

Note: Do not apply the curphyhost keyword to your virtual server's Client resource. This can cause unexpected behavior, such as bootstraps and online indexes being written to the local storage node for the virtual clients, rather than to a local device on the server.

Directing a virtual client backup to a local storage node

To direct a virtual client to back up to a local storage node:

1. Start the **Sun StorageTek EBS Administration Console** software.
2. From the **Administration** window, click **Configuration**.
3. In the left pane, select **Clients**.
4. From the **File** menu, select **Properties** and then select the **Globals 2 of 2** tab.
5. For the **Storage** nodes attribute, add the **curphyhost** keyword. Position the keyword in the list based on the required priority. The top of the list is the highest priority.

Configuring a virtual client to be recovered from a local storage node

The procedure to configure a virtual client to be recovered from a local storage node is similar to the procedure describe in "Configuring a virtual client to back up to a local storage node" on page 684, except that this features is specific to a recover operation.

To configure a virtual client to be recovered from a local storage node, type the following in the Storage Nodes attribute or Recover Storage Nodes attribute of the virtual client:

```
currechost
```

With this keyword set, a request is made that the required volume be mounted on the storage node where the recover operation is running.

The following restrictions apply when configuring a virtual client to be recovered from a local storage node:

- Ensure that there are no hosts or machines named *currechost* on the network.
- The currechost keyword can not be set in the Clone storage nodes attribute of a virtual client.
- Do not apply the currechost keyword to the virtual server's Client resource.

Directing a virtual client recovery from a local storage node

To direct a virtual client to be recovered from a local storage node:

1. Start the **Sun StorageTek EBS Administration Console** software.
2. From the **Administration** window, click **Configuration**.
3. In the left pane, select **Clients**.
4. From the **File** menu, select **Properties** and then select the **Globals 2 of 2** tab.
5. For the **Storage nodes** attribute or the **Recover storage nodes** attribute, add the **currechost** keyword. Position the keyword in the list based on the required priority. The top of the list is the highest priority. Ensure that this keyword is not the only keyword in the list.

Recovering a virtual client's data by using the currechost keyword

To recover a virtual client's data by using the currechost keyword, run the recover operation on the node in the cluster where the virtual client resides. The NetWorker server requests that the volume be mounted on the local storage node.

Be aware of the following considerations when recovering the virtual client's data:

- Ensure that the local storage node has access to the required volume.
- If the required volume is in use by another drive, the recover operation waits until the volume is unmounted. When the required volume becomes available, the volume is mounted and the recover operation continues.

Configuring an external client to a virtual server

Before a Sun StorageTek EBS server can back up a client, the client must grant the server access. Granting access is controlled by these files:

- On UNIX: /nsr/res/servers
- On Windows: <*Sun StorageTek EBS_install_path*>\res\servers

Only the specific servers listed in this file can back up the client. "Adding or changing the Sun StorageTek EBS servers for a client" on page 673 provides more information about the servers file.

Recovering data in a UNIX cluster environment

This section discusses how to recover data from private disks belonging to a physical client and shared disks belonging to a virtual client.

Recovering data from a private disk that belongs to a physical client

Recovering data backed up from a *private disk* that belongs to a physical client in a cluster is similar to recovering data on a computer that is not part of a cluster.

Chapter 13, "Recovering Data" provides detailed instructions.

Recovering data from a shared disk that belongs to a virtual client

To recover data backed up from a *shared disk* that belongs to a virtual client:

1. Ensure that you have properly configured the virtual client for remote access from your current node:
   a. From the **Administration** window, click **Configuration**.
   b. In the left pane, select **Clients**.
   c. From the **File** menu, select **Properties** and then select the select the **Globals 2 of 2** tab.
   d. Make sure that the **Remote Access** attribute of the virtual client for which you want to recover the data contains the name of the physical host you are using to recover the data. For example:

      root@*physical_hostname*
2. Make sure that the file systems of the virtual client are mounted.
3. Recover data from a shared disk that belongs to a virtual client.
   - To recover the data by using the **nwrecover** program, make sure that you select the virtual client.

- To perform a command line recovery, enter the **recover** command with the **-c** option at the command prompt. The **-c** option must be followed by the name of the client you are attempting to recover. For example:

  ```
  recover -s server_name -c virtual_client
  ```

The **recover** man page or the *EMC NetWorker Command Reference Guide provide information*.

Recovering the Console server database

The Console database contains management data such as report information. The *EMC NetWorker Disaster Recovery Guide* provides information about recovering a corrupted Console database.

Tracking scheduled saves

To verify that the Client and Group resources are properly configured, run a test probe for each client from the node where the Sun StorageTek EBS server is running:

```
savegrp -pv -c client_name group_name
```

If the test probe does not display all the scheduled save sets you want, do one of these and run a second test probe to verify the configuration:

- Check the cluster configuration by using the cluster management software. If necessary, reconfigure the cluster.
- Make sure that the client owns the save sets defined for that client. If necessary, redistribute the client save sets to the appropriate Client resources.

Note: Misconfiguration of the cluster resources might cause scheduled save sets to be dropped from the backup. "Path-ownership rules" on page 689 provides more information about configuring clusters for scheduled backups .

To override scheduled save rules, create an empty file named pathownerignore on each node in the cluster by running this command:

```
touch /usr/sbin/nsr/pathownerignore
```

This allows any path to be backed up for a client, whether it is owned by the virtual or physical node.

Note: Use the pathownerignore file with caution.

The pathownerignore file does not override the default path-ownership rules. It causes the path-ownership rules to be ignored when determining if a file system should be backed up during a scheduled save.

Forcing save sets to use a specific index

If you use pathownerignore, check that the Sun StorageTek EBS scheduled save uses the correct client file index. If it uses the wrong index, you can force save sets to go to a specific index.

To force save sets to use a specific index:

1. From the **Administration** window, click **Configuration**.
2. In the left pane, select **clients**.
3. In the right pane, select the client that is using the incorrect index.
4. From the **File** menu, select **Properties** and then select the **Apps & Module** tab.
5. For the **Backup Command** attribute, enter the name of a backup script that contains **save -c** *client_name*.

"Using the save command with a customized backup script" on page 110 provides detailed information about the Backup Command attribute.

Path-ownership rules

In a cluster environment, the Sun StorageTek EBS software must distinguish between:

- File systems associated with a physical client.
- File systems managed by a resource group (a virtual client).

The criteria used are called the path-ownership rules. These rules determine which client file index a save set is written to.

By default, the Sun StorageTek EBS software might not back up scheduled save sets due to conflicts with path ownership. This

prevents a virtual Sun StorageTek EBS client from writing to multiple client file indexes. The Sun StorageTek EBS software views the client (which owns the file system) as not having matched the client of the current scheduled save set. Writing to multiple indexes might cause recovery issues.

These conditions cause a file system to be omitted and *not backed up* during a scheduled save:

- If a file system owned by a virtual client is defined in the Save Set attribute for a physical Client resource.
- If a file system owned by a physical client is defined in the Save Set attribute for a virtual Client resource.

Administrative privileges

A user must have administrative privileges to query the cluster database. If a nonadministrative user performs a save of a file system that belongs to a virtual node, path ownership resolution fails and the files are saved under the physical node's client file index, instead of the virtual node's index. To force the Sun StorageTek EBS software to save the files under the virtual node's client file index, use the **-c** option. "Forcing save sets to use a specific index" on page 689 provides details.

Checking path-ownership rules

To check the Sun StorageTek EBS path-ownership rules:

1. Type this command at the command-prompt on the Sun StorageTek EBS server:

   ```
   savegrp -p -c client_name
   ```

2. Review which file systems *client_name* owns. This procedure is part of the normal cluster installation setup. The *Sun StorageTek EBS Installation Guide* provides details.

   To test for misappropriated save sets, run a test probe with the **verbose** option set. The command output indicates to which client file indexes a save set is saved. For example:

   ```
   savegrp -pv -c client_name group_name
   ```

To ignore Sun StorageTek EBS default path-ownership rules, you can create the pathownerignore file in the NetWorker bin directory:

/usr/sbin/nsr/bin

This file causes the Sun StorageTek EBS software to back up the file system in question. However, the file system is saved under the client file index of its correct owner.

"Forcing save sets to use a specific index" on page 689 provides information if a situation occurs when the scheduled save uses the wrong index.

Overriding path-ownership rules

Overriding the path-ownership rules causes a save set to be saved under a client file index different from the save set's default owner.

How to override path-ownership rules during a manual backup

To override the default path-ownership rules for a manual backup, type the following at the command-prompt:

```
save -c client_name
```

How to override path-ownership rules during a scheduled backup

To override the path-ownership rules during a scheduled backup:

1. From the **Administration** window, click **Configuration**.
2. In the left pane, select **Clients**.
3. In the right pane, select the client that is using the incorrect index.
4. From the **File** menu, select **Properties** and then select the **Apps & Module** tab.
5. For the **Backup Command** attribute, enter the name of a backup script that contains **save -c** *client_name*.

"Using the save command with a customized backup script" on page 110 provides detailed information about the Backup Command attribute.

Using a physical storage node and virtual server on UNIX platforms

This section describes how to access a physical storage node with a virtual Sun StorageTek EBS server.

A Sun StorageTek EBS storage node can be installed and configured on nodes in the cluster. However, the Sun StorageTek EBS storage node is not a highly available application. These instructions apply in

the cases where you are using storage nodes in the cluster and storage nodes that are outside of the cluster.

Note: Sun StorageTek EBS release 6.0 and later does not officially support remote virtual storage nodes.

To access a physical Sun StorageTek EBS storage node from a virtual server (rd=*physicalhostname:devxxx*):

1. Shut down the virtual Sun StorageTek EBS server by using the cluster management software.
2. Modify the Sun StorageTek EBS virtual cluster **start** script to recycle the remote Sun StorageTek EBS storage node.

   On each node to which the virtual Sun StorageTek EBS server can fail over:

   a. Open the Sun StorageTek EBS virtual cluster script, ebs_bin_dir, for editing.

   b. Before the line that starts the **nsrd** daemons, add these:

   rsh remote_hostname "ebs_bin_dir**/nsr_shutdown -q &&**

   ebs_bin_dir**/nsrexecd"**

   Note: Replace the ebs_bin_dir variable with the appropriate path where the Sun StorageTek EBS binaries are installed on the Sun StorageTek EBS storage node.

   c. If necessary, resynchronize the virtual Sun StorageTek EBS server **startup** script with the cluster management system.
3. Restart the Sun StorageTek EBS server.

Using autochangers and stand-alone tape devices

This section describes how to use autochangers and stand-alone tape devices with a highly available Sun StorageTek EBS server. Device-sharing infrastructure is defined as hardware, firmware, and software that permit several nodes in a cluster to share access to a device.

Configurations for autochangers and stand-alone tape devices

If the cluster has a device-sharing infrastructure that supports autochangers and stand-alone tape devices. To configure your system to use:

- Shared autochanger devices and shared stand-alone tape devices, use "Configuration option 1" on page 693.
- Shared autochanger devices, but *not* shared stand-alone tape devices, use "Configuration option 2" on page 695.
- Nonshared devices, or any configuration, use "Configuration option 3" on page 697, which does not require device-sharing infrastructure.

Note: Make sure that the device-sharing infrastructure is certified by the cluster software vendor.

Configuration option 1

This configuration option offers a seamless backup solution that is hands-free. It allows you to access shared tape devices and shared autochanger devices with a highly available Sun StorageTek EBS server. Figure 34 on page 694 provides a graphical view of this configuration option.

Note: Only one autochanger is required, although you can add more.

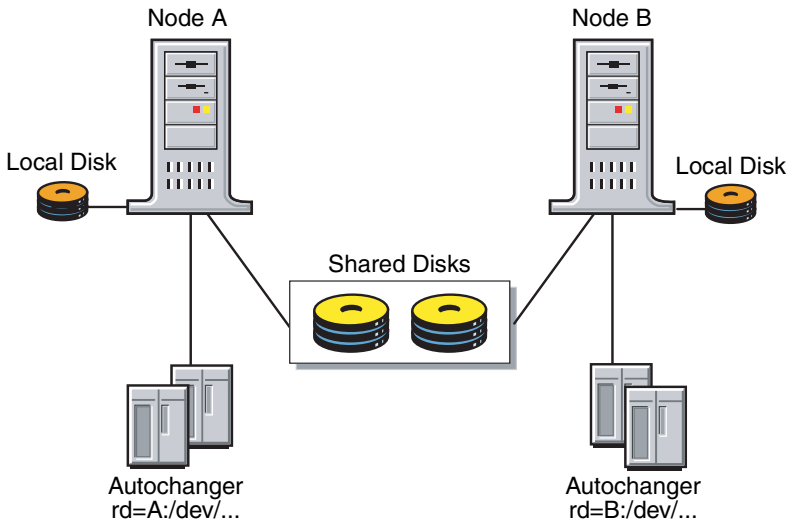


Figure 34 **Configuration option 1**

To configure shared tape devices and shared autochanger devices to work with a highly available Sun StorageTek EBS server:

1. Interface all the autochanger devices to all the nodes in the cluster to which the Sun StorageTek EBS server can fail over. Use a working SCSI or Fibre Channel switch, and appropriate software as required.
2. Make sure that the shared device infrastructure (SAN bridge, modular data router) supports complete isolation and protection of the path session between the autochanger and the node with the Sun StorageTek EBS server currently running. The path must be protected from stray bus signals and unauthorized session access from the other nodes.

   Note: If processes on nodes other than the one that is running on the Sun StorageTek EBS server can access the tape devices, data corruption might occur. The Sun StorageTek EBS software might not detect the data corruption.

3. Modify the virtual Sun StorageTek EBS server startup (/usr/sbin/networker.start) script by adding any required node device-reassignment reservation commands (commands for the bridge) before the line that starts the **nsrd** daemon.

   Note: If you modify the **networker.start** script, it is no longer a Sun Cluster certified failover solution.

4. Test that the device-sharing infrastructure supports path . isolation. For UNIX platforms, use the appropriate operating system commands for your platform against the device. For example, **tar** or **cpio**:

   a. From Node A, write a significantly large amount of data to tape. During the write session, turn node B on and off a couple of times.

   b. Verify that the data was properly written by reading the data from tape and comparing it to the original data. Table 104 explains the verification results.

Table 104 Data verification

Indication	Description
Failure	The switching infrastructure does not properly support path isolation. Try "Configuration option 2" or "Configuration option 3", or contact the device-sharing infrastructure vendor.
Success	Reverse the node roles: • From Node B, write a significantly large amount of data to tape. During the write session, turn node A on and off a couple of times. • Verify that the data was properly written by reading the data from tape and comparing it to the original data. • If the verification fails, try "Configuration option 2" or "Configuration option 3", or contact the device-sharing infrastructure vendor.

5. Configure the autochanger and tape devices.

Configuration option 2

This option offers a seamless backup solution for accessing nonshared tape devices and shared autochanger devices from a highly available Sun StorageTek EBS server. However, some intervention might be required during the data recovery process. Figure 35 on page 696 provides a graphical view of this configuration option.

Note: If *N* nodes are used, *N-1* storage node licenses are required. *N* indicates the number of failover nodes for the virtual Sun StorageTek EBS server.

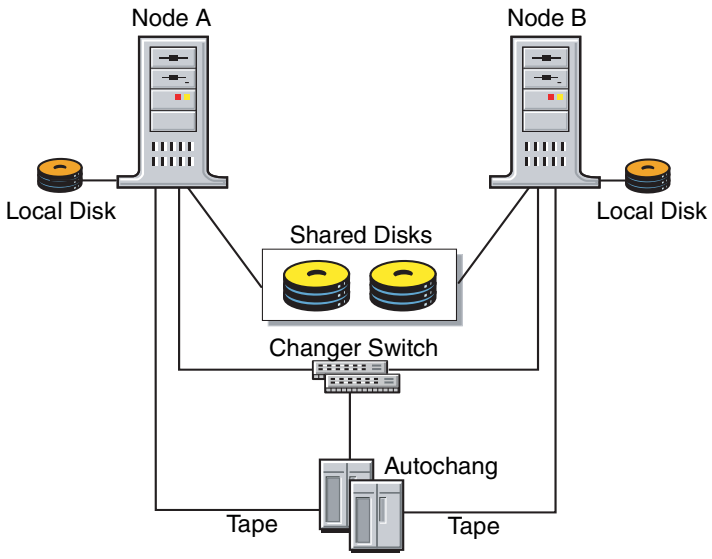


Figure 35 **Configuration option 2**

To configure the devices to work with a virtual Sun StorageTek EBS server:

1. Interface the autochanger device to the Sun StorageTek EBS failover nodes. Use a working SCSI or Fibre Channel switch and appropriate software if required.
2. Privately attach each tape device to a separate node within the cluster. For example:
   - Attach tape drive 1 exclusively to Node A.
   - Attach tape drive 2 exclusively to Node B.
3. Configure the tape devices as remote physical storage nodes. For example:

```
rd=nodeA:/dev/tape#1
rd=nodeB:/dev/tape#2
```

Note: When defining each tape device, assign ownership to the node exclusively connected to that tape device.

Configuration option 3

This option offers a seamless backup solution for accessing nonshared tape devices and nonshared autochanger devices from a highly available Sun StorageTek EBS server. However, some intervention might be required during the data recovery process. Figure 36 on page 697 provides a graphical view of this configuration option.

Note: *N* indicates the number of failover nodes for the virtual Sun StorageTek EBS server.

- If *N* nodes are used, *N-1* storage node licenses are required.
- At least *N* autochangers are required, one per failover node.

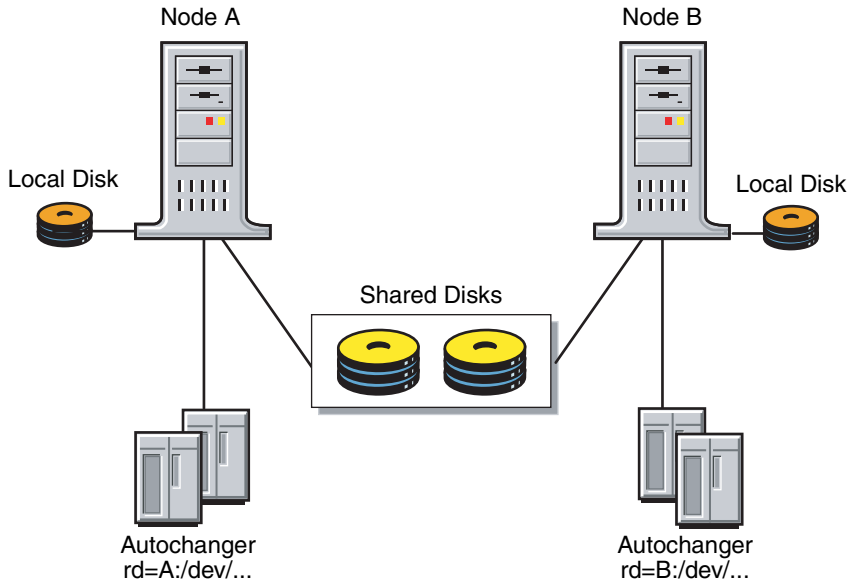


Figure 36 Configuration Option 3

To configure tape autochanger devices to work with a virtual Sun StorageTek EBS server:

1. Attach the autochanger exclusively to the Sun StorageTek EBS server.
2. Use this command to configure the nonshared tape and nonshared autochanger devices as remote physical storage node devices. For example:

```
jbconfig -s virtual_server
rd=nodeA:/dev/tape#1
rd=nodeB:/dev/tape#2
rd=nodeA:/dev/changer
```

3. Repeat these steps on each Sun StorageTek EBS server failover node.

DDS within a cluster

There are two options for using the DDS feature within a cluster:

- "Controlling the robotics inside the cluster configuration" on page 699
- "Controlling the robotics from a stand-alone host" on page 701

"Dynamic drive sharing" on page 132 provides information explaining how to activate the DDS feature.

These sections explain the supported cluster platforms and the two methods for using DDS within a cluster.

High availability and DDS

The Sun StorageTek EBS software relocates and restarts operations that were in progress when a failure occurs on a cluster node. Currently, the only operations that are highly available are the Sun StorageTek EBS server and savegroups.

The nsrjb program high availability limitations

If the Sun StorageTek EBS server fails over from one node to a new target node, standard jukebox operations (such as performing an inventory, labeling, mounting or unmounting a volume) will not automatically restart on the new target node.

Example 46 Host crash requires user intervention

You have this scenario: two physical hosts, A and B, with DDS enabled, sharing the drives on an autochanger.

Physical host A mounts a tape in a shared drive on the autochanger. If physical host A subsequently crashes, the volume is held in that shared drive until you issue the **nsrjb -H** reset command (from host B, in this example).

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 Sun StorageTek EBS software is able to access the same path through which the initial loading of the volume was completed.

Controlling the robotics inside the cluster configuration

To use DDS in a cluster, with the robotics controlled within the cluster configuration:

- One node must be a virtual Sun StorageTek EBS server. This virtual Sun StorageTek EBS server controls the robotics from within the cluster.
- Physical storage nodes of the cluster are configured to share the same drive.
- You must run the **jbconfig** program on the Sun StorageTek EBS server to install the autochanger.

Note: To configure a highly available cluster, the special device file for the library robotics must be installed on all the physical nodes in the cluster. The special device file for the library robotics is provided by the cluster vendors. You must use the same naming convention across all the cluster nodes. Some operating systems already have the special device file for the library robotics bound in. Essentially, these special device files represent a pointer to a generic tape or autochanger device driver. To determine whether the special device file has been bundled in with your operating system, run the **inquire** command.

If the special device files are not bundled in with the operating system, refer to the appropriate *Sun StorageTek EBS Installation Guide*.

Note: If you do not have matching special device files across cluster nodes, you may need to install fibre HBAs in the same PCI slots on all the physical nodes within the cluster.

With most cluster environments, the HBA can be configured to create the same scsi/FC path on each physical node of the cluster to make the Sun StorageTek EBS robotic control highly available. As a result, the scsidev@bus.target.lun will be identical for the robotics from each node in the cluster. Figure 37 illustrates how DDS can be enabled within a cluster.

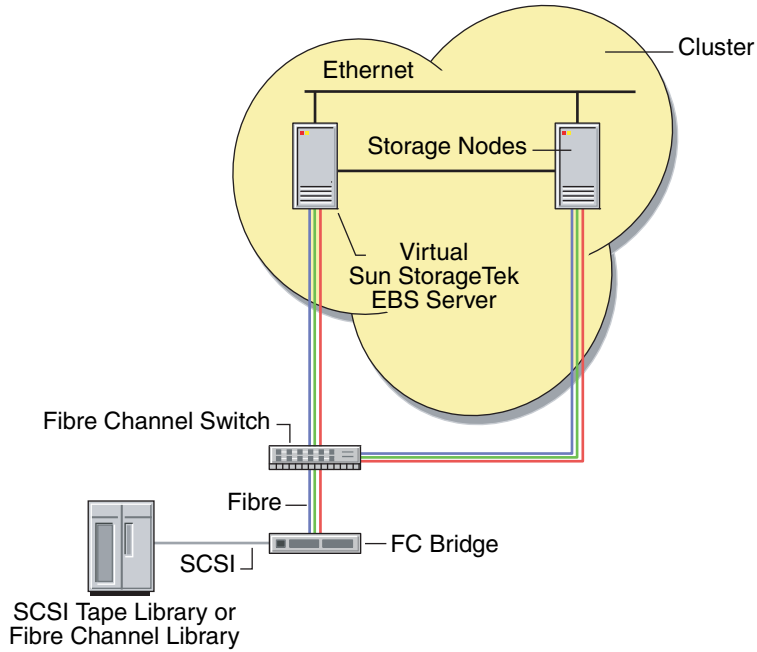


Figure 37 DDS enabled with robotics controlled in the cluster

Using the jbconfig program to control the robotics

To control the robotics, you must use the **jbconfig** program to configure the devices within the cluster. There are two options for using the **jbconfig** program to configure devices with a cluster.

- Using the **jbconfig** program *with* special device files
- Using the **jbconfig** program *without* special device files

The configuration process is similar in both cases, differing only at the prompt:

```
enter the control port of the jukebox in the following
format
```

The prompt at which the response differs is illustrated in Example 47 on page 701.

This is an example of the **jbconfig** script with appropriate responses:

```
jbconfig
1) Configure an AlphaStor/SmartMedia Jukebox.
2) Configure an Autodetected SCSI Jukebox.
3) Configure an Autodetected NDMP SCSI Jukebox.
4) Configure an SJI Jukebox.
5) Configure an STL Silo.
What kind of Jukebox are you configuring? [1] 4
Enter the number corresponding to the type of jukebox
you are installing:
1) ADIC-1200c/ADIC-1200d
2) ADIC-VLS
3) ARC DiamondBack
4) Sun 20Gb 4mm Tape Loader
.
.
.
45) Digital TL800
46) Digital TL810
47) Digital TL820
48) Digital TL893
Choice? 48
Installing an 'Digital TL893' jukebox.
What name do you want to assign to this jukebox device?
juke_box_1\
```

At this point, the **jbconfig** program differs for configurations without vendor-supplied drivers and for configurations with vendor-supplied drivers.

Example 47 On Sun Cluster with no special device file

Using scsidev@2.3.0. for the device, enter the following when using the **jbconfig** program:

```
Enter the control port of the jukebox in the following
format:scsidev@3.0.0
Pathname of the control port for the jukebox device?
scsidev@2.3.0
```

Controlling the robotics from a stand-alone host

You can set up a stand-alone physical host outside the cluster to control the robotics, as follows:

- If you are unable to get a uniform nomenclature for the vendor-supplied drivers and you cannot match bus target LUNs across the cluster nodes.
- If you do not have a Sun StorageTek EBS server within the cluster.

Controlling the robotics from outside the cluster creates the opportunity for a single point of failure. Therefore, evaluate your requirements if you choose this method.

If you choose to have a stand-alone physical host controlling the robotics, you must ensure that the following binaries are installed:

- Client software
- Storage node software
- Driver software

You would configure the **jbconfig** program on the stand-alone physical host to control the robotics.

Figure 38 on page 703 illustrates one manner in which a stand-alone physical host outside the cluster can be used to control the robotics.

Note: The stand-alone physical host can control the robotics through a Fibre Channel or SCSI connection.

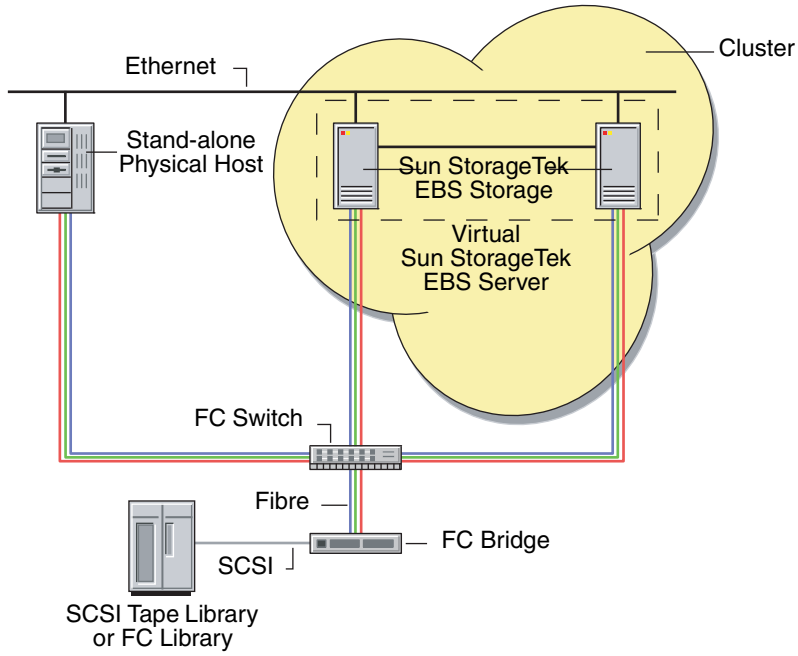


Figure 38 DDS enabled with the robotics controlled outside the cluster

Considerations when using jbconfig for a virtual server

While running the **jbconfig** command with the **-s** *virtual_server* option on the cluster node that is also running the virtual server, the name given to the *Prefix to the jukebox name* attribute has the following implications:

- If the virtual server name (default) is used for the Prefix to the jukebox name attribute, the jukebox's devices will be local to the virtual server.
- If the physical cluster hostname is used for the Prefix to the jukebox name attribute, the jukebox's devices will be local to the physical cluster host.

Controlling robotics from a Sun StorageTek EBS server on another cluster

A Sun StorageTek EBS server in a cluster can also be configured to control the robotics of the nodes of another cluster. This configuration is illustrated in Figure 39. The virtual Sun StorageTek EBS server of cluster B controls the robotics for the physical storage nodes in cluster A.

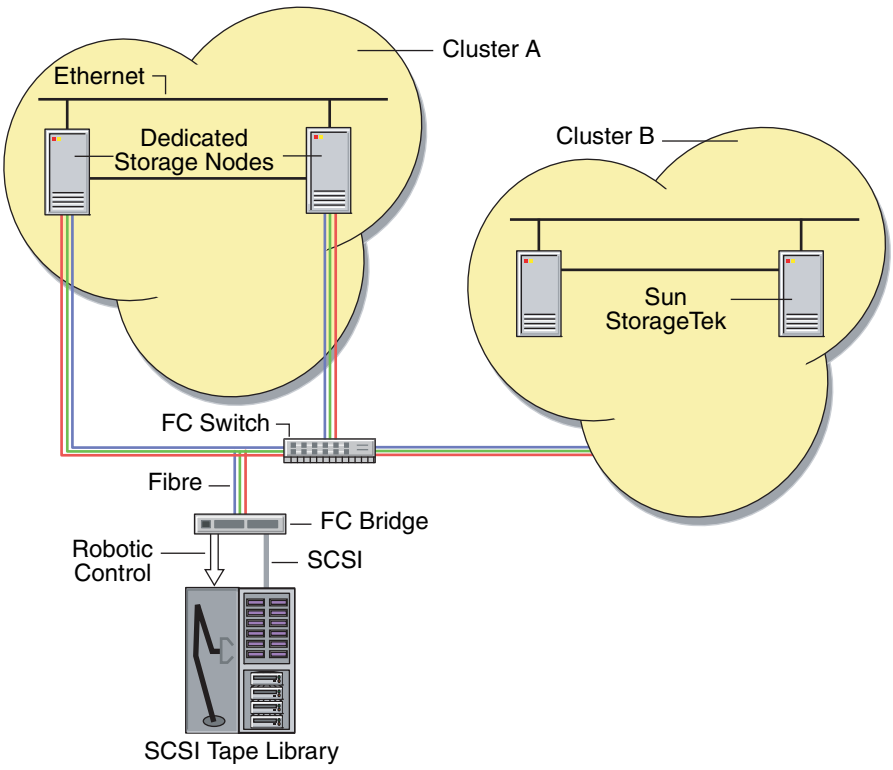


Figure 39 Robotics controlled between clusters

Client Backup Configuration Wizard support requirements

This section contains the requirements to use the Client Backup Configuration Wizard in a cluster environment.

- On a highly available Sun StorageTek EBS Server, the shared servers file must be updated with the nsrwizcon.res and nsrwizclnt.res files from the local servers file:

- On UNIX:

  cp */nsr.NetWorker.local/res/nsrwiz*res /nsr/res*

- On Windows:

  copy c:*program files**legato**nsr**res**nsrwiz*res*
  <shared drive>: *nsr**res*

◆ The host from which the wizard is being run and the Sun StorageTek EBS server host must be listed in both the local and shared servers file on the client host that is being configured for a scheduled backup.

"Using the Client Backup Configuration Wizard" on page 75 provides additional general considerations that apply to the Client Backup Configuration Wizard.

NetWorker cluster performance issues

The NetWorker program, **lcmap**, queries cluster nodes to creates a map that includes information such as path ownership of resource groups. In large cluster configurations, **lcmap** may take a long time to complete and thus make certain operations occur slowly. This is most often noticed in slow NetWorker server (**nsrd**) start times or in very long backup times. In these situations, consider adjusting the NSRLA attribute named cluster cache timeout. This attribute specifies a time, in seconds, in which to cache the cluster map information on a NetWorker client.

Edit the cluster cache timeout attribute with caution. Values for the attribute can vary from several minutes to several days, depending on how often the cluster configuration changes, the possibility of resource group failover, and the frequency of NetWorker operations. Too large a value can result in an out-of-date cluster map, which can cause incorrect path resolution. For example, if the cluster cache timeout value is set to 86400 (one day), any changes to the cluster map will not be captured for up to one day. If cluster map information changes before the next refresh period, some paths may not be resolved correctly. Also be aware that too small a value could negatively affect performance because the cache is being updated too frequently. Experiment with one physical cluster node to find a satisfactory timeout value. If a significant improvement in performance cannot be obtained by adjusting this attribute, reset the attribute value to 0 (zero) so that the attribute will not be used.

Editing the cluster cache timeout attribute

The cluster cache timeout attribute is available only on NetWorker clients in a VERITAS Cluster or Sun Cluster environment. Additionally, the cluster cache timeout attribute resides in the NSRLA database of the NetWorker client and is visible only when NetWorker is configured for a cluster. For example on UNIX, a NetWorker client is configured for a cluster when the networker.cluster script is executed and **nsrexecd** is restarted.

Note: It is recommended that you edit the initial cluster cache timeout value before bringing the highly available NetWorker server online.

To edit the cluster cache timeout value:

1. Log in as root or as Windows Administrator on a physical node that is running the Sun StorageTek EBS client.
2. Type this at the command prompt:

   nsradmin -p nsrexecd

   The **nsradmin** prompt appears.
3. Display the current settings for attributes in the **NSRLA** resource, for example:

   print type:NSRLA
4. Change the value of the cluster cache timeout attribute:

   update cluster cache timeout:*value*

   where *value* is the timeout value in seconds. A value of 0 (zero) specifies that the cache is not used.

   Note: 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.

   ---
5. Type **Yes** when prompted to confirm the change.
6. Perform the previous steps on each physical node in the cluster.
7. Bring the highly available NetWorker server online.

   Unlike a standalone environment, there are two types of NSRLA databases when a highly available NetWorker server is configured in a cluster: a local NSRLA database on each physical node, and a shared NSRLA database on the highly available NetWorker server's shared disk. The previous steps set the attribute in the local NSRLA database on each physical node. Later when the highly available NetWorker server is brought online for the first time, it will get the cache timeout value from the local NSRLA database on the physical node where the highly available server is running.
8. The following steps are optional and can be used if you want to experiment with different cluster cache timeout values on the highly available NetWorker server.

a. Change the cluster cache timeout value in the shared NSRLA database by issuing the previous **nsradmin** commands (steps 2 to 5) on the physical node where the highly available NetWorker server is running.

b. The updated value will take effect after the next cache update, which is based on the previous timeout value. To make the timeout value take effect immediately, delete the cache file on the physical node. On UNIX, the cache file is located at /tmp/lcmap.out. On Windows, the cache file is located at C:\Program Files\Legato\nsr\bin\lcmap.out.

c. Once you have found a satisfactory value for the highly available server, you can also update the cluster cache timeout value on each physical node.

   To update the local cluster cache timeout value on the physical node that is currently hosting the highly available NetWorker server, specify the location of the NSRLA database when using the **nsradmin** command, for example:

   nsradmin -d *directory*

   where *directory* is the location of the local NSRLA database. On UNIX, the directory is located in /nsr.NetWorker.local/res/nsrladb. On Windows, the directory is located in C:\Program Files\Legato\nsr\res\nsrladb.

Setting NetWorker environment variables in a Sun cluster

In a Sun Cluster environment, the networker.start script is used instead of the /etc/init.d/networker script to start the **nsrd** process.

NetWorker environment variables must be added to the networker.start script instead of /etc/init.d/networker script. Place the environment variables in the networker.start script before **nsrd** is started.

The networker.start script is located in the /usr/lib/nsr directory.

20

Support for VMware

This chapter covers these topics:

- Introduction to VMware support 710
- Licensing Sun StorageTek EBS support for VMware..... 710
- Configuring Sun StorageTek EBS clients for virtual machines . 710
- Automatic discovery of VMware environments 722
- Notifications of changes to VMware environments 723
- Visual Representation of VMware environments 724
- Recovering VMware Consolidated Backups 728

Introduction to VMware support

The Sun StorageTek EBS software provides support for backup and recovery of VMware virtual clients using traditional Sun StorageTek EBS file system backup and recovery or VMware Consolidated Backups (VCB), which provides LAN-free backups of VMware virtual machines. Additionally, the Sun StorageTek EBS software also provides automatic discovery of VMware environments and notification of changes to those environments, and provides both a graphical map and tabular view of VMware environments.

The following sections describe each of these in detail.

Licensing Sun StorageTek EBS support for VMware

Chapter 28, "Licensing" provides information on licensing Sun StorageTek EBS support for VMware.

Configuring Sun StorageTek EBS clients for virtual machines

Sun StorageTek EBS clients can be configured for virtual machines to perform either standard Sun StorageTek EBS file system backups, or VMware Consolidated Backups (VCB).

To configure the Sun StorageTek EBS software to backup a virtual machine you can either:

- Use the Client Backup Configuration Wizard to create a Client resource, or modify an existing Client resource.
- Manually create a new Client resource, or manually modify an existing Client resource.

Considerations for VCB backups

Be aware of the considerations described in this section before implementing

Install VM Consolidated Backup 1.5 Update 1 on a Proxy host

NetWorker 7.6 and later requires that the VM Proxy host has VCB 1.5 Update 1 installed. *VMware Consolidated Backup 1.5 Update 1 Release Notes* provides instructions for upgrading from a previous version of VCB. *VMware Consolidated Backup 1.5 Update 1 Administrators Guide*

provides installation and configuration instructions for a first time VCB installation.

Install NetWorker on the Proxy host

The NetWorker 7.6 client software must be installed on the Proxy host.

Considerations for hotadd transport mode

When the VCB_TRANSPORT_MODE is set to hotadd, the ESX server must be running ESX 3.5 update 4 or later and the proxy host must be running VCB 1.5 Update 1. The value for VCB_PREEXISTING_MOUNTPOINT and VCB_PREEXISTING_SNAPSHOT must be the same because it is not possible to reliably remove a pre-existing snapshot in this case without cleaning up the mount point.

Provide memory for the nsrvcb_save command

When VCB backups run, **nsrvcb_save** uses a large amount of memory. For example, to backup a VM with a 200 GB vmdk, the nsrvcb_save can use up to 1.15 GB of RAM. A minimum of 1 GB or more of RAM should be available for each **nsrvcb_save** binary.

Note: For every VM backup, one **nsrvcb_save** binary runs. Memory consumed by the **nsrvcb_save** binary is released when the VM backup completes.

VM should be powered on for the first backup

When running the first backup, the VM should be powered on whether the VCB_VM_LOOKUP_METHOD is based on the IP address or the display name.

Note: NetWorker shall identify and back up VMware virtual machine clients, configured as part of the VCB, using the IP address of the virtual machine or the VMname as displayed in the VMware VC or ESX server. The virtual machine is identified by its IP address, by default.

Support for directives

The VCB keyword directive is the only global directive that is supported for VCB backups. However, all local directives are supported for ALLVMFS backups and specific save set backups. No directives, global or local, are supported when using the *FULL* keyword directive and the backup level is 0 (full).

Task 1: Configure the Hypervisor resource

If VirtualCenter is configured in the environment, there must be a Hypervisor resource for the VirtualCenter server that hosts the Virtual Machines that use VCB. Before creating a Hypervisor resource for Virtual Center, ensure that the NetWorker client software is installed on the VirtualCenter server to allow the Virtual Map of the environment to be generated with the auto-discovery feature.

Note: The Sun StorageTek EBS software supports auto-discovery of VMware environments with VMware VirtualCenter only. It does not support auto-discovery with an ESX server.

If VirtualCenter is not configured in the environment, there must be a Hypervisor resource created for each ESX server in the environment. It is not necessary to install the NetWorker client on an ESX host that will be configured as a Hypervisor resource.

To configure a Hypervisor resource:

1. From the **Administration** window, click **Configuration**.
2. In the expanded left pane, select **Virtualization**.
3. In the right pane, click **Enable Auto-Discovery**.
4. In the Auto-Discovery dialog box:
   a. In the **Hostname** field, enter the fully qualified domain name (FQDN) or IP address of the VirtualCenter server or the ESX server.
   b. In the **Username** and **Password** fields, enter the credentials required to log onto the server. The username and password must belong to an account that has permission to perform VCB backups, snapshots and registering/creating a new VM.
   c. If the VirtualCenter or ESX server is not configured for the default port 443 for communications, click the **Advanced** tab and specify the correct port in the **endpoint attribute of NSRhypervisor** field.

      For example, if the VirtualCenter or ESX web service uses port 2000, define the endpoint attribute of NSRhypervisor attribute as:

      `https://server_hostname:2000/sdk`

Where *server_hostname* is the FQDN name of the VirtualCenter or ESX web service host.

d. If a VirtualCenter server is configured, right-click **Virtualization** and select **Run Autodiscovery** to generate the topology map.

Note: If auto discovery fails with the error "Falling back to rsh, but RUSER not provided," ensure that the NetWorker server and the VirtualCenter server can resolve each other's IP / FQDN name.

e. Click **OK**.

Task 2: Configure the VM proxy host

It is not required that the VM Proxy host be backed up, however it is required that a NetWorker client is created for the Proxy host before configuring the virtual clients. The Proxy NetWorker client will be referred to by VM clients during VCB backup and recovery operations.

To create a NetWorker client for the Proxy host:

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 Proxy client.
5. The browse and retention policy fields can remain empty, as they are set for the virtual clients.
6. If the Proxy client must be backed up, ensure that **Scheduled Backups** is selected.

   Note: It is not mandatory to backup the Proxy client.

7. In the **Save Set** attribute, type the name of the files or directories to be backed up:
   a. To specify a file or directory for backup such as C drive, type **c:**.
   b. To back up a specific directory such as Documents and Settings, type **c:\Documents and Settings**.

c. To backup all file systems and VSS/System save sets, type **ALL**.

Note: If the Proxy client will not be backed up use the default selection.

8. In the **Directive** field, select a directive from the list, if desired.

   In NetWorker 7.6, a new directive, **vcb directive** is supported in the following two scenarios:

   - When file level incremental backups are performed instead of FULL image level backups.
   - When FULL file level or incremental file level backups are performed when the save set is ALLVMFS.

   The vcb directive skips the following files and folders:

   - pagefile.sys
   - hiberfil.sys (Hibernation file)
   - WINDOWS\system folder
   - WINDOWS\System32 folder

9. Click **Apps and Modules**.

10. In the Application Information field type:

    VCB_HOST=`any.vc-or-esx.com`

    where *any.vc-or-esx.com* is the hostname of the VirtualCenter or ESX server configured as the NSR Hypervisor resource.

11. In previous releases of NetWorker, the VCB Framework config.js was read to define certain VCB variables. In 7.6 and later, these variables are set and read from the Application Information section. For more information about these variables, refer to the following table.

Table 105 Application information values (1 of 3)

Attribute name	Description	Default value
VCB_BACKUPROOT	• Directory where all the VM backup jobs are supposed to reside in. Note: Ensure that the directory already exists or VCB backup jobs will fail with "directory does not exist" error. • For each backup job, a directory with a unique name derived from the * backup type and the VM name will be created here. • "If omitted, BACKUPROOT defaults to c:\\mnt. Example: VCB_BACKUPROOT=C:\\mnt"	C:\mnt
VCB_HOST **Note:** This attribute is mandatory.	Specify the hostname of the VC or ESX server configured as part of the NSR Hypervisor resource.	

Table 105 Application information values (2 of 3)

Attribute name	Description	Default value
VCB_TRANSPORT_MODE	Specify the transport mode to transfer data from VMFS data store to VCB proxy server. The following options are supported: • SAN – Virtual disk data is read directly off a shared storage device that the virtual disk resides on. This requires VMFS storage on SAN or iSCSI and the storage device has to be accessible from both ESX and the VCB proxy • NBD – VCB will use an over-the-network protocol to access the virtual disk. Data is read from the storage device by the ESX host and then sent across an unencrypted network channel to the VCB proxy. Please note that this mode does not provide the offload capabilities of the "san" mode (since data is still transferred from the ESX host across the network). However, "nbd" does not require shared storage and also enables VCB to be run inside a virtual machine. • NBDSSl – This mode is the same as "nbd" except that the data transferred over the network is encrypted. The data transfer in "nbdssl" mode can be slower and use more CPU than in the "nbd" transport mode. The "nbdssl" mode will fail when connecting directly to ESX 3.0.x or VC 2.0.x host. • hotadd – This mode can be used when Consolidated Backup is used in a virtual proxy. Because it uses the ESX I/O stack to move data, hotadd is more efficient than the transport mode LAN. Before using hotadd, refer to the Virtual Machine Backup Guide to configure the following roles: VMware Consolidated Backup User and VMware Consolidated Backup Proxy. Note when hotadd is used **VCB_PREEXISTING_MOUNTPOINT** and **VCB_PREEXISTING_SNAPSHOT** must be the same, because it is not possible to reliably remove a pre-existing snapshot in this case without cleaning up the mount point. Example: VCB_TRANSPORT_MODE=san	SAN
VCB_VM_LOOKUP_METHOD	Specify the lookup method to be used to identify virtual machines: • ipaddr – VMs are searched based on their DNS name or IP address • name – VMs are searched based on their name displayed in the VirtualCenter. Note that backup will fail for VMs with non-unique names. Only change the default value if you have multiple sets of VMs with the same IP address, but with different names. Example: VCB_VM_LOOKUP_METHOD=ipaddr	ipaddr
VCB_PREEXISTING_MOUNTPOINT	If the mount point already exists, the backup is considered on the following options: • Fail – Backup is made to fail if there is a pre-existing mount point • Delete – Attempts to delete the pre-existing mount point (and backup snapshots associated with it). If this operation fails, then backup fails Example: VCB_PREEXISTING_MOUNTPOINT=fail	Fail

Table 105 Application information values (3 of 3)

Attribute name	Description	Default value
VCB_PREEXISTING_SNAPSHOT	Behavior for backup job when a previous backup snapshot of a virtual machine already exists: • Fail – Backup is made to fail if there is an pre-existing backup snapshot • Delete – Attempts to delete the pre-existing backup snapshot. If this operation fails with errors, then the backup fails Example: VCB_PREEXISTING_VCB_SNAPSHOT=delete	Fail
VCB_MAX_RETRIES	Number of times an operation is re-tried after it fails. Use this option if you see a large number of backup jobs fail with "resource busy" errors. Usually, backup software will retry failed jobs, but it might be hours until the backup software retries. Example VCB_MAX_RETRIES=1	0
VCB_MAX_BACKOFF_TIME	Number of seconds to wait before retrying a failed operation. If you change this default, also change the default for MAX_RETRIES (because this setting only applys if MAX_RETRIES is larger than 0). VCB_BACKOFF_TIME=20	10

Note: Only the automatic snapshot policy is supported as part of the configuration.

Example 48 VCB Configuration

The following example displays all the possible attribute values used for VCB configuration:

```
VCB_HOST=any.vc-or-esx.com
VCB_BACKUPROOT=G:\mnt
VCB_TRANSPORT_MODE=hotadd
VCB_VM_LOOKUP_METHOD=ipaddr
VCB_PREEXISTING_MOUNTPOINT=delete
VCB_PREEXISTING_SNAPSHOT=delete
VCB_MAX_RETRIES=2
VCB_MAX_BACKOFF_TIME=15
```

Task 3: Configure a virtual client for backups

Complete the steps in one of the following topics depending on your environment:

- "Configuring a virtual client if VirtualCenter is configured and auto-discovery has been run:" on page 718

- "Configuring a virtual client if VirtualCenter is not in the environment" on page 720

VMware clients can also be configured as deduplication clients. After creating a VMware client, following instructions for "Deduplication client" on page 656 to configure the client for deduplication.

After the virtual client has been backed up, its client index can be browsed, and data can be recovered directly to the virtual client or can be recovered onto a different virtual client using directed recovery. For recovery of VCB backups, *FULL* image backups can be restored directly on the VCB proxy or CIFS mount point. VMware provided utilities VMconverter or VCBrestore can be used for disaster recovery of the full virtual machine using the full image. "Recovering VMware Consolidated Backups" on page 728 provides recovery procedures.

Configuring a virtual client if VirtualCenter is configured and auto-discovery has been run:

To configure a virtual client if VirtualCenter is configured:

1. In the Virtualization map, right click on the Virtual Machine and select **Client Backup Configuration -> New**.
2. In the Specify the Client Name section, confirm the **client name** field is populated and **Virtual Client** is enabled. Click **Next**.
3. In the Specify the VMWare Backup Type section, the physical host field will be populated with the Physical Host for the Virtual Machine.
4. Select **VCB backup** and from the **Proxy backup host** list, select the name of the Proxy Host server.
5. In the Select the VCB Options section, specify the desired backup:
   - Image level backup (this is equivalent to saveset *FULL*).

     Note: NetWorker directives, including encryption and compression directives, are not supported when backing up a VCB image using the Image Level backups.

   - Backup all files (this is equivalent to saveset ALLVMFS).
   - Backup Specific files and folders.

     Note: Due to limits with the VCB, only one entry is allowed for the Save Set attribute.

- To specify a file or directory for backup such as C: drive, enter **c:** or **c**.
- To back up a specified directory, such as Documents and Settings, enter **c:\Documents and Settings**.

6. Click **Next**.
7. In the Select Client Properties section, select the Browse and Retention policies from the drop down menus.
8. If desired, select the Backup schedule for this client.

   Note: If a backup schedule is also defined for the group that this client will be added to, the group schedule will override the client schedule.

9. If the NetWorker server and VCB Proxy client are two different machines, in the **Remote access** field specify:

   `user=system, host=VCB proxy host`

   Where *system* is the system account of the Windows VCB Proxy and *VCB proxy host* is the name of the Proxy host.
10. Click **Next**.
11. In Choose the Backup Group, choose the desired group or select **Create a new group** and provide a group name and desired number of client retries.
12. If a new group is created, in the Schedule Options section, specify the desired time for the group to start in the **Schedule backup start time** field and enable **Automatically start the backup at the scheduled time**.
13. Click **Next**.
14. In the Specify the Storage Node Options section, ensure that the correct host has been selected for the devices that the backups will be directed to and that recoveries will be performed from.
15. Click **Next**.
16. Review the summary of the client and click **Create**.

    In NetWorker 7.6, a new directive, **vcb directive** is supported in the following two scenarios:

    - When file level incremental backups are performed instead of FULL image level backups.

- When FULL file level or incremental file level backups are performed when the save set is ALLVMFS.

The vcb directive skips the following files and folders:

- pagefile.sys
- hiberfil.sys (Hibernation file)
- WINDOWS\system folder
- WINDOWS\System32 folder

To enable the directive on the VM:

a. Click **Clients**, right-click the newly created VM client, and select **Properties**.

b. From the **Directive** list, select **vcb directive**.

c. Click the Apps and Modules tab and ensure that **nsrvcb_save** is in the **Backup command** field.

d. Click **OK**.

Configuring a virtual client if VirtualCenter is not in the environment

To configure a virtual client when VirtualCenter is not in the environment:

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 client.
5. In the **Browse Policy** field, select a browse policy from the list.

   Note: If the browse policy is set at the client level, it will override the browse policy specified for any groups to which this client is a member.

6. In the **Retention Policy** field, select a retention policy from the list.

   Note: If the retention policy is set at the client level, it will override the retention policy specified for any groups to which this client is a member.

7. Ensure **Scheduled Backups** is selected.
8. In the **Save Set** attribute, type the name of the files or directories to be backed up.

Note: Due to limitations with VCB, only one entry is allowed for the Save Set attribute.

a. To specify a file or directory for backup such as C drive, type **c:**.

b. To back up a specific directory such as Documents and Settings, type **c:\Documents and Settings**.

c. To backup all virtual machine file systems, type **ALLVMFS**.

d. To backup up the entire VM image, type ***FULL***.

9. From the **Directive** attribute, select a directive from the list, if desired.

Note: NetWorker directives, including encryption and compression directives, are not supported when backing up a VCB image using the ***FULL*** keyword in the Save Set attribute of the client resource.

In NetWorker 7.6, a new directive, **vcb directive** is supported in the following two scenarios:

- When file level incremental backups are performed instead of FULL image level backups.
- When FULL file level or incremental file level backups are performed when the save set is ALLVMFS.

The vcb directive skips the following files and folders:

- pagefile.sys
- hiberfil.sys (Hibernation file)
- WINDOWS\system folder
- WINDOWS\System32 folder

10. Click the **Apps and Modules** tab.
11. Select **VCB** for the **Proxy backup type** field.
12. In the **Backup Command** field, type **nsrvcb_save**.
13. If the NetWorker Server and VCB Proxy client are on two different machines:

a. Click on the **Globals (2 of 2)** tab.

b. In the **Remote access** field specify:

user=*system*, **host=***VCB proxy host*

Where *system* is the system account of the Windows VCB Proxy and *VCB proxy host* is the name of the Proxy host.

14. Click **OK**.

Automatic discovery of VMware environments

The Sun StorageTek EBS software can automatically discover VMware environments, and, on a scheduled or on-demand basis, can detect changes to those environments. Automatic discovery is performed by contacting one or more VMware VirtualCenters which host a Web Services server. VMware VirtualCenter is an infrastructure management tool that provides a central point for configuring, provisioning, and managing virtualized IT environments, and is part of VMware Virtual Infrastructure package. For autodiscovery of VMware environments within Sun StorageTek EBS, VMware VirtualCenters must also have Sun StorageTek EBS client software installed. The VMware Infrastructure documentation provides information about configuring VMware VirtualCenter.

Note: The Sun StorageTek EBS software supports auto-discovery via VMware VirtualCenter only. It does not support auto-discovery via an ESX server. "Task 1: Configure the Hypervisor resource" on page 712 describes how to enable auto-discovery.

A binary, nsrvim, is used to facilitate communication between the Sun StorageTek EBS software and the VMware VirtualCenter. The nsrvim binary can communicate with the Web Services server on the VMware VirtualCenter using the secure HTTPS protocol. The nsrvim binary is available only on Windows x86 platforms and is included in the Sun StorageTek EBS for Windows 32-bit installation. All Sun StorageTek EBS servers regardless of platform must contact a Sun StorageTek EBS client running in a Windows environment for auto-discovery. By default, the Sun StorageTek EBS server contacts the Sun StorageTek EBS client running on the VMware VirtualCenter.

The Sun StorageTek EBS software uses auto-discovery for two purposes:

- Notification of changes to the VMware environment.
- Creating and updating the visual view of the VMware environment.

"Notifications of changes to VMware environments" on page 723 provides more details on notifications of changes to the VMware environment. "Visual Representation of VMware environments" on page 724 provides more details on the visual view of the VMware environment.

Performing on-demand auto-discovery of VMware environments

To perform an auto-discovery of VMware environments at any time, right-click the **Virtualization** node in the left pane of the **Configuration** screen and select **Run Auto-Discovery**. Individual elements in the Virtualization node can be selected to limit the auto-discovery task to the selected element.

After selecting Run Auto-Discovery, either from the right-click menu or from the Auto-Discovery dialog, the Running Autodiscovery Now window allows you to monitor the auto-discovery process. Clicking the Stop Monitoring button will close the Running Autodiscovery Now window, but the auto-discovery process will continue.

Note: If auto discovery fails with the error, "Falling back to rsh, but RUSER not provided", ensure that the NetWorker server can resolve the IP/FQDN of the Virtual Centre server, and that the Virtual Centre server can resolve the NetWorker server.

Notifications of changes to VMware environments

After auto-discovery has been performed, the Sun StorageTek EBS software compares the new virtual machine list with the previous one that was created for each VirtualCenter. If the Sun StorageTek EBS software determines that there are any new unprotected virtual machines identified by the Virtual Center that do not have Sun StorageTek EBS Client resources associated with them, a notification will be triggered. A notification is also sent if auto-discover fails.

A default Notification resource, named New Virtual Machine, is included with the Sun StorageTek EBS installation. You can also create additional Notification resources by selecting Hypervisor for the Event attribute of the Notification resource.

"Notifications" on page 619 provides more information.

Visual Representation of VMware environments

After performing auto-discovery of VMware environments, the Sun StorageTek EBS console provides a graphical or tabular view of your VMware environments. This view is available in the right pane when the Virtualization node is selected in the left pane of the Configuration screen.

If auto-discovery has not been configured and the Virtualization node is selected, the right pane displays the Enable Auto-Discovery button. If auto-discovery has been configured and an auto-discovery has been performed, the right pane will display a graphical map of the VMWare environment that was in place during the last auto-discovery. Automatic discovery is performed by contacting one or more VMware VirtualCenters which host a Web Services server. VMware VirtualCenter is an infrastructure management tool that provides a central point for configuring, provisioning, and managing virtualized IT environments, and is part of VMware Virtual Infrastructure package. For autodiscovery of VMware environments within NetWorker, VMware VirtualCenters must also have NetWorker client software installed. See the VMware Infrastructure documentation for information about configuring VMware VirtualCenter.

Virtualization node hierarchical display of the VMware environment

After an auto-discovery has been performed, the Virtualization node in the left pane of the Configuration screen can be expanded to provide a hierarchical display of the VMware environment. Four elements are displayed, in hierarchical view:

1. VirtualCenters
2. DataCenters within the VirtualCenter
3. Clusters within the DataCenter
4. ESX servers

Virtual machines, Sun StorageTek EBS clients associated with those virtual machines, and Sun StorageTek EBS groups performing backups of those clients are not displayed in the Virtualization node hierarchical display. They are displayed in the right pane only.

Clicking on any element in the hierarchical tree will provide a detailed map view of that element and all of its children in the right

pane. For example, selecting the top level virtualization node will display a complete view of your VMware environment across all VirtualCenters that are configured for auto-discovery, while selecting an individual ESX server in the hierarchy will display all child elements associated with that ESX server including virtual machines, Sun StorageTek EBS clients associated with those virtual machines, Sun StorageTek EBS groups performing backups of those clients, and the proxy node for VMware clients.

Two right-click menu operations are available from the Virtualization node:

- Enable Auto-Discovery will open the Auto-Discovery dialog to configure auto-discovery, as described in "Performing on-demand auto-discovery of VMware environments" on page 723.
- Run Auto-Discovery will perform an on-demand auto-discovery of your VMware environment. Individual elements in the Virtualization node can be selected to limit the auto-discovery task to the selected element.

Graphical display of the VMware environment

After an auto-discovery has been performed, elements of the VMware environment are displayed in the right "details" pane of the Sun StorageTek EBS Console. Objects displayed in the details pane vary depending on what is selected in the Virtualization node in the left pane. Several operations are available from the details pane, such as configuring new Sun StorageTek EBS clients to protect virtual machines.

Note: In order for a NetWorker Client resource to appear in the details pane, the name of the virtual machine and the name of the NetWorker Client resource must be identical.

Two views are available:

- Map view
- Tabular view

Map view of the VMware environment

Items displayed in the map view of the VMware environment will vary depending on what is selected in the Virtualization node in the

left pane. If the Virtualization node is selected, the map view displays all VirtualCenters that have had an auto-discovery performed and all child elements of those Virtual Centers, beginning with VirtualCenters on the left side of the pane. Lines connect each child element to its parent, with child elements proceeding hierarchically from left to right in the display.

Items displayed in the right details pane can be refined by selecting child elements in the Virtualization node hierarchy in the left pane. For example, if an individual ESX server is selected in the Virtualization node, only child elements associated with that ESX server are displayed.

Displaying Sun StorageTek EBS clients associated with virtual machines

By default, Sun StorageTek EBS clients associated with individual virtual machines are displayed. Rather, a Client icon will indicate whether the virtual machine has one or more Sun StorageTek EBS clients configured to protect it. Sun StorageTek EBS groups performing backups of those clients will be displayed with lines connecting the groups to the virtual machine.

If the virtual machine is being protected by the Sun StorageTek EBS software, you can double-click on the virtual machine to expand the display to view Sun StorageTek EBS clients configured to protect the virtual machine, with a line connecting the client to the Sun StorageTek EBS group that performs the backup of that client. You can also expand all virtual machines in the display by right-clicking anywhere in the right pane and selecting **Expand**>**All VMs**.

Creating new Sun StorageTek EBS clients for unprotected virtual machines

If a virtual machine displayed in the right details pane is unprotected, this is indicated by the lack of a Client icon for that virtual machine. You can create a new Sun StorageTek EBS client for that virtual machine by right-clicking on the virtual machine and selecting **Client Backup Configuration**>**New** to open the Client Backup Configuration Wizard, or by selecting **New** to manually create a new client. "Configuring Sun StorageTek EBS clients for virtual machines" on page 710 provides information about creating clients to protect virtual machines.

Other operations available from the map view

You can also perform typical Sun StorageTek EBS operations on clients and groups from the map view. For example, by right-clicking

on an existing Sun StorageTek EBS client, you can edit, delete, and copy clients, as well as initiating a recovery. You can also right-click on a Sun StorageTek EBS group displayed in the map view and perform typical group operations, such as editing or copying the group with all clients.

Unscheduled clients for backup are displayed in dotted-line within the configuration. Multiple instances of the same client within the savegroup are represented by their backup type, client name, and saveset name.

Navigating within the Map view

Several operations are available to facilitate navigation within the map view:

- Zoom: You can zoom in and out of the map view by selecting the zoom icons on the map view icon bar or by clicking on the right details pane and scrolling with the middle mouse wheel. You can also select an area to zoom into by selecting the **Zoom Area** button, or fit the entire display into the right details pane by selecting the **Fit Content** button. These operations are also available from the right-click menu in the details pane.
- Moving the display: You can move the graphical display by left-clicking in the details pane and dragging the mouse cursor.
- Expanding and collapsing elements: You can expand or collapse any element in the map view to display or hide the child elements associated with the element by double-clicking the element. Additionally, you can expand or collapse all elements of a certain type by right-clicking anywhere in the details pane and selecting **Expand** or **Collapse**, and then selecting the element type.
- Overview: You can open the Overview dialog by selecting the Overview icon on the map view icon bar or by right-clicking anywhere in the details pane and selecting Overview. The Overview dialog is particularly useful for large maps and allows you to quickly drill down to specific areas in the map.
- Show and Find: The Show and Find functions allow you to limit items displayed in the map, and to search for specific items.
- Tabular view: You can also switch to viewing the VMware environment in tabular view, rather than map view, by selecting the Table icon on the map view icon bar or by right-clicking anywhere in the details pane and selecting Table.

Tabular view of the VMware environment

The right details pane can display the VMware environment in tabular form, rather than map form, by selecting the Table icon on the map view icon bar or by right-clicking anywhere in the details pane and selecting Table. The tabular view functions like other tabular views in the Sun StorageTek EBS Console.

"Sun StorageTek EBS Management Console interface" on page 49 provides general information on using tables in the Sun StorageTek EBS Console.

Recovering VMware Consolidated Backups

This section outlines the following procedures for recovering data backed up using VMware Consolidated Backup (VCB):

- "Performing a file-based recovery on the local host" on page 728
- "Performing a file-based recovery using CIFS share" on page 728
- "Performing a file-based recovery using directed recovery" on page 729
- "Performing a single step recover of the full virtual machine" on page 730

"Configuring Sun StorageTek EBS clients for virtual machines" on page 710 provides information about VCB support and configuration.

Performing a file-based recovery on the local host

File-based recovery on the local host running a VM client requires that the Sun StorageTek EBS client is installed on the VCB proxy.

To perform a file-based recovery on the local host:

1. Launch the Sun StorageTek EBS User program on the VM client.
2. Follow the procedure outlined in "Recovering data on Windows systems" on page 426.

Performing a file-based recovery using CIFS share

To perform a file-based recovery using the CIFS share:

1. Launch the Sun StorageTek EBS User program on the Sun StorageTek EBS server or VCB proxy.

   Note: The remote access list of the VM client must include either user@server or user@proxy.

2. Browse the file system for the VM client and select file to recover, as outlined in "Recovering data on Windows systems" on page 426.
3. Set the destination directory to the CIFS share of the VM client.
4. Recovery the files onto the CIFS share.
5. At the VM client, move the files from the CIFS share to the appropriate directory.

Performing a file-based recovery using directed recovery

File-based recovery using directed recovery requires that the Sun StorageTek EBS client is installed on the VCB proxy.

1. Launch the Sun StorageTek EBS User program on the Sun StorageTek EBS server or VCB proxy.

   Note: The user must have the Remote Access All Clients privilege.

2. Select the VM client as source client.
3. Select the target client as VM-client.
4. Select a destination folder.
5. Follow the procedure in "Recovering data on Windows systems" on page 426 to select files for recovery and perform the recovery.

The following considerations apply when performing a file-level recovery of a image level backup:

Note: Support for file-level recovery from image based backup is available only on Windows.

- File-level recovery is supported only on VMs having Windows OS with NTFS5 file system.

- File-level incremental backup of a VM is not supported after a hardware change in the VM. Perform a full image level backup after every change in the VM hardware.

Performing a single step recover of the full virtual machine

There are two methods of recovering a full virtual machine:

- "Perform a single step recover from the NetWorker User program" on page 731
- "Perform a single step recover from the command line" on page 732

The following considerations apply when performing a single step recover of a full VMware virtual machine:

Note: In case of a remote VCB proxy client, single step recover requires the members of the VCB proxy clients administrator group to be part of the remote access list of the VM clients or should have the "Remote access all clients" privilege.

- The VMware Consolidated Backup (VCB) proxy system must be running Microsoft Windows 2003 (with at least SP1 installed) or Microsoft Windows 2008.
- Recover of the full virtual machine is only supported using save set recovery.
- The user must have the required VMware privileges to register or create virtual machines.
- The VMware converter must be installed on the VCB proxy host machine. If the VMware converter is not installed, the save set of the full virtual machine (FullVM) can be recovered using a traditional NetWorker recovery.

  Note: Single step recovery is only supported with VMware stand-alone converter version 3.0.3.

- The ESX server must be at version 3.x or later.
- The VMware virtual machine will recover to the same VMware ESX server or VMware Virtual Center (VC) taken at the time of backup. Specifying another VMware ESX server or VMware VC server will cause the recover operation to fail.

- A recover of the VMware virtual machine will fail if the VMware virtual machine already exists in the specified VMware ESX or VMware VC server.

Note: During a FULLVM restore using the GUI or Command Line Interface, restored virtual machine will start as in forceful powered off state because of VCB snapshot limitation.

Perform a single step recover from the NetWorker User program

This procedure is supported on Windows XP and later Windows platforms only.

To perform a single step recover of a full VMware virtual machine to the VMware ESX server or VMware Virtual Center server:

1. Launch the Sun StorageTek EBS User program on the Sun StorageTek EBS client or VCB proxy.
2. From the **Operation** menu, select **Save Set Recover**.
3. In the **Source Client** dialog box, select the virtual machine client from where the save set originated and click **OK**.
4. In the **Save Sets** dialog box, select the **Save Set** name for the full virtual machine backup client (FULLVM) and click **OK**.

   If the VMware converter is installed on the VCB proxy host machine:

   a. In the **VCB Restore** dialog box, type the required information.

      – **ESX or VC server name** — Specify the VMware ESX server or VC server name to which the VMware virtual machine will be recovered. You must specify the server's IP address or its fully qualified domain name (FQDN).
      – **User name and Password** — Specify the user name and password of the VMware ESX server or VMware VC server.
      – **Restoring staging location** — Specify a staging location to recover the backed up FULL VM image before recovering it to the VMware ESX server or VMware VC server.

      The Restore staging location adds the "FULLVM" directory to the specified staging location.

   b. Select **Delete staging location on restore** to automatically delete the temporary staging location.

c. (Optional) **Data store name** - Use this field only if you want to recover the VMware virtual machine to a data store other than the one that was used at the time of backup. To do so, specify the name of the shared data store that is accessible by the VMware ESX server or VMware VC server.

d. Click **Restore**.

If NetWorker cannot access the VMware converter, a traditional save set recovery is attempted:

a. In the **Recover Options** dialog box, enter a valid local path in the **Relocate recovered data to** attribute.

b. Click **OK**.

Perform a single step recover from the command line

To perform a command line recover of a full VMware virtual machine to the VMware ESX server or VMware Virtual Center (VC) server:

1. Use the **mminfo** command to determine the save set ID of the full virtual machine (FULLVM) client, for example:

   mminfo -avot -q "name=fullvm"

2. Recover the full VMware virtual machine using the **recover** command, for example:

   recover -S *ssid* **-d** *staging-location* **-o** *VCB:host*=*ESX/VC host-name*;*VCB:user*=*ESX/VC username*;*VCB:password*=*ESX/VC password*[;*VCB:datastore*=*data-store-name*;*VCB:delstaging*=**<0|1>**]

   where:

   - *ssid* is the save set identifier of the FULLVM.
   - *staging-location* is the staging location path to recover the backed up FULLVM image. This is a temporary recover location and will be deleted if the VCB:delstaging flag is set to 1.
   - *VCB:host* is the VMware ESX server machine name where the VMware virtual machine needs to be recovered.
   - *VCB:user* is the Username of the VMware ESX/VC server.
   - *VCB:password* is the Password of the VMware ESX/VC server.
   - *VCB:datastore* is the disk name that is accessible by the ESX/VC server. This is an optional value. If specified, the virtual machine will be imported to the specified data store disk name.

- *VCB:delstaging* is an optional flag that specifies whether to delete the staging location after a recover operation. If this flag is set to '1' (True) the staging location will be deleted after recovering the virtual machine. The default value of this flag is 0 (False).

The following example depicts a command to recover the FULLVM with ssid 413546679 to the ESX server named esxDemo1.emc.com. The staging location is to be deleted after recovering it to the ESX server:

recover.exe -S 413546679 **-d** e:\mnt\ **-o VCB:host=**esxDemo1.emc.com**;VCB:user=**root**; VCB:password=**w5dft9s1**;VCB:delstaging=**1

If the recover operation fails with an error indicating that there was an issue with the VMware converter, attempt to recover the FULLVM save set by following the procedure described in the next procedure, "Traditional command line recovery of the full VMware (FULLVM) save set".

Traditional command line recovery of the full VMware (FULLVM) save set

If the steps listed in "Perform a single step recover from the command line" on page 732 do not work, attempt to recover the FULLVM save set using a traditional command line recovery, for example:

recover -s *server* **-S** *ssid* **-d** *local-path*

where:

- *server* is the NetWorker server.
- *ssid* is the save set identifier of the FULLVM.
- *local-path* is a valid local path.

Note: If the traditional recovery option fails with an error message indicating that the VMware converter failed to start, the most likely cause is that the VMware converter was uninstalled incorrectly. Either re-install the VMware converter on the VCB proxy host machine or remove any remnants of the VMware converter program from the host machine. Then, re-attempt the single step recover operation.

21

Sun StorageTek EBS Server Support for NDMP

This chapter covers these topics:

- Overview of NDMP 736
- Supported NDMP configurations 737
- Requirements for and limitations of NDMP support 743
- Configuring the Sun StorageTek EBS server for NDMP operations 750
- Performing NDMP backups 763
- Performing NDMP recoveries 770
- NDMP cloning 780
- Ancillary product information 785

Overview of NDMP

"NDMP (Network Data Management Protocol)" 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.

NDMP enables disparate vendors to use a common NDMP protocol for the backup architecture. Data is copied from disk to tape or disk media, regardless of the operating system or platform. Tracking information is also passed to and from the backup software by using the same protocol, regardless of the software type.

NDMP works with the Sun StorageTek EBS software for backups, cloning, recoveries, and updates or queries to resource files. All communications are through NDMP.

The objective of NDMP is to address problems associated with backing up heterogeneous networks for operating system vendors, backup developers, and Network Attached Storage (NAS) devices.

Sun StorageTek EBS NDMP support includes library sharing and library management in a NAS environment. The term library is synonymous with autochanger, autoloader, carousel, datawheel, jukebox, and near-line storage.

Licensing

The Sun StorageTek EBS NDMP interface is an optional feature that must be licensed separately from the Sun StorageTek EBS base product according to a tiered licensing structure.

Chapter 28, "Licensing" provides information on Sun StorageTek EBS licensing support for NDMP.

NDMP connection support

With the Sun StorageTek EBS NDMP interface, you can connect to hosts that have an active "NDMP service" or an NDMP data module installed. The hosts, clients, or devices (except for the Sun StorageTek EBS server) do not need to have the Sun StorageTek EBS software installed. NDMP allows a NAS device to back up data to other NDMP-controlled tape or disk devices on the network. The tape or disk device does not have to be directly connected to the NAS device.

NDMP passes control of data and file metadata to and from the Sun StorageTek EBS program.

The respective vendor's documentation provides information on how to install an NDMP data module or enable NDMP support on an application server.

Supported NDMP configurations

The Sun StorageTek EBS server with NDMP Connection supports these configurations:

- Local backup when using an NDMP Tape Server — The backup is performed to a tape device attached to the server where the data resides.
- Sun StorageTek EBS Storage Node — The backup is performed from an NDMP host to a Sun StorageTek EBS tape device, disk, optical, and other media type attached to a Sun StorageTek EBS storage node.

  Sun StorageTek EBS storage nodes use a data server agent (DSA). DSA acts as an agent in the performance of backups between a Sun StorageTek EBS server and any client without Sun StorageTek EBS software. An example of a DSA is an NDMP host that generates proprietary save data and sends that data to a Sun StorageTek EBS storage device to have a save set associated to it. Similarly, a DSA will act as an agent in the recovery between a Sun StorageTek EBS server and any client without Sun StorageTek EBS software.

  Note: Any backup performed with the DSA feature has a **D** flag.

- Three-party that uses NDMP Tape Server — The backup is performed from one NDMP server to another NDMP server.

Local backup by using an NDMP tape server

A local backup with an NDMP tape server means the data being backed up is sent from the NDMP tape/data server (both on the same host) to its directly attached tape device or library. The data does not traverse the network, thus preventing network congestion.

With the local backup type, only metadata and NDMP control information (shown in Figure 40 on page 739 as file history) traverse the network between the backup server and the NDMP tape and data server. The Sun StorageTek EBS program uses the file history information to maintain appropriate indexes and media databases.

Figure 40 illustrates a local backup type configuration. The NDMP host is both an NDMP data service provider (DSP) and tape service.

The Sun StorageTek EBS server, or data management application (DMA) performs these tasks:

- Initiates the backup or recovery request through the NDMP connection.
- Receives the file history information from the data service and tape service.

The DSP performs these tasks:

- Receives the backup requests.
- Reads the backup data from disk.
- Produces a datastream for backup.

As a result, these actions occur:

- During a backup, the tape service reads a datastream from the DSP and writes it to tape.
- During a recovery, the tape service reads data from tape and generates an NDMP datastream.

  The the following occurs:

  1. The datastream is sent to the DSP.
  2. The DSP reads from the datastream, and it to disk.

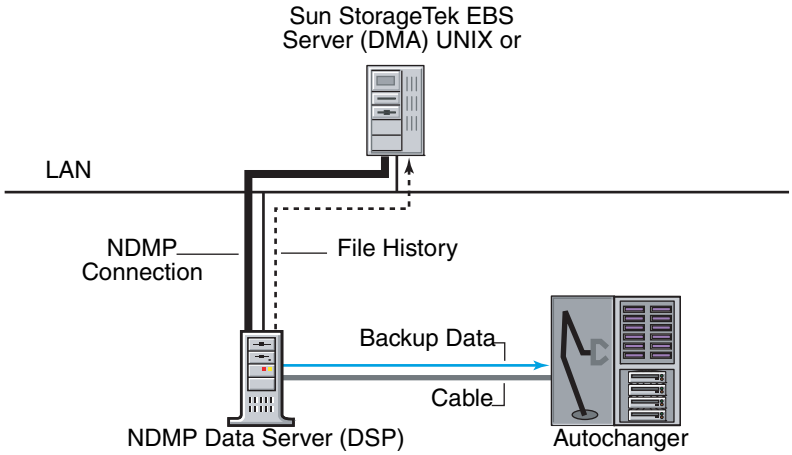


Figure 40 NDMP local configuration

Backup by using Sun StorageTek EBS storage node support

This section features the supported NDMP client backup configurations which use non-NDMP based Sun StorageTek EBS devices, such as a tape, disk, or optical.

Figure 41 on page 740 illustrates a Sun StorageTek EBS storage device attached directly to the Sun StorageTek EBS server. The NDMP backup that uses DSA, is initiated from the Sun StorageTek EBS server. NDMP save set data may be written to either a file type device or tape. Metadata, including file history, is processed in the Sun StorageTek EBS server (**nsrmmd**).

Table 106 shows how DSA differs in use from NDMP tape server.

Table 106 Distinctions between NDMP tape server and DSA

NDMP Tape Server	DSA
Supports only NDMP type of tape device	Supports any type of Sun StorageTek EBS supported device.
Does not support backup to disk.	Supports backup to disk.
Does not support multiplexing.	Multiplexing is supported.
Record size is fixed for all device types, such as DLT7000, SDLT, Ultrium.	Record size varies with the type of device, which improves the write performances.

"Data write technologies" on page 769 provides information on storage technologies under Sun StorageTek EBS storage node support for NDMP clients.

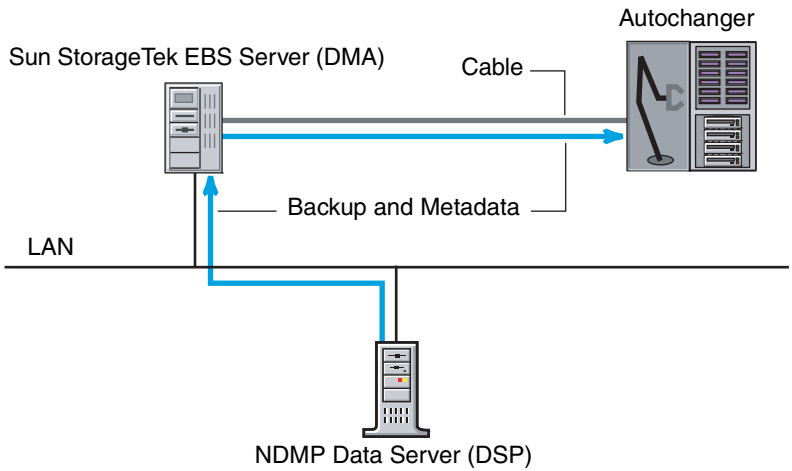


Figure 41 Backup Initiated from a Sun StorageTek EBS server with an attached storage device

Figure 42 on page 741 illustrates a Sun StorageTek EBS configuration in which the file history is backed up to the Sun StorageTek EBS server, and the data is backed up to the Sun StorageTek EBS client

with an attached storage node. The backup command, **nsrndmp_save**, is run directly on the Sun StorageTek EBS client that has a Sun StorageTek EBS storage node locally attached.

Data is processed and saved to the Sun StorageTek EBS storage node, while the file history is processed on the designated Sun StorageTek EBS server (**nsrmmd**).

After the data is backed up and sessions with the NDMP server and the Sun StorageTek EBS server are closed, the file history is committed to the Sun StorageTek EBS server.

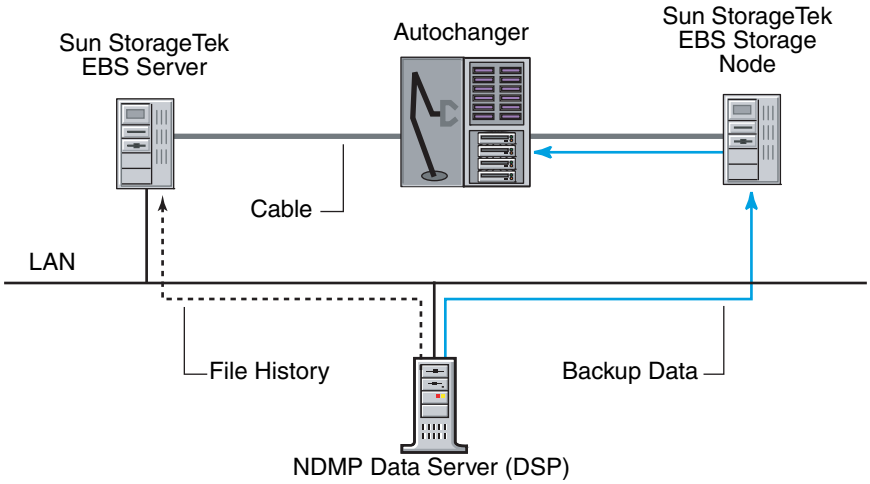


Figure 42 File history and backup data to Sun StorageTek EBS storage node

In Figure 43 on page 742, the file history and data are backed up to the Sun StorageTek EBS server, which is also an EMC NetWorker SnapImage™ Module NDMP data server. In this configuration, the data server does a block-level backup.

The storage technology is the same as in Figure 41 on page 740.

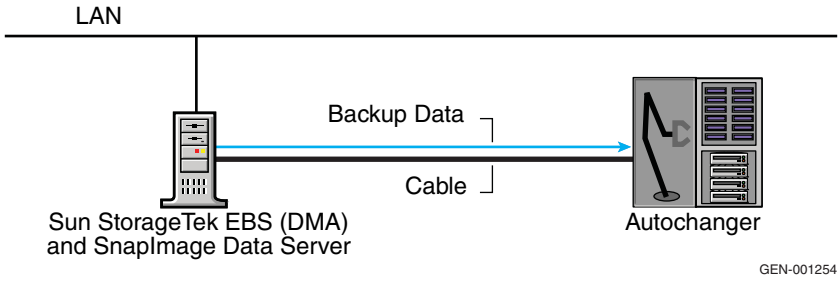


Figure 43 Backup that uses SnapImage module

Three-party backup with NDMP tape server

A three-party backup, also known as a three-way backup, is completed by using two NDMP servers. The data flows from the NDMP host (data service provider DSP) to the NDMP tape server, and then to the library that is locally attached to the tape server. The DSP might be an EMC, NetApp Filer, Auspex, or SnapImage Module.

Note: The two NDMP servers must be of the same type. For example, backups between an EMC NAS device and a NetApp server are not supported. A NetApp server must back up from, and recover to, another NetApp server.

Figure 44 on page 743 demonstrates a three-party configuration, which enables backup and recovery to a Sun StorageTek EBS device that is attached to another NDMP server.

- One server is the data server.
- The second server is the tape server.
- The third party is the Sun StorageTek EBS server (DMA).

This configuration is similar to the flow of data between a Sun StorageTek EBS client and a Sun StorageTek EBS server or storage node, except that it is not necessary to install special software on either end. Data flows from the NDMP data server over the network to the tape server and then to tape. Metadata is sent from the data server to the backup server.

Supported tape servers

The Sun StorageTek EBS software with NDMP supports these tape servers:

- Sun StorageTek EBS NDMP TapeServer 2.0 software installed on the Sun StorageTek EBS server or another server.
- NetWorker SnapImage Module software installed on the Sun StorageTek EBS server or another server.
- DinoStor TapeServer connects one or more libraries to the network. This allows any NDMP host to be backed up to one location instead of requiring a local backup device for each server.

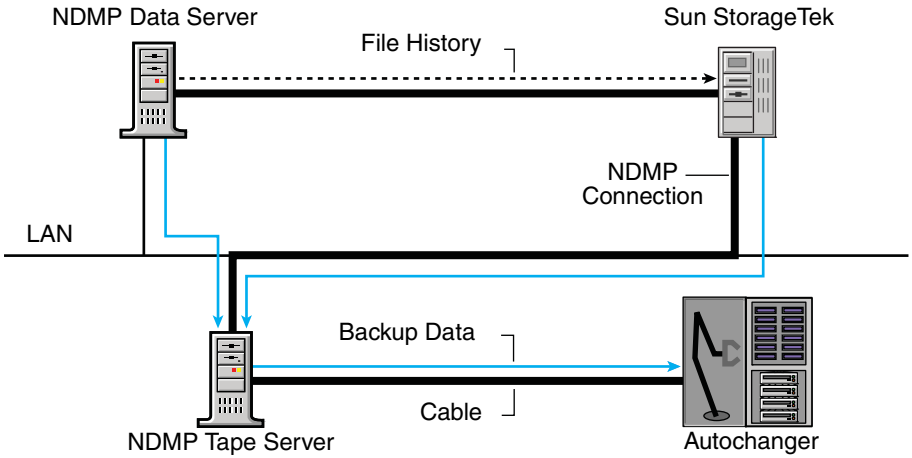


Figure 44 NDMP tape server

Requirements for and limitations of NDMP support

These sections contain information about requirements, support, and limitations of the Sun StorageTek EBS software for NDMP, including "Sun StorageTek EBS storage node support for NDMP clients" on page 747.

Requirements for Sun StorageTek EBS software with NDMP support

These requirements apply when running the Sun StorageTek EBS software with NDMP support:

- To back up the bootstrap file, a non-NDMP device must be attached to the Sun StorageTek EBS server. The bootstrap includes the media database and configuration files needed to recover the Sun StorageTek EBS server after a disk crash.

  This requirement does not apply when the backup is to a Sun StorageTek EBS storage node with a non-NDMP device. In this case, the same non-NDMP device could be used for a bootstrap device.

- To perform either an NDMP backup or recovery operation, you must be one of these:
  - Logged in as root
  - Windows Administrator
- The media device handle must be a nonrewind device handle.
- To perform parallel backups, perform one of these actions.
  - Specify multiple save sets in a single Client resource and verify multiple tape drives are available.
  - Specify multiple Client resources in separate groups (with one client/save stream in each group) are configured with the same start time for each volume to be backed up.
  - Specify Save Set All.
- When recovering data in a localized NDMP environment, the Index Recover status window shows the process in English, rather than the localized language.

Limitations of the Sun StorageTek EBS software with NDMP support

The addition of NDMP support to a Sun StorageTek EBS datazone environment renders these limitations:

- Other NDMP hosts cannot be browsed for backups. However, you can browse for save set recoveries by using the command prompt, the **Sun StorageTek EBS User** (Windows), or **nwrecover** (UNIX) program.

- Running the **scanner** command on a volume that has both NDMP and non-NDMP save sets.
  - Use a non-NDMP device path, an NDMP device path is not supported.
  - Use a local device for improved performance.
- The following features are not supported:
  - Archiving
  - Save set consolidation

    Save set consolidation interprets save records on the media. The DSA does not write save records, hence the save set consolidation is not available. The DSA acts as an agent in the performance of backups between a Sun StorageTek EBS server and any client without Sun StorageTek EBS software that generates data and sends it to a Sun StorageTek EBS storage node or device to have a save set associated to it.
  - Cloning from NDMP tape device to non-NDMP tape device and vice versa
  - The **jbexercise** utility
- Sun StorageTek EBS Server resource attributes (**nsrmmd** polling interval, **nsrmmd** restart interval, and **nsrmmd** control timeout) do not apply to NDMP hosts, but they do apply when using Sun StorageTek EBS storage nodes for backup and recovery of data.
- The Sun StorageTek EBS software does not support the recovery options **Rename Recovered File** and **Discard Recovered File** for NDMP operations. For both save set and file-by-file recoveries, *current files are always overwritten*.
- The Sun StorageTek EBS software does not support incremental backup settings for NDMP backups. If you select the incremental setting, the Sun StorageTek EBS server performs a full backup. However, you can schedule level backups to function like incremental backups. For example, a weekly backup schedule of full on day 1 and incrementals on days 2 through 7 is the same as 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.
- The **scanner -i** command fails for DSA save sets and skips them. The **scanner –i** command is not supported for NDMP backups.

- The **SIZED** flag is not supported with NDMP. This flag is set in the device configuration attribute, **Device tape flags**, which is displayed only in Diagnostic mode.
- The NDMP device block size should not exceed 256k for the NDMP backups to complete successfully.

Protocol limitations

These protocol limitations apply to versions 1 and 2 of the NDMP protocol:

- The Save Set attribute default All is supported for NDMP servers with NDMP version 3 or later.
- The library handle must be determined to run **jbconfig**.

Important notes on index processing and memory requirements

These notes concern the index processing and memory requirements and other information about NDMP backups that use Sun StorageTek EBS software.

Index Database

During an NDMP backup, the NDMP data server sends the file history metadata information to the Sun StorageTek EBS server. The **nsrndmp_2fh** binary starts processing the file history metadata immediately. The generation of the index database occurs at the end of the backup by the **nsrdmpix** binary. The file history metadata processing scales linearly with the increase in the number of entries (files) in the file system (save sets).

Especially during failover of a clustered Sun StorageTek EBS environment, if NDMP indexing is interrupted, the following results occur:

- The **nsrdmpix** binary will not automatically restart.
- The generation of the index database may be incomplete.
  - This means that a file-by-file recovery will not be possible.
  - A save set recovery may still be performed.

Memory requirements

The **nsrndmp_2fh** and **nsrdmpix** binaries interact with the raw database, instead of virtual memory, to process the file history metadata. Memory requirements for this process are minimal.

Temporary space required during Backup

To determine the required physical space for the default \nsr\tmp directory per file entry, use this formula.

2 * (144 + average filename length) * number of entries in the file system

Example 49 **Calculating temporary space requirements**

For one million file entries with an average filename length of 128, the required temporary swap space is computed with this formula.

2 * (144 + 128) * 1,000,000 = 544 MB approximately

Changing the location of the temporary space directory

To change the location of the temporary space directory, provide a valid directory name as a value to the NSR_NDMP_TMP_DIR variable:

1. Change the application information in the Client window.
2. Set an environment variable.

This change moves the creation of the temporary files to that directory.

Potential error messages

If a backup or recovery fails due to an inadequate amount of swap space, error messages similar to the following appear:

Failed to store index entries

When the **nsrndmp_2fh** binary runs out of temporary space, this error message appears.

IO_WritePage write failed - No space left on device (28): No space left on device

The communication error between **nsrndmp_save** and **nsrndmp_2fh** results in this error:

```
error in reading the FH entries from save through stdin
```

Sun StorageTek EBS storage node support for NDMP clients

Sun StorageTek EBS storage node functionality includes NDMP save streams. Sun StorageTek EBS storage node software can be used for

both NDMP and non-NDMP based backups in a unified fashion. For NDMP backups, Sun StorageTek EBS uses the DSA feature.

Clients that generate data have backup history in the form of manageable Sun StorageTek EBS save sets. These clients are referred to as NDMP clients.

Note: The DSA feature listens on a port number in the "service port range" of that system (DSA running host).

NDMP backups to non-NDMP based Sun StorageTek EBS devices

With this extension, Sun StorageTek EBS supports backup of an NDMP client to a non-NDMP based Sun StorageTek EBS device such as:

- Tape
- Optical
- Disk
- Other media types

Features with Sun StorageTek EBS storage node support for NDMP

Sun StorageTek EBS storage node support for NDMP backups include these features:

- Multiplexing
- Staging
- Coexistence of NDMP save sets with non-NDMP based Sun StorageTek EBS save sets
- Firewall support
- Availability on wide variety of platforms
- Auto media verification
- Backup to disk
- True Sun StorageTek EBS client support
- Cloning
- Direct Access Restore (DAR)

Save set flags in media database

The **mminfo** program (with the **-p** option to display a report on the browse and retention times for save sets) is used to view save sets.

Sun StorageTek EBS storage node support for NDMP clients uses different flags with this program.

Table 107 on page 749 lists the save set flags that are displayed by **mminfo** output and their meanings.

Table 107 DSA feature save set flags

Flag	Description
N	The client of the save set is an NDMP client.
s	The backup was performed to a Sun StorageTek EBS storage node by using DSA.
Ns	The backup was performed by an NDMP client to a Sun StorageTek EBS storage node by using DSA.

Performance

An increase in NDMP backup performance can be seen in several instances.

The NDMP tape server is limited to use 60 KB as the record size to write to the tape. For backup to non-NDMP devices, the size of the record depends on the type of the media involved. For instance, for an dlt8000 tape device, the record size is 128 KB. The backup speed increases when compared to the NDMP tape devices.

The **nsrndmp_save** command does not query the DSA every second to obtain the backed-up size and update the Sun StorageTek EBS Console, Sun StorageTek EBS User or **nwrecover** and **nsrmmd**, as it does the NDMP tape server while backing up to an NDMP tape device.

- Backup to disk improves the performance.
- Staging and cloning do not need to involve the NDMP tape server.

During a volume selection process, loading and positioning the volume does not occur. This is due to an exchange between **nsrmmd** and the **nsrndmp_save** or **nsrndmp_recover** commands. This exchange is an overhead situation, which is avoided by backing up to the Sun StorageTek EBS storage node.

Vendor-specific information

For information regarding the limitations, library configuration, and application information variables of each vendor, see Table 108 on page 750.

Table 108 Vendor-specific information

Vendor	Refer to these links
Network Appliance	"Network Appliance" on page 786
Celerra	"EMC Celerra" on page 791
EMC IP4700	"EMC IP4700" on page 795
Auspex	"Auspex" on page 796
Procom NetFORCE	"Procom NetFORCE" on page 799
Mirapoint	"Mirapoint" on page 800
BlueArc	"BlueArc" on page 800
DinoStor	"DinoStor TapeServer" on page 802
SnapImage	"SnapImage or NDMP TapeServer Module" on page 803

Configuring the Sun StorageTek EBS server for NDMP operations

Before performing an NDMP backup, configure the Sun StorageTek EBS server for NDMP operations. Use the Sun StorageTek EBS Console Administration window to configure the resources listed in Table 109 on page 751 on the Sun StorageTek EBS server before backing up data. Note that certain attributes of the resources are configured differently than for non-NDMP clients. "Additional considerations for NDMP operations" on page 762 provides information on other configuration issues.

Note: Resources not listed in Table 109 on page 751 should be configured the same as for non-NDMP operations.

Table 109 Configuring resources for NDMP

Resource	Configuration
Autochanger	Configure a library with the Sun StorageTek EBS Console Administration window. "Configuring library support for NDMP" on page 751 provides instructions. "Ancillary product information" on page 785 provides vendor-specific requirements.
Device	To use a stand-alone device, create a Device resource for NDMP backups. "Configuring an NDMP device resource" on page 757. "Configuring library support for NDMP" on page 751 provides information on how to use a library.
Advanced File Type Device	Configure an advanced file device with the Administration window of the Sun StorageTek EBS Console.
Client	Create a Client resource with specific NDMP attributes. "Configuring an NDMP Client resource" on page 758 . "Ancillary product information" on page 785 provides vendor-specific requirements. To perform NDMP-based backups and recoveries, the Client resource must be set up with specific NDMP attributes.
Group	Add the NDMP client to a backup group. If the backup contains more than five million files, set the Inactivity Timeout attribute to zero (0). "Additional considerations for NDMP operations" on page 762 provides information on additional issues.
Schedule	Select a schedule for the NDMP client. NDMP supports level full and levels 1-9.
Policy	Specify the browse and retention policies for the NDMP client.

Configuring library support for NDMP

Before performing an NDMP backup or recovery, configure the libraries by using the server's Sun StorageTek EBS Administration interface. Depending on the environment, the library handle may not be required to configure the library:

- If the NDMP host is running NDMP version 3 or later and supports NDMP_CONFIG interfaces, the library handle *is not* required to configure the library.
- If the NDMP host is running NDMP version 2 or earlier or does not support NDMP_CONFIG interfaces, the library handle *is* required to configure the library.

To determine if the NDMP host supports NDMP_CONFIG interfaces, refer to the appropriate vendor documentation.

Note: "Configuring an NDMP device resource" on page 757 provides instructions on creating a Device resource for a stand-alone device.

Obtaining information about NDMP autochangers

Before the library can be configured for an NDMP backup and recovery, obtain the pathname of the media device. "Ancillary product information" on page 785 provides instructions on obtaining the device pathname for a particular vendor.

Using the Sun StorageTek EBS Administration window with NDMP devices

To use the Sun StorageTek EBS Administrator window to search for available devices on the NDMP tape server:

1. From the **Console** window, click **Enterprise**.
2. Click the Sun StorageTek EBS server hostname.
3. Double-click **Sun StorageTek EBS**. The **Sun StorageTek EBS Administration** window opens.
4. Click **Devices**.
5. Right-click the **NetWorker Server**.
6. Select **Scan for Devices**.
7. Click the **Monitoring** button.
8. Click the **Devices** tab to view the device lists.

Using the inquire command with NDMP devices

Use the **inquire** command to search for available devices on the NDMP tape server.

The NDMP library information can be obtained two ways:

- From the location on the local server.
- On a specific NDMP host.

To obtain the NDMP library information:

1. Log in as root or as a member of the Windows Administrator's group.
2. Type this command.

inquire -N *ndmp_hostname*

3. Respond to the NDMP prompts as they occur.
   a. Type the username.
   b. Type the password for the NDMP host *hostname* (characters will not be echoed). The list of devices appears.

The *EMC NetWorker Command Reference Guide* or the UNIX man page provide more information on the **inquire** command.

Configuring an NDMP autochanger

The Sun StorageTek EBS server's Administration interface and the Sun StorageTek EBS Autochanger resource configuration tool, **jbconfig**, are used to configure libraries for NDMP operations. Once the library is configured, the same interfaces are used to add or remove drives and devices.

Chapter 3, "Configuring Libraries and Devices" provides information on how to configure a library with the Sun StorageTek EBS server's Administration interface.

To configure a library with the **jbconfig** tool:

1. Determine the device pathname information for the NDMP library. "Ancillary product information" on page 785 provides information on how to obtain this information.
2. At the command prompt, change to the <Sun StorageTek EBS_install_path>\bin directory.
3. As root or as a member of the Windows Administrator's group, type the following command:

   jbconfig

   The use of **jbconfig** requires an NDMP password for the computer with the attached tape device.
4. When prompted, select one of these options and answer the questions that appear:
   - If the library is connected to an NDMP tape server and the robot is controlled by the NDMP tape server, select the Autodetected NDMP SCSI Autochanger option.

- If the Sun StorageTek EBS backup server or storage node controls the robot and shares the library with an NDMP tape server, select the Autodetected SCSI Autochanger option. Configure the tape drives to be used by the NDMP tape server.

Example 50 How to configure an NDMP autochanger for a NetApp computer

This partial example demonstrates how to configure an NDMP library for a NetApp filer by using the **jbconfig** program. The script uses the term *jukebox* instead of *library*. The terms are interchangeable.

Note: This example might not be valid for some NDMP hosts. To identify valid device names, refer to the NDMP host documentation.

```
jbconfig
1) Configure an AlphaStor/SmartMedia Jukebox.
2) Configure an Autodetected SCSI Jukebox.
3) Configure an Autodetected NDMP SCSI Jukebox.
4) Configure an SJI Jukebox.
5) Configure an STL Silo.
What kind of Jukebox are you configuring? [1] 3
Enter NDMP Server name: ? netapp-760-01
Enter NDMP user name: ? root
Enter NDMP password (characters will not be echoed):
Communicating to devices on NDMP Server 'netapp-760-01',
this may take
a while...
```

When using Sun StorageTek EBS server's Administration interface or the **jbconfig** program to configure an autochanger, a new resource is created with the name specified. View the new Library resource in the Libraries section of the Sun StorageTek EBS Console window.

The *EMC NetWorker Command Reference Guide* or the UNIX man page provide information on the **jbconfig** command.

Chapter 3, "Configuring Libraries and Devices" provides more information on the Sun StorageTek EBS server's Administration interface.

Sharing libraries among NDMP clients

The Sun StorageTek EBS software permits the control and use of individual devices within a library to different Sun StorageTek EBS hosts within a datazone. If only NDMP hosts have access to a library, the autochanger has to be defined as NDMP. In this situation, the tape

drives within the library can still be shared and controlled by one host, and shared by the other hosts.

Robotic arm control

The library is configured for sharing during the initial library setup by using Sun StorageTek EBS server's Administration interface or the **jbconfig** program.

The **jbconfig** program must be run from the Sun StorageTek EBS server if the robotic arm is owned by the server or an NDMP host, or from a standard storage node if that node owns the robotic arm. The **jbconfig** program provides the following options.

- If the Sun StorageTek EBS server owns the robotic arm, run **jbconfig** from that server and choose **Autodetected SCSI Autochanger** when prompted.
- If an NDMP host owns the robotic arm, run **jbconfig** from the Sun StorageTek EBS server and choose **Autodetected NDMP SCSI Autochanger**.
- If a standard storage node owns the robotic arm, run **jbconfig** from that node and choose **Autodetected SCSI Autochanger**.

While running **jbconfig**, identify the following when prompted.

- The Sun StorageTek EBS hostname and device path of the server that controls the first library device.
- The controlling host for *each* device in the library.

Example 51 Robotic arm control

In a three-drive library shared among two NDMP hosts and a Sun StorageTek EBS server, Host1 controls the arm and the first drive.

- When prompted for the owner of Host 1, specify the hostname:device path for Host 1 as the host controlling the library device.
- When prompted for the owner of Drive 2, specify the hostname:device path for Host 2 for the owner of Drive 2.
- When prompted for the owner of Drive 3, specify the Sun StorageTek EBS server hostname:device path.

Chapter 3, "Configuring Libraries and Devices" provides more information on sharing libraries.

DDS support with NDMP

The Sun StorageTek EBS software includes support for DDS, which provides the capability to recognize shared drives. DDS enables Sun StorageTek EBS software to skip the shared drives that are in use and route the backups or recoveries to other available shared drives. For a list of NDMP disk and tape services that support DDS with the Sun StorageTek EBS software, refer to the *EMC NetWorker Software Compatibility Guide*.

The DDS feature may be enabled when configuring a library through the **jbconfig** program. Chapter 11, "Device Operations" provides more information about DDS Support with NDMP.

Enable the DDS feature

To enable DDS:

1. When prompted by the **jbconfig** program whether the drive is to be shared, type:

   `Yes`

2. When prompted, assign a hardware ID.
3. Respond to the prompts as required for the first host to have access to the shared drive.
4. When prompted whether this drive is to be shared with another host, type:

   `Yes`

5. When prompted, type the hostname and device path of the second host to have access to the shared drive.
6. Complete the prompts for the second device.

All the components of a SAN configuration must be compatible when DDS is enabled with the Sun StorageTek EBS 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.

Chapter 3, "Configuring Libraries and Devices" for additional information about standard DDS

Running nsrjb to load tapes

If **nsrjb** is run to load a tape, include the **-f** option and **rd** syntax in the device name. For example, if the NDMP hostname is *nas1* and the device is *nrst0a*, and a tape is to be loaded from slot 1, this syntax for loading a tape into that drive is used.

```
nsrjb -l -f rd="nas1:nrst0a (NDMP)" -S 1
```

Using nsrjb to load volumes in a jukebox connected to an NDMP server

To load the volume in slot 1 of jukebox *mylibrary* (connected to NDMP server *10.31.32.220*), use the **-J** and **-j** options.

```
nsrjb -J 10.31.32.220 -j mylibrary -l -S 1
```

Silo support

A Sun StorageTek EBS server with NDMP enabled supports silos controlled by either ACSLS or DAS controllers. Any silo behind the ACSLS or DAS is supported, since all robotic control and communication is done by ACSLS and DAS, which receive commands through TCP/IP from the Sun StorageTek EBS server.

Configuring an NDMP device resource

Chapter 3, "Configuring Libraries and Devices" provides details on how to configure the Device resource in the standard, non-NDMP Sun StorageTek EBS software. This section provides instructions for configuring the Device resource information when using a stand-alone device with NDMP.

The **jbconfig** program creates the NDMP device for NDMP devices in a library.

To configure an NDMP device (for non-storage node NDMP support) in the **Administration** window of the Sun StorageTek EBS Console:

1. In the **Administration** window, click **Devices**.
2. Right-click **Devices** in the navigation tree, and select **New**.
3. Complete these attributes in the **Create Device** window.
   a. For the **Name** attribute, specify the NDMP device by using this syntax:

      rd=*NDMP_computer:tape_device_pathname* **(NDMP)**

Note: You must add `(NDMP)` after pathname. If you do not add it, you will receive the following error message:
`NDMP device name shall be in rd=<snode>:<devname (NDMP)> format`

b. Set the **Target Sessions** attribute to **1**. Multiplexing is not supported on NDMP computers.

c. Complete the other attributes in the same manner as a non-NDMP client.

4. On the Configuration tab, complete these new device attributes.

   a. Select the checkbox to set the **NDMP** attribute. This attribute can only be set when creating the device. The NDMP attribute cannot be changed after the device has been created. To change the device configuration, first delete the configuration, then re-create it.

   b. For the **Remote User** attribute, type the NDMP account username for the computer where the NDMP library robotics or tape device resides.

   c. For the **Password** attribute, type the password for the NDMP user account. This password should match the password for the NDMP user account for the computer where the NDMP library robotics or tape device is installed.

5. Complete any other attributes in the same manner as a non-NDMP Sun StorageTek EBS device configuration.

Note: When configuring an NDMP Device resource, the Dedicated Storage Node attribute must be set to its default value of **No**. A dedicated storage node cannot be used with the new NDMP DSA feature, as a dedicated storage node will not work with **nsrdsa_save**. "Tips for troubleshooting storage nodes" on page 131 provides information about the Dedicated Storage Node.

Configuring an NDMP Client resource

"Task 6: Create a backup Client resource" on page 81 provides information on creating a Client resource for non-NDMP operations. This procedure covers issues related to specific attributes of the Client resource that are affected by NDMP.

To create an NDMP client in the Client resource:

1. From the **Administration** window, click **Configuration**.
2. In the expanded left pane, select **Clients**.
3. On the **General** tab, complete the following attributes.
   a. For the **Name** attribute, type the client name.
   b. For the **Save Set** attribute, list the file systems to be backed up. To back up all of the client's file systems, specify **All**.

      Note: If the NDMP server is running a version earlier than 3, the value **All** is not supported. The file systems must be listed manually.

      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.
   c. Complete the other attributes in the same manner as for a non-NDMP based backup.
4. On the **Globals (2 of 2)** tab, specify the appropriate storage node in the **Storage Nodes** attribute, depending on the type of backup to be performed.
   - For a local backup, type the NDMP server name to which the tape device is attached, listing this server first, unless **-M** is a parameter, in which case list the storage node hostname first.
   - For a three-party backup, list the destination server for the data being backed up first. The client backs up the data to each NDMP server or storage node in the order in which they appear in the attribute.

   Complete the other attributes in the same manner as for a non-NDMP based backup.
5. On the **Apps and Modules** tab, complete the following attributes.
   a. For the **Remote Access** attribute, list all computers that will store and retrieve entries for the client. In addition, list any computer that is used to browse entries for the NDMP computer. Since the Sun StorageTek EBS server stores the client file index entries for the NDMP computer, be sure that administrator@server is listed.

b. For the **Remote User** attribute, type the NDMP account username for the computer where the NDMP data module is installed. If this is administrator, be sure that the administrator has a password.

c. For the **Password** attribute, type the password for the NDMP user account. This password should match the password for the NDMP user account for the computer where the NDMP data module is installed.

d. For the **Backup Command** attribute, set the value as follows:

– For use with NDMP Tape Server:

```
nsrndmp_save -T backup_type
```

– For use with storage node:

```
nsrndmp_save -s server -c client -M -T backup_type
```

– To back up an NDMP client to a non-NDMP device, use the **-M** option.

The value for the NDMP backup type depends on the type of NDMP host. For example, NetApp, EMC, and Procom all support **dump**, so the value for the Backup Command attribute is:

```
nsrndmp_save -T dump
```

For a list of backup values supported by a particular NDMP host, refer to the appropriate vendor documentation.

e. For the Application Information attribute, list the NDMP environment variables. "Specifying application information" on page 761. provides a list of the variables and rules for typing their values.

The *NetWorker Command Reference Guide* or the UNIX man page provide information on the **nsrndmp_save** command.

6. On the **Miscellaneous** tab, select **Yes** for the **NDMP** attribute.

Note: Unless Yes is selected, the client will not be able to perform backups and recoveries. Any NDMP operations involving a non-NDMP client would have to be run from the Sun StorageTek EBS server. A non-NDMP client has no **ndmpd** service (or daemon) to connect to the server or another NDMP target client for a directed recovery.

7. Complete any other attributes appropriate for a Sun StorageTek EBS client configuration.

Specifying application information

In the Application Information attribute, list environment variables that the Sun StorageTek EBS server uses for an NDMP backup. For the backups to be successful:

- Separate the environment variable name from its value with an equal (=) sign, for example:

  ```
  HIST=y
  ```

- Type each environment variable in the Application Information attribute on a separate line.
- Click **Add** after each entry.

The values typed depend on the NDMP host.

Required variables for all NDMP hosts

These application variables are required for all (or most) NDMP hosts.

- HIST=y

  HIST determines whether a file history is created and required for index recovery. This value must be set to **y**. This is also the default setting.

  Note: For Mirapoint, this value must be set to *n* to deactivate creation of a file history during a backup. The (fromimagefull=) option informs level backups about the date of the last full backup. "Mirapoint" on page 800 provides more information.

Additional considerations for HIST=Y

When HIST=n, NDMP backups with Sun StorageTek EBS 7.2 and later are always run at level FULL, even with a correct schedule and UPDATE=y (for more information, refer to "UPDATE=y" on page 762).

A level FULL scheduled backup is performed if there is no previous fully browsable version of the same save set. Furthermore,

- If HIST=y, full and successive level backups perform properly.

- If HIST=y and at least one browsable save set exists, then future level backups of the same save set will back up at the specified level correctly.
- If HIST=n and no previous browsable backups of the save set exist, then all backups will be at level full, regardless of level specified. If there was at least one time this save set was backed up with HIST=y, all future HIST=n backups will be backed up at the correct level, not otherwise.
- With Sun StorageTek EBS 7.2 and later, if HIST=n (other variables being DIRECT=y and UPDATE=y), then all backups are performed at a level FULL, regardless of the level specified. Note that the index backup is performed at LEVEL 1, but the data is always at level FULL.
- The behavior is the same with NDMP DSA and non-DSA NDMP 3. If the level 1 backup is forced outside of Sun StorageTek EBS using **nsrndmp_save -T** with the **-l1** option, then a LEVEL 1 backup is performed.

◆ UPDATE=y

UPDATE forces the timestamp information to be updated with the time, date, level, and file system that is being backed up. This value must be set to **y**. If a file system is backed up at the same level as its previous backup, the old timestamp information entry is overwritten with the new information.

Note: The environment variables are case-sensitive. Check with the vendor for the latest information on supported values.

"Ancillary product information" on page 785 provides application information variables for a specific vendor, see the particular vendor section. These values were correct at the time of publication and *must* be set before a backup is performed. For the latest information, check with the NDMP vendor.

Additional considerations for NDMP operations

These sections provide information about various Sun StorageTek EBS settings to modify for NDMP operations.

Force incremental and interval attribute settings for NDMP tape server

If the Force Incremental attribute is set to Yes, and the Interval attribute is set to a value less than 24 hours, execution of a level

backup from the Console, NetWorker User program, or **save** program results in a full backup.

When the Interval attribute is set to less than 24 hours, the Sun StorageTek EBS server checks the value specified for the Force Incremental attribute. If the Force Incremental attribute is set to Yes, an incremental backup is being directed to be performed, regardless of the level specified for the backup. Since NDMP does not support incremental backups, a full backup is performed. To mitigate this problem, set the Interval attribute to 24 hours or later, or set the Force Incremental attribute to No.

Savegroup parallelism setting for NDMP tape server

By default, savegroup parallelism is set to zero, so the Sun StorageTek EBS server uses the value of the client parallelism. Rather than depending on client parallelism, set the Savegrp parallelism attribute for NDMP group backups to equal the number of available NDMP drives. If the Savegrp parallelism attribute is set to a higher value, there will not be enough drives to support all the save sets that are queued for a backup. Large save sets may fail due to the Inactivity Timeout limit. In addition, if more than one group is being backed up simultaneously, divide the number of drives by the number of groups. Use that value for each of the Savegrp parallelism attributes.

Performing NDMP backups

This section describes backups which use configurations with the NDMP Tape Server and storage node support for NDMP.

Tape server support for NDMP client backup

These sections describe how to perform an NDMP backup by using the **nsrndmp_save** command, Sun StorageTek EBS User, or the **save** program. The *EMC NetWorker Command Reference Guide* or the UNIX man pages provides details. The **nsrndmp_save** command is used in NDMP operations. This command supports most of the same arguments as **save** by referencing the UNIX man page. The **-T** *backup_type* option within **nsrndmp_save** identifies the NDMP backup type.

Prerequisites

Before performing a test backup, ensure that the following actions have been performed.

- Enable the NDMP Connection.
- Add the NDMP host to a backup group.
- Select a backup schedule with browse and retention policies.

  Note: Incremental backup settings are not supported. If the incremental setting is selected, the Sun StorageTek EBS server performs a full backup. However, differential level 1 backups may be scheduled to function like incremental backups. For example, complete a full backup on day 1 and level 1 backups on days 2 through 7.

- Select a volume pool for the data.
- Install an NDMP data module or enable the application server for NDMP operations. File systems on hosts without an NDMP data module cannot be backed up through the Sun StorageTek EBS NDMP Connection.
- Configure a non-NDMP device to back up the bootstrap and index information. If data is to be cloned, a second non-NDMP device must be configured to receive the cloned bootstrap and index information. Note that the index and bootstrap information is generated only if the scheduled group is backed up manually or through the schedule.

Perform the test backup by running a scheduled group backup.

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.

Performing a local backup by using the NetWorker User program

Use this program to perform a local backup that is not scheduled. A three-party backup *cannot* be performed with this program.

To perform a local backup:

1. In the NetWorker **User** program, click **Backup**.
2. In the **Backup** browse window, select the directories or files to be backed up.
3. To mark the files or directories for backup, click **Mark**.
4. To start the backup, click **Start**.

Performing an NDMP backup from the command prompt

Initiate NDMP backups from the command prompt by using the **nsrndmp_save** command. To browse files or directories that have been backed up, use the NetWorker User program**, nwrecover**, or the **recover** command.

Note: Configure the client in the Administration window of the Sun StorageTek EBS Console before performing a backup from the command prompt.

To perform a backup from the command prompt, use this syntax.

```
nsrndmp_save -T backup_type -s server_name -c client_name
-l backup_level -N name -g save_group -W width -LL local
path
```

where the **-T** option is the type of backup on the NDMP server, for example, *celestra*. The *EMC NetWorker Command Reference Guide* or the UNIX man page provide information on the **nsrndmp_save** command.

During backup and recovery operations, the NDMP host provides information on the amount of data moved. This status information is passed back to the Sun StorageTek EBS server and can be accessed through the Sun StorageTek EBS Console.

Note: When backups are performed through the **nsrndmp_save** command, the bootstrap is not automatically backed up. Without the bootstrap, a disaster recovery of an NDMP host cannot be performed. To perform a bootstrap recovery of an NDMP host, use the **savegrp -G** *group_name* command to back up the data. This command runs the specified backup group as if it were initiated from within the Sun StorageTek EBS server. All environment settings specified for that group, such as clients, pools, and schedules, are used.

Performing an NDMP DSA backup from the command prompt

Initiate NDMP DSA backups from the command prompt by using the **nsrndmp_save** command with the **-M** option.

To perform a backup from the command prompt, use the following syntax.

```
nsrndmp_save -T backup_type -s server_name -c client_name
-I hostname -l backup_level -M -N name -P hostname
-g save_group -W width -LL local path
```

where

- The **-M** option designates a backup to a non-NDMP device. The Sun StorageTek EBS software selects a storage node that has non-NDMP devices only. If there is more than one, ascending order is used.
- The **-P** *hostname* option designates a host for the non-NDMP based backup.
  - This option is valid only if the **-M** option is also present. If this option is absent, then the host is assumed to be the host that is local to the **nsrndmp_save** command. This *hostname* must be a valid Sun StorageTek EBS client and should be listed in the remote access list of the NDMP Client resource.
  - Unless the hostname of the client initiating the backup is added to the host for the non-NDMP based backup's /nsr/res/servers file (UNIX example), the backup will fail.

Performing a three-party backup from the command prompt

Three-party backups involve two NDMP servers. "Three-party backup with NDMP tape server" on page 742 provides details.

Note: For instructions on performing a remote backup, see "How to perform a remote backup with the SnapImage/NDMP TapeServer software" on page 805.

To perform a three-party backup:

1. Set up the servers as demonstrated in Figure 44 on page 743.
2. Install the Sun StorageTek EBS software on the server designated as the Sun StorageTek EBS server.

Note: If the Sun StorageTek EBS storage node is involved in a configuration which uses a conventional tape or disk device, then the bootstrap and file history can be saved conventionally. There is no need then for additional setup and explicit definition for the NDMP storage device.

3. Configure the Sun StorageTek EBS server.
   a. For the **Name** attribute of the **Client** resource, specify the hostname of the NDMP host that contains the data to be backed up.
   b. For the **Storage Nodes** attribute of the **Client** resource, verify that the destination server for the data being backed up is listed first.

      The client directs the data to the first storage node in the list that has an enabled device and is able to receive the data.
   c. Complete the rest of the required **Client** resource attributes and configure the Sun StorageTek EBS server in the same manner as for any other NDMP backup.

"Configuring the Sun StorageTek EBS server for NDMP operations" on page 750 provides instructions on configuring a Sun StorageTek EBS server.

Performing DAR backups

DAR is an optimized data recovery operation that allows data to be recovered in the middle of a tape set without the need to parse the tape set sequentially. This significantly reduces recovery time of large backups.

DAR is supported with SnapImage and the Sun StorageTek EBS DSA feature. Any data server that supports DAR is able to use this functionality when:

- DIRECT=y is set as an information variable
- DSA enabled storage node is used.

DSA acts as an agent in the performance of backups between a Sun StorageTek EBS server and any client without Sun StorageTek EBS software. An example of a DSA is an NDMP host that generates proprietary save data and sends that data to a Sun StorageTek EBS storage device to have a save set associated to it. Similarly, a DSA acts

as an agent in the recovery between a Sun StorageTek EBS server and any client without Sun StorageTek EBS software.

When performing a DAR recovery from a tape device, the tape device needs to support variable length records, refer to the tape device vendor's manual. If the tape device is not configured for variable length records, then DAR performance will degrade. This is true for any tape server as well.

The *EMC NetWorker Software Compatibility Guides* provides information on how to determine whether DAR is supported with a particular NDMP vendor.

To specify whether to use DAR, set the application information variables before the backup, specify these values:

```
DIRECT=y
HIST=y
```

DAR is supported in these NDMP configurations.

- Local
- Three-party

A DAR backup is performed in the same manner as a non-NDMP based backup. A DAR recovery is performed through the **recover** command, Sun StorageTek EBS User, or **nwrecover**.

Note: The **nsrndmp_recover** program cannot be used for a recovery with DAR.

Storage node support for NDMP client backup

In addition to being able to back up to an NDMP tape device, the **nsrndmp_save** program enables NDMP clients to back up to a Sun StorageTek EBS non-NDMP storage device (Sun StorageTek EBS storage node).

The *EMC NetWorker Command Reference Guide* or the UNIX man page provides more information on the **nsrndmp_save** command.

These sections describe how to perform an NDMP backup by using the **nsrndmp_save** command and the NetWorker **User** program. The *EMC NetWorker Command Reference Guide* or the UNIX man page provide information on the **nsrndmp_save** command.

The **nsrndmp_save** program starts the **nsrdsa_save** program (which is part of the Sun StorageTek EBS client software package) to back up data to a Sun StorageTek EBS storage device.

The **nsrndmp_save** program receives the file history, then processes and stores the indexes. The host on which the **nsrndmp_save** runs is the *index host*. The data to be backed up is generated by the NDMP data server. The host on which the **nsrdsa_save** program is running acts as a proxy to the NDMP host, which is the *proxy host*.

The **nsrndmp_save** program and **nsrdsa_save** program can be disassociated to run on different index and proxy hosts. This action helps improve performance and is highly scalable. For example, **nsrndmp_save** can be dedicated to run on a high-end host for processing millions of indexes, while **nsrdsa_save** can run on another host that has devices attached to it.

Data write technologies

The **nsrdsa_save** program, on execution, establishes a save session with the Sun StorageTek EBS server and listens on a port number. After a DSA receives data from the NDMP client, there are three different technologies that a DSA can use to write the data to a Sun StorageTek EBS storage device.

- Immediate save technology

  The DSA reads the data by using TCP/IP into shared memory. The **nsrmmd** program writes the data to a storage device from shared memory. Virtually all the data from the NDMP client goes directly to **nsrmmd**.

- Direct file access technology for a backup to disk

  The DSA directly writes the data to disk by passing the **nsrmmd** program. The **nsrmmd** program only plays a role in loading the volume. This is a highly efficient and high performance approach.

- Nonimmediate save technology (or remote save)

  The DSA reads the data by using TCP/IP into a local buffer and transmits it to the **nsrmmd** program that is also using TCP/IP. The **nsrmmd** program writes it to the storage device. This approach is inefficient and has slow performance.

Rather than nonimmediate save technology, configure the setup so that the storage device is local to the **nsrndmp_save** program, and that an immediate save technology is supported. If an immediate save technology is not supported, such as when the attached storage

device is remote, the DSA will fall back to a nonimmediate save technology.

Cloning and staging

After a backup is performed to a non-NDMP device, cloning and staging operations can be performed by using non-NDMP based Sun StorageTek EBS cloning and staging features and without involving the NDMP data or tape server. When NDMP save sets are marked as non-NDMP based Sun StorageTek EBS save sets and the volume data format is compatible with non-NDMP based Sun StorageTek EBS save sets, cloning and staging operations can be performed normally.

Multiplexing

Based on the server and device parallelism values, simultaneous multiple NDMP backup sessions can use a common non-NDMP device so that the device is kept busy. Similarly, multiple simultaneous NDMP recover sessions can read from a common volume if the data is multiplexed on that volume.

If the data is not multiplexed, then one NDMP save instance waits for the other to finish in the event that there are no available resources.

Multiple NDMP save instances can coexist with non-NDMP Sun StorageTek EBS save instances on the same volume. Any DSA save stream can join a non-NDMP Sun StorageTek EBS save stream. Both can be sent to the same volume.

Performing NDMP recoveries

NDMP save set and directory-level recoveries are performed with Sun StorageTek EBS User (Windows), **nwrecover** (UNIX), or the **nsrndmp_recover** command. NDMP file-by-file recoveries are performed through **nwrecover** or the **recover** command.

The **nsrndmp_recover** command supports most of the same command line operations as **recover(1m)**, in addition to others that specifically support NDMP. When data is retrieved with **nwrecover** (UNIX) or Sun StorageTek EBS User (Windows), the **nsrndmp_recover** command is selected automatically.

Note: NDMP does not support the recover options **Rename Recovered File** and **Discard Recovered File**. For both save set and file-by-file recoveries, *current files are always overwritten*.

Storage node support for NDMP client recovery

The **nsrndmp_recover** program recovers data from a Sun StorageTek EBS storage device when **nsrdsa_recover** is started on the same host. The **nsrndmp_recover** program is spawned by the following programs.

- **recover**
- **nwrecover**
- **winworkr** (Sun StorageTek EBS User graphical interface)
- **command line**

The **nsrndmp_recover** program sends data to the NDMP data server. The **nsrndmp_recover** and **nsrdsa_recover** programs cannot disassociate themselves as with the backup programs, **nsrndmp_save** and **nsrdsa_save**. The **nsrndmp_recover** program identifies that the save set is stored on a non-NDMP device and invokes the **nsrdsa_recover** program to perform a recovery from the non-NDMP device. There are no changes to the command line options for **nsrndmp_recover** when using storage node support for NDMP client recovery.

Note: Data written to a storage device by the **nsrdsa_save** program can be recovered only by using the **nsrdsa_recover** program. A tape that has NDMP data backed up with the **nsrdsa_save** program cannot be used by any other recovery tool.

Optimizing recoveries

When a backup is performed with the **-P**, **-I**, or **-M** option, and the intention is to optimize recovery time, launch a recovery on a host that has locally attached Sun StorageTek EBS devices. If a recovery is launched on a host other than the device host, then the recovery performance will degrade as data flows from the device host to the recovery host, and then from the recovery host to the NDMP data server.

Performing a save set recovery

These sections describe how to perform a save set recovery from the Sun StorageTek EBS User or **nwrecover** programs, or from the command prompt. For , see "Performing a directed recovery" on page 777 provides more information about directed recoveries.

How to perform a save set recovery with the NetWorker User program

To perform a save set recovery from the User program:

1. From the **Operation** menu, select **Save Set Recover**.
2. Select the **NDMP** client with the save set data to be recovered and click **OK**. The save sets available for recovery appear in the **Save Set** dialog box.
3. From the **Save Set Name** list box, select the specific save set with the data to be recovered. After a save set is selected, that save set's backup instances appear in the **Version** window.
4. Select the particular instance of the save set's backup to recover.
5. If needed, specify a subdirectory in the **Path to Recover** attribute, and click **Start**.

Note: The **Save Set Recover** window contains a **Recover Options** window for NDMP recoveries. This window is used if you are recovering the data to a path other than the save path. "How to perform a directed save set recovery" on page 777 provides instructions on using this window.

How to perform a save set recovery with the Sun StorageTek EBS Recover program

To recover save sets by using **nwrecover**:

1. In **nwrecover** program, select **Save Set Recover** from the **Options** menu.
2. In the **Save Set** window, click the save set to be recovered.

   You can select an unlimited number of versions for the same save set for simultaneous recovery.

   - If you do not choose a save set from the list, the Sun StorageTek EBS server automatically selects the last save set on the list.
   - If you do not choose a save set instance to recover, the Sun StorageTek EBS server automatically selects a save set instance.
3. When all the files to recover are selected, click **Recover**.

Perform a save set recovery from the command prompt

To perform a save set recovery from the command prompt, type:

```
nsrndmp_recover -c client -s server -S ssid/cloneid
```

For example, to recover data to the original location client named *venus*, with a server named *mars*, from the command prompt, type:

```
nsrndmp_recover -c venus -s mars -S 2485094913
```

Destructive save set recoveries

When NDMP data is backed up using the SnapImage Module, save set recoveries are generally destructive. With a destructive save set recovery, data can recovered to the original location or an alternate location, and in either case all existing data is overwritten. Essentially, destructive save set recoveries overlay the data at File system level and re-impose the saved image of the File system.

Destructive save set recoveries are performed using the **nsrndmp_recover** command with the **-r** *raw_device* and **-m** *mountpoint* options. When **-r** *raw_device* is specified, **-m** *mountpoint* is mandatory (however, when **-m** *mountpoint* is specified, **-r** *raw_device* is not required). For example, on a Microsoft Windows system, the following command would perform a destructive save set recovery to the P:\ drive:

```
nsrndmp_recover -s mars -c venus -m P:\ -r P:\ -S
2674606849
```

On a UNIX system, the following command would perform a destructive save set recovery to the /dev/dsk/c1t1d0s0 device, mounted at /:

```
nsrndmp_recover -s mars -c venus -r /dev/dsk/c1t1d0s0 -S
2674606849 -m /
```

Note: Excluding the **–r** option from the above commands, and then identifying the pathname of the file to be restored, will make the SnapImage save set recoveries non-destructive and operate at the file/directory level rather than the filesystem level. This will still overwrite existing files on the destination that have the same names as those in the restore list, but will leave other data untouched on the file system. This may be useful when doing a directory level recovery in a high density filesystem environment to restore everything under one directory when it is not practical to pass all the file names that use the NetWorker User or **nwrecover** graphical interface.

The *EMC NetWorker Command Reference Guide* or the UNIX man page provide information on the **nsrndmp_recover** command.

Performing a directory-level recovery

These sections describe how to perform a directory-level recovery from Sun StorageTek EBS User or from the command prompt. "Performing a directed recovery" on page 777 provides information about directed recoveries.

How to perform a directory-level recovery from the NetWorker User program

To perform a directory-level recovery from the User program:

1. From the **Operation** menu, select **Save Set Recover**.
2. Select the **NDMP** client from the **Source Client** window with the save set data to be recovered and click **OK**. The save sets available for recovery appear in the **Save Set** dialog box.
3. In the **Save Set Name** list box, select the save set with the data to recover. The save set's backup instances appear in the **Version** window.
4. Select the particular instance of the save set's backup to recover.
5. Click **Start**.

Note: The **Save Set Recover** window contains a **Recover Options** window for **NDMP** recoveries. This window is used to recover data to a path other than the save path. "How to perform a directed save set recovery" on page 777 provides instructions on using this window.

How to perform a directory-level recovery from the command prompt

To perform a directory-level recovery:

- From a UNIX command prompt, type:

```
nsrndmp_recover -c client -s server -S ssid -v
verify_flag paths
```

- From a Microsoft Windows command prompt, type this command:

```
nsrndmp_recover -s server -c client -R recover_target
{-r rawdev -S ssid -m mntpnt|-F} -v verify_flag paths
```

The *EMC NetWorker Command Reference Guide* or the UNIX man page provides information on the **nsrndmp_recover** command.

Performing a file-by-file recovery

A file-by-file recovery only recovers a select group of files or directories. These sections describe how to perform a file-by-file recovery.

How to perform a file-by-file recovery from the recover program

To recover files by using **nwrecover**:

1. To open the Sun StorageTek EBS Recover window, type the following at the command prompt:

   ```
   nwrecover
   ```

   You can select a file that is listed in the Sun StorageTek EBS Recover window for recovery. To perform a recovery:

   a. Browse the index for a backups.

   b. Select various versions of the file to recover.

   c. View your marked files.

   d. Relocate files upon recovery (optional).

   e. Recover files from another authorized system (optional).

   An authorized system is a client from which you have permission to recover data.

2. If necessary, select a server by clicking **Change Server** on the speedbar.

3. If necessary, change the level of displayed file information.

   - To expand the directory display one level, double-click the directory folder in the directory display.
   - To expand a directory branch to display subdirectories, click a directory folder and select **Expand Branch** from the **tree** menu.
   - To collapse a branch, double-click the directory folder in the display.
   - To view detailed information for files in a directory, select **File Details** from the **View** menu.

4. Select the files (or directories) to be recovered, and click **Mark**.

5. From the **View** menu, select **Volumes** to see whether the volumes are online for the recovery.

6. Click **Start** to begin the recovery.
7. In the **Conflict Resolution** dialog box, click **OK**.
8. When the recover completion time message in the **Recover Status** window appears, click **Cancel** to close the window.

How to perform a file-by-file recovery with the NetWorker User program

To perform a file-by-file recovery with the User program:

1. Click **Recover** to open the **Source Client** dialog box.
2. Select the source client with the data to recover, and click **OK**.

   The local client is the default selection.
3. In the left pane of the **Recover** window, select the save set that contains the required files.
4. In the right pane, select the file or files to recover.
5. Click **Mark** to select these files for recovery.
6. To recover the files to another path on the same host:
   a. From the **Options** menu, select **Recover Options**.
   b. In the **Relocate Recovered Data To** text box, specify the path.
   c. Click **OK**.
7. Click **Start**.

How to perform a file-by-file recovery from the command prompt

To perform a file-by-file recovery:

1. From the command prompt, type:

   recover -s *server*
2. When prompted, type the directory to browse, for example:

   cd \.
3. Issue the appropriate commands to add the necessary files to the recover list. The *EMC NetWorker Command Reference Guide* provide a complete list of command options.
4. When all the required files have been added, type:

   recover

On a Microsoft Windows a Windows system program named **recover.exe** is also located in *%SystemRoot%\System32*. To avoid using the Windows version of **recover.exe**, do one of the following:

- Include the full path at the command prompt.
- Ensure that <Sun StorageTek EBS_install_path>*bin* is listed before *%SystemRoot%\System32* in the PATH environment variable.

Performing a directed recovery

Sun StorageTek EBS software with NDMP supports directed recoveries that enable you to recover backup data to a host or path other than the one used to back up the original data.

Note: The operating system of the target host (*hostname*) must be the same as the one used to back up the data, and the target host must be a client of the Sun StorageTek EBS server. If it is not a client, the server cannot obtain the correct username and password, and the data cannot be recovered to that host. For example, if the original data was backed up to a NetApp filer, the data must be recovered to the same NetApp filer or another NetApp filer that is a Sun StorageTek EBS client.

Note: To perform a directed recovery with an Auspex filer, Auspex requires that you specify an existing directory as the destination directory. If the directory does not already exist, the recovery will not complete.

How to perform a directed save set recovery

You can perform a directed save set recovery by using the NetWorker **User** program or the **nsrndmp_recover** command.

- When using Sun StorageTek EBS User:
  1. From the **Operation** menu, select **Save Set Recover**.
  2. In the **Source Client** window, select the source client.
  3. In the **Save Sets** dialog box, select the backup version to recover.
  4. Click **Recover Options**. The **NDMP Recover Options** window appears.
  5. In the **Destination Client** text box, specify the client to receive the recovered data.

6. In the **Relocate Recovered Data To This Path** text box, specify the recovery path for the volume. If recovering to a Windows platform, this attribute should contain the same drive letter as the **Raw Device** attribute.

   Note: If you are recovering to a filer or UNIX platform, specify *target_hostname::/mount_point*.

7. In the **Relocate Recovered Data To This Raw Device** text box, specify the drive letter of the destination drive. If you do not specify the raw device name of the destination, the save set is recovered and the files are overwritten, but the current file system structure is retained.
8. Click **OK**.
9. In the **Save Set** window, click **OK**.

- When using **nsrndmp_recover**, type:

  ```
  nsrndmp_recover -c client -m
  target_hostname::/mount_point -S ssid
  ```

  For the mount point, specify the NAS computer hostname. If you are using Data ONTAP, you might need to add a slash (/) after the mount point. For example, *target_hostname::/mount_point/*.

Note: In these command sequences, the host specified is the target host to which the data is recovered. If a host is not specified, the data is recovered to the original application server. The *mount_point* variable remains the previously defined mount point of the application server.

The *EMC NetWorker Command Reference Guide* or the UNIX man page provide more information on the **nsrndmp_recover** command.

How to perform a directed file-by-file recovery

You can perform a directed file-by-file (Windows) recovery by using either the NetWorker User program or the **recover** command.

- When using the NetWorker **User** program:
  1. Click **Recover** to open the **Source Client** dialog box.
  2. Select the source client with the data to be recovered, and click **OK**.
  3. Select the source client with the data to recover, and click **OK**.

4. In the left pane of the **Recover** window, select the save set that contains the files to recover.
5. In the right pane, select the files to recover.
6. Click **Mark**, then **Start**.

- When using the interactive **recover** program:

1. From the command prompt, type:

   recover -s *server*

2. When prompted, type the directory to browse, for example:

   ```
   cd \.
   ```

3. Run the appropriate commands to add the necessary files to the recover list. The *EMC NetWorker Command Reference Guide* or the UNIX man page provide a complete list of command options.
4. To recover the files to a different location, type:

   relocate *target_hostname*::*mount_point*

   Note: In these command sequences, the host specified is the target host. If a host is not specified, the target is the original application server. For the mount point, specify the NAS computer hostname. If you are using Data ONTAP, you might need to add a backslash (\) after the mount point. Here is an example, *target_hostname::\mount_point*.

5. When all the required files have been added, type this command:

   ```
   recover
   ```

The files are recovered to the specified location.

How to perform a directed index recovery

You can perform a directed index (UNIX) recovery using the **recover** program.

1. From the command prompt, type:

   recover -s *server*

2. When prompted, type the directory to browse, for example:

   ```
   cd /.
   ```

3. Run the appropriate commands to add the necessary files to the recover list.

   The *EMC NetWorker Command Reference Guide* or the UNIX man page provide a complete list of command options.

4. To recover the files to a different location, type this command:

   relocate *target_hostname::/mount_point*

   Note: In these command sequences, the host specified is the target host. If a host is not specified, the target is the original application server. For the mount point, specify the NAS computer hostname. If you are using Data ONTAP, you might need to add a slash (/) after the mount point, for example, *target_hostname::/mount_point/*.

5. When all the required files have been added, type:

   recover

   The files are recovered to the specified location.

NDMP cloning

NDMP cloning is set through the same Sun StorageTek EBS clone interface as non-NDMP based backups. Chapter 9, "Cloning" provides information on how to clone data.

Note: To support cloning, the source NDMP host must be running NDMP version 3 or later. The destination server can be running any version of NDMP, but a volume cloned with NDMP earlier than version 3 cannot be cloned again to another volume.

Cloning between storage nodes

Cloning between storage nodes is accomplished by an NDMP tape server on the source node by reading from a volume, and another NDMP tape server on the target node by writing to a volume. However, cloning from a NDMP tape device to a non-NDMP tape device, and vice-versa, is not supported.

If the save set to be cloned was backed up by **nsrndmp_save** by using **nsrdsa_save** (that is, the DSA feature save set flags are used to differentiate the save set records, which include *N* and *s*), then use

nsrclone, and not **nsrndmp_clone,** to clone these save sets. The **nsrdsa_save** generated save sets can be cloned to any Sun StorageTek EBS storage device other than an NDMP tape device.

The source node is determined by the location of a source volume. The location of the source volume is determined either by its currently mounted location or, if unmounted, by its location field (see the **mmlocate** man page). The target node of a clone is found by the Clone Storage Nodes attribute of the Client resource in descending priority.

Cloning functionality

The following functionality applies to NDMP cloning.

- The Sun StorageTek EBS Console and the command line interface support volume cloning, save set cloning, and the cloning of all save sets as part of a scheduled group.
- NDMP tapes from one NDMP host can be cloned to another NDMP host of the same type. For example, tapes from a NetApp filer that has a library attached can be cloned to another NetApp filer or to the same filer.
- NDMP cloning can copy data from one type of tape device to another, for example data from DLT to AIT.
- Once the cloning is complete, the Sun StorageTek EBS server verifies that the cloned copy was successfully created. In the Volumes window that contains the save set information, the cloned save set is marked cloned.
- A cloned volume or save set can be automatically or recovered manually. "Volume cloning" on page 291 and "Recovering cloned data" on page 301 provide more information.

NDMP index and bootstrap cloning

Index and bootstrap cloning is performed only through a non-NDMP-based cloning device. Index and bootstrap cloning can be achieved either through the command prompt by using **nsrclone** or by using the Sun StorageTek EBS Console through a scheduled group with automatic cloning.

"Cloning a save set automatically" on page 295 provides more information on automatic cloning.

A full device setup for NDMP cloning involves four devices when automatic cloning is enabled.

- An NDMP tape device that is used for NDMP backups and is the source device for NDMP cloning.
- An NDMP cloning device that is used to clone an NDMP save set or volume.
- A non-NDMP device that is used for backing up bootstrap information and indexes.
- A non-NDMP cloning device that is used for cloning indexes and bootstrap from an NDMP backup.

Note: The indexes and bootstrap are non-NDMP save sets, so a non-NDMP device is needed to clone those non-NDMP save sets. The indexes and bootstrap both go to the same volume, unless they are configured to go to separate devices.

When cloning is performed manually, a minimum of two devices are required, one to read the existing data and one to write the cloned data. "Cloning a save set manually" on page 292 provides more information.

NDMP Path-to-Tape cloning for regular save sets

NetWorker 7.5 and higher supports NDMP path-to-tape cloning of regular save sets with the **nsrndmp_clone** command.

Using NDMP cloning, regular NetWorker save sets from a DL3D 1500 or 3000 Virtual Tape Library (VTL) running version 1.1.2, can be copied directly to a physical tape library, which is attached to the VTL. This process, also known as path-to-tape, frees up network bandwidth and offloads NetWorker storage node I/O resources because save set data does not go through the storage nodes. NetWorker manages the data movement and catalogs the physical tape copy in the media database as a clone instance. Both the VTL and the physical tape devices are visible to the NetWorker storage node and restores can be performed directly from either copy.



IMPORTANT

This feature is available with EMC DL3D models 1500 and 3000 at version 1.1.2, Version 1.1.2 is available on request from your EMC DL3D account representative.

The path-to-tape feature can also be set up to move data from one physical tape library to another physical tape library. Additionally, all NDMP clone options work with the path-to-tape feature.

Requirements and limitations for NDMP cloning regular save sets

Ensure that the following are configured before starting the NDMP cloning process:

- An NDMP storage node resource is required for the VTL. The VTL must have drives configured for NDMP.
- An NDMP storage node resource is required for the physical tape library. The physical tape library must have drives configured for NDMP.
- The Client resource that corresponds to the source save sets or volumes must have its **Clone Storage Node** attribute set to the destination storage node host, which is where the clone data is to be written. The destination storage node is usually the hostname of the NDMP storage node for the physical tape library.
- Do not use the "device tape flags" variable (if you are in Diagnostic Mode) to restrict the volume size when performing NDMP (including path-to-tape) cloning.

Path-to-Tape cloning for tapes with different block sizes

From release 7.6 onwards, NetWorker supports path-to-tape cloning between tapes of different block sizes.

Note that, however, NetWorker does not support cloning of a path-to-tape clone instance that is cloned from a tape of a different device file size than that of the original source tape.

For example, consider the case in which you create a back up save set (referred to as clone instance #1 by default) on an LTO-3 tape that has a block size of 128KB and a device file size of 4GB. Using path-to-tape, you then create a clone save set (clone instance #2) to another LTO-1 tape that has a block size of 64KB and a device file size of 1GB. If you use path-to-tape to clone "clone instance #2" to another LTO-1 tape, the cloning attempt will fail with the following error messages:

```
nsrndmp_clone: Error - Incorrect tape file size
<current_tape_file_size>
nsrndmp_clone: Set device file size to
<recommended_value> in source device properties
You will have to change the source device file to the
value suggested by the nsrndmp_clone error messages.
```

In this scenario, you could perform one of the following options, however, it is recommended that you perform option 2:

- Option 1
  Change the source (clone instance #2) device file size to the value suggested in the error message and restart the clone operation.



IMPORTANT

After the cloning operation, change the source device file size back to its original value; otherwise, subsequent cloning attempts using this source device would create clone instances that may not be recoverable.

- Option 2
  Perform the second clone operation from the original source, which is clone instance #1 in this example.

Using the nsrdmp_clone command for cloning regular save sets

For NDMP cloning of regular save sets, use the following command:

`nsrndmp_clone -J` *ndmp_storage_node* **`-p -b`** *clone_pool* **`-S`** *ssid1 ssid2...*

where:

- The **-J** option is used to specify the ndmp clone to load the source device on a specified host.
- The **-p** option is used to create a clone of regular save sets.
- The **-b** option is used to specify the pool type.
- The **-S** option is used to identify the save set ID(s).

Note: NDMP supports cloning of opaque save sets with same block or device file sizes from the original size.

Recover can be done from the original backup or from the clone copy. You can perform a regular file-by-file or save set recover. A regular storage node resource is required both for the VTL and the physical library. The VTL and the physical tape library must have drives configured for SCSI. The recover process will automatically load the volume from an NDMP drive to a non-NDMP drive (SCSI path).

Note: The recovery process automatically selects a SCSI path for recovery. It will not select the NDMP path because the save sets to be recovered are non-NDMP save sets.

Ancillary product information

These sections cover limitations and configuration requirements for specific vendors and other NetWorker modules:

- "Filers that act as data servers" on page 785
- "Filers that act as tape servers only" on page 802

Filers that act as data servers

This section contains configuration information for these vendor products that act as NDMP data servers:

- "Network Appliance" on page 786
- "EMC Celerra" on page 791
- "EMC Centera" on page 795
- "EMC IP4700" on page 795
- "Auspex" on page 796
- "Procom NetFORCE" on page 799
- "Mirapoint" on page 800
- "BlueArc" on page 800

Backup on NDMP servers is dependent on the data server provider type. Servers use the commands listed in Table 110 on page 785.

Table 110 NDMP server examples (1 of 2)

Server	Command	Function
EMC	dump	Used for backing up filer data to tape. Also used to back up directories, quota trees, or entire volumes.

Table 110 NDMP server examples (2 of 2)

Server	Command	Function
NetApp	dump	Used for backing up filer data to tape. Also used to back up directories, quota trees, or entire volumes.
Auspex	gtar	Used for archiving the contents of a directory to a file or tape.
EMC SnapImage Module	SIDF or dd	SIDF (System Independent Data Format) backups allow users to interchange media with operating system backup utilities and other tape backup products from vendors that comply with the SIDF standard. A dd-compatible command backs up all blocks in the file system. Data is written in dd image.

Only those variables listed in the sections corresponding to a given NAS filer are supported by the Sun StorageTek EBS software. Variables documented in NAS filer reference guides do not necessarily work for Sun StorageTek EBS backups. In these cases, the results may be unpredictable.

For example, the variable FILESYSTEM is supported by NetApp. With Sun StorageTek EBS 7.2.*x*, this value is passed to NetApp, which translates the path as required. This is the path that is backed up, not the one specified by the save set.

Network Appliance

The following sections provide information about limitations, autochanger configuration, and application information variables for the NetApp filer. For a complete list of supported Data ONTAP versions, refer to the *EMC NetWorker Software Compatibility Guide.*

Features and limitations

These sections cover NDMP issues that are specific to NetApp hosts.

Note: NetApp passwords are limited to eight characters.

Backup and recovery of large number of files

On Network Appliance clients with an operating system of Data OnTap 6.4 and later, all metadata is read from tape before the files are recovered. As such, with large save sets (save sets with 20 million files or more), the recovery time for a file can exceed three hours. This also applies for backup, as the metadata for the whole volume is recorded onto tape even for a single file backup.

UTF-8 versions supported

The Sun StorageTek EBS software supports the UTF-8 format with CIFS clients. NFS clients of a NetApp filer are only supported if the NFS clients can generate UTF-8 encoded data. Support for UTF-8 is specified in the Application Information attribute of the Client resource.

UTF8=n is the default setting and should remain set.It is no longer required to set the UTF8=y variable in the Application Information attribute during a Sun StorageTek EBS client configuration for NDMP clients. If the UTF8=Y variable is set during an NDMP client backup and the backup contains pathnames with non-ascii characters, an index based recovery of this backup will fail with the error "RESTORE: could not create path <*pathname*>".

If upgrading from any Sun StorageTek EBS release previous to 7.3 Service Pack 3, remove the **UTF8=y** variable from the client configuration in order to ensure successful backup and recovery operations:

1. Connect to the Sun StorageTek EBS server using the Sun StorageTek EBS Management Console (NMC).
2. Click on the **Configuration** button.
3. Click on **Clients**.
4. In the clients window, right-click on the NDMP client and select **Properties**.
5. Click the **Apps and Modules** tab.
6. In the Application Information section, remove **UTF8=Y** if defined.
7. Click **OK**.

For existing backups that were run with **UTF8=y** in Sun StorageTek EBS 7.3 Service Pack 3 or later, pathnames with non-ascii characters must be recovered using either a save set recovery from the command line, or NDMP Directory restore.

Note: Ensure that the NAS filer is configured for UTF8 characters sets. Contact your NAS vendor for configuration requirements."

Single save sets support only one code set

A single save set supports data that belong to only one code set. If you have data in multiple code sets, you must create multiple save groups.

The 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, so filenames from these languages can be backed up in a single save set.

Required NDMP version when running DAR on a NetApp host

To run DAR on a NetApp host, you must have NDMP version 3 or later. Recoveries fail if an earlier NDMP version is used.

How to verify the NDMP version

To verify the NDMP version:

1. Log in to the NetApp host as root or as a Windows Administrator.
2. Type this command:

   ```
   ndmpd version
   ```

   The NDMP version is displayed.

How to change the NDMP version

To change the NDMP version:

1. Log in as root or as Windows Administrator.
2. Stop the **NDMP** process by typing:

   ```
   ndmpd off
   ```

   The NetApp host confirms that the NDMP process has stopped.
3. To change the **NDMP** version, type:

   ```
   ndmpd version 4
   ```

   The NetApp host confirms that the **NDMP** version is now set to 4.

   Note: The NDMP version can also be set to 3.

4. To restart the **NDMP** process, type:

   ```
   ndmpd on
   ```

The NetApp host confirms that the NDMP process has started, and DAR can now be run.

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.

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.

How to find the device pathname for a NetApp host

The pathname of the media device is required before the autochanger for NDMP backups and recoveries can be configured. Use the pathname to configure the autochanger with the **jbconfig** program, as described in "Configuring library support for NDMP" on page 751.

Note: This procedure is only an example. For the correct syntax and usage, refer to the NetApp documentation.

To identify the media device pathname:

1. Log in as root or as Windows Administrator
2. Verify that the tape device is installed on the NetApp filer by typing this command.

   ```
   sysconfig -v
   ```

   The host responds with:

   ```
   NetApp Release 6.2: Sun Mar 12 13:29:37 PST 2009
   System ID: 0016784666 (NetAppHost1)
   slot 6: SCSI Host Adapter 6 (Qlogic ISP 1040B)
   Fast/Wide, Differential
   Firmware Version 4.65 Clock Rate 60MHz.
   4: QUANTUM Powerstor L200  001F
   5: QUANTUM DLT7000         2255
   ```

3. Determine the device pathname by typing this command.

   sysconfig -t

   The host responds with a list of media device names.

   ```
   Tape drive (6.5)  Digital DLT7000
   nrst0a -  no rewind device, format is: 85937 bpi 70 GB
   (w/comp)
   nrst0m -  no rewind device, format is: 85937 bpi 35 GB
   ```

   where:

   - **(6.5)** indicates slot 6 and SCSI ID 5. This information must match the output in the **sysconfig -v** command.
   - **nrst0a** is the media device name.

   If the NDMP tape server is running version 2 or does not support NDMP_CONFIG interfaces, determine the autochanger handle by typing this command.

   sysconfig -m

   The host responds with the devices on the host.

   ```
   Media changer (6.4)  QUANTUM  Powerstor L200 mc0 - media
   changer device
   ```

   where **mc0** is the autochanger handle.

Note: When **jbconfig** prompts for the hostname and device path for a NetApp device, do *not* type a slash before the device name. Although the **jbconfig** program completes without errors if the slash is typed, NetApp will not recognize the tape device or autochanger, and you will be unable to label the tape.

Application information variables

In addition to the standard application information variables, HIST=y and UPDATE=y, that are required for all vendors, Table 111 lists variables specific to NetApp that may be required. Type these variables in the Application Information attribute of the Client resource, as described in "Configuring an NDMP Client resource" on page 758.

Table 111 NetApp application information variables

Variable	Required	Values and meaning
FILESYSYTEM=*path*	No	This variable, if passed to NetApp, will supersede the save set specification.
DIRECT= y/n	Yes for DAR operations	This value must be set to **y** if you are performing a DAR recovery. If you are not performing a DAR recovery, the value is set to **n**. DAR is supported on Data ONTAP 6.1 and later. "Performing DAR backups" on page 767 provides additional information about DAR.
EXCLUDE=	No	This string value determines which files are excluded from backup. These rules apply. • The string is a filename. Only filenames are accepted, not absolute paths. • The asterisk (*) is accepted as a wildcard, but it must be 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 filename. Spaces are not allowed. • You can specify up to 32 strings.
EXTRACT_ACL= y	Yes to recover ACLs	To recover access control lists (ACLs) when using DAR with a NetApp host, you must set the value to **y**. The default is **n**.
UTF8=n	For UTF-8 formatted data	The default is **n** and should remain set

EMC Celerra

These sections provide information about limitations, autochanger configuration, and application information variables for the Celerra host. The *EMC Software Compatibility Guide* provides a complete list of supported versions.

Celerra considerations with NDMP

These sections discuss NDMP issues that are specific to Celerra filers.

Celerra backup of root file system

Instead of specifying a save set of All for an Celerra backup operation, the separate file systems should be named in the save set definition, excluding the root, or "/" file system. Including the root file system causes the client index to fail.

Using an earlier version of NDMP on Celerra Data Movers

Celerra data movers are run in the default configuration, which enables support for the latest version of NDMP. Servers autonegotiate

with other NDMP devices so that the NDMP versions are compatible. At times, the autonegotiation function fails. In this case, Celerra can be forced to use an earlier version of NDMP to test whether the devices can still function. This can also help to isolate NDMP issues by determining if the problem happens with different versions.

Lower the NDMP version on Celerra

The NDMP version must be set individually on each data mover that requires the modification.

To set the NDMP version on Celerra data movers so that the data mover negotiates only at the specified NDMP version or an earlier version:

1. Log in as root or as a Windows Administrator to the Celerra Control Station.
2. At the command prompt, type:

   cd /nas/server/*slot_#*

   where *slot_#* refers to the server number.

   For example, /nas/server/slot_2/netd contains the configuration files for server_2.
3. Open the netd file in a text editor such as vi.
4. At the end of the NDMP string in the netd file, type:

   proto=*ndmp_version_#*

   For example, the resulting NDMP string is this when setting the NDMP version to 3.

   ndmp port=10000 proto=3
5. Save the file.
6. Reboot the data mover by typing this command.

   server_cpu *server_name* **-reboot -m now**

How to set the ntape parameter

To help avoid tape drive issues, set the *ntape* parameter on the Celerra filer. For every tape drive discovered on a particular data mover there must be a corresponding value set for the NDMP *ntape=#* value. For example, if a data mover has five tape drives configured on it, set the parameter to NDMP ntape=5.

To modify the parameter:

1. With a text editor, open the /nas/server/*slot_#*/*param* file, where *slot_#* correlates directly to the server number. For example, /nas/server/slot_2/netd contains the configuration files for server_2.
2. Add or modify this line.

   NDMP ntape=n

   Note: The value of *n* cannot be greater than 8.

3. In order for the changes to take effect, reboot the data mover.

How to find the device pathname for a Celerra host

The pathname of the media device is required before the autochanger for NDMP backups and recoveries can be configured. Use the pathname to configure the autochanger with the **jbconfig** program, as described in "Configuring library support for NDMP" on page 751.

Note: This procedure is only an example. For the correct syntax and usage, refer to the Celerra documentation.

To identify the device pathname, perform *one* of these actions.

- Re-create all devices and view all nondisks by typing the following commands at the command prompt:

  ```
  Celerra server_devconfig movername -create -scsi -all
  Celerra server_devconfig movername -list -scsi
  -nondisks
  ```

- Obtain the device pathname through the Celerra Administrator program.

  To access the **Celerra Administrator** program:

  a. In the browser window, type the following URL:

     ```
     http://EMCcontrolhostIP:8000/top_level.htm.
     ```

  b. Type your username and password.

  c. Use the **Celerra Administrator program** to list and/or determine SCSI devices.

Note: If the NDMP tape server is running version 2 or does not support NDMP_CONFIG interfaces, locate the autochanger handle in the Celerra Administrator program.

Application information variables

In addition to the standard application information variables, HIST=y and UPDATE=y, that are required for all vendors, Table 112 lists Celerra-specific variables that may be required. Type these variables in the Application Information attribute of the Client resource, as described in "Configuring an NDMP Client resource" on page 758.

Table 112 Celerra variables

Variable	Required	Values and meaning
DIRECT= y/n	Yes for DAR operations	This value must be set to **y** if you are performing a DAR recovery. If you are not performing a DAR recovery, set the value to **n**. "Performing DAR backups" on page 767 provides additional information about DAR.
EMC_EDIR*nn*=	No	This string value identifies a directory to be excluded from the backup. The asterisk (*) is accepted as a wildcard as the last character. Multiple directories an be included by incrementing the number; for example: • EMC_EDIR01=/fsX/DIRx • EMC_EDIR02=/fsX/DIRy This variable is supported with Celerra version 5.5 and greater.
EMC_EFILE*nn*=	No	This string value determines which files are excluded from backup. The asterisk (*) is accepted as a wildcard, but it must be the first or last character in the string (or both). Multiple files an be included by incrementing the number; for example: • EMC_EFILE01=*mp3 • EMC_EFILE02=temp* This variable is supported with Celerra version 5.5 and greater.
OPTIONS=NT	Yes for ACL restoration	This value must be set to make sure all ACLs are recovered when NT data is contained in the backup path. In addition to setting this variable, **axttrp** must exist in the /nas/server/slot_#/netd file.
SNAPSURE=Y	No	This option is used to decrease the backup time for an Celerra host when DX is integrated in the Sun StorageTek EBS environment.

Char and Block Special Files

During recovery, the Celerra filer skips char and block special files. The following error message will be displayed:

```
Warning: /fs1/SPE_REL/my.char_file has an unknown file
type, skipping
```

EMC Centera

The *EMC Centera Backup and Recovery Module Release 2.01 Installation and Administrator's Guide* and the *EMC Centera Backup and Recovery Module Release Supplement* provide information about the EMC Centera® Backup and Recovery Module software .

The *EMC NetWorker Software Compatibility Guide* provides a complete list of supported versions.

EMC IP4700

These sections provide information about autochanger configuration and application information variables for the EMC IP4700 host. The *EMC NetWorker Software Compatibility Guide* provides a complete list of supported versions.

How to find the device pathname for an EMC IP4700 host

The pathname of the media device is required before the autochanger for NDMP backups and recoveries can be configured. The device pathnames can be found through the IP4700 Administrator web interface. Use the pathname to configure the autochanger with the **jbconfig** program, as described in "Configuring library support for NDMP" on page 751.

To obtain the autochanger and tape drive device pathnames:

1. In the browser window, access the **EMC IP4700 Administrator** program at this URL.

   http://*IP4700name***/**

2. Click **Tape Drives**.
3. Type the **tape list** command.
4. Use the resulting values to configure the autochanger and tape drive devices for the Sun StorageTek EBS software.

Note: If the NDMP tape server is running version 2 or does not support NDMP_CONFIG interfaces, locate the autochanger handle in the EMC IP4700 Administrator program.

Application information variables

In addition to the standard application information variables, HIST=y and UPDATE=y, that are required for all vendors, Table 113

on page 796 lists IP4700-specific variables that may be required. Type these variables in the Application Information attribute of the Client resource, as described in "Configuring an NDMP Client resource" on page 758.

Table 113 EMC IP4700 variables

Variable	Required	Values and meaning
FAST_ INCREMENTAL= y	No	Performs a fast incremental backup. This function works only if fast incremental is enabled on the volume. You must have the SnapView/IP option to perform a fast incremental backup. The default is **y**.
SNAPSHOT= y	No	Backs up a snap copy of the file system. This function only works if snap copy is enabled on the volume. You must have the SnapView/IP option to use snap copy. The default is **n**, unless FAST_INCREMENTAL is set to **y**. Then, SNAPSHOT is set to **y**.
NO_ACLS=y	Yes	On Microsoft Windows systems, set this value to **y** when backing up a Windows system to ensure that ACLs are not backed up. Do not set this value to n when backing up Windows files with ACLs. The default is **n**.

Auspex

These sections provide information about limitations, autochanger configuration, and application information variables for the Auspex host. The *EMC NetWorker Software Compatibility Guide* provides a complete list of supported versions.

Auspex-Specific considerations with NDMP

These are NDMP issues that are specific to Auspex filers.

Auspex NS3000 requirements

In order to support Auspex NS3000:

- Connect the robot to channel 6 of the SCSI controller. To confirm that the robot is connected to the correct channel, inspect the device filename under the /dev/raxac directory. It should display a channel similar to *fsp0c6t0*.
- To use snapshots, a cache device must be reserved. The appropriate Auspex documentation provides more information about snapshots.

How to back up large files with Auspex

By default, support for files larger than 2 GB is not enabled for an Auspex backup. This allows for use of a standard **gtar** tape format.

To enable support for files larger than 2 GB:

1. Open the /usr/AXndmp/etc/config file on the NDMP host.
2. Remove the comment marks for the line containing this text.

   ```
   ALLOW_TAR_EXTENSION
   ```

3. Save the modified file.

Note: If you modify the system to support large files, the tapes will be in nonstandard **gtar** format.

DD-Type backup limitations and requirements

DD-type backups have these limitations and requirements.

- HIST=y is not supported. Set this value to HIST=n.
- DD-type backups are only supported with NetOS V4.1P3 or later.

Autochanger configuration tips

Configure the Autochanger resource to have these minimum timeout values.

- Eject Sleep: 10
- Cleaning Delay: 60
- Unload Sleep: 10
- Load Sleep: 90
- Deposit Timeout: 15
- Withdraw Timeout: 15
- Port Polling Period: 3

Auspex limitations

These limitations apply to Auspex backups and recoveries.

- Root partition backups are not supported. Do not back up the root partition.
- File-by-file recoveries should not exceed 10,000 files per recovery.

  To recover a larger number of files, perform *one* of these steps.

- Set up multiple save sets for a given client to break the file system into directory-level save sets rather than whole volumes.
- Perform a save set recovery by using the NetWorker User program, **nwrecover**, or the **nsrndmp_recover** command from the command prompt.

How to find the device pathname for an Auspex host

The pathname of the media device is required before the autochanger for NDMP backups and recoveries can be configured. The device pathname for an Auspex computer is located in the *messages* file at /var/adm/messages. Use the pathname to configure the autochanger with the **jbconfig** program, as described in "Configuring library support for NDMP" on page 751.

The Auspex documentation provides information about adding a SCSI device.

The *messages* file contains the list of devices, for example:

```
Jan 18 10:32:22 nightcap FSP_0[1611]:fsp0c7t0:STK
9714
[ Medium_Changer ][ Fast ] [ min:1792 ]
Jan 18 10:32:22 nightcap FSP_0[1611]:fsp0c7t1:Quantum
DLT4000  [ Tape ]
[ Not Negotiated ] [ min:1808 ]
Jan 18 10:32:22 nightcap FSP_0[1611]:fsp0c7t2:Quantum
DLT4000  [ Tape ]
[ Not Negotiated ] [ min:1824 ]
```

Note: If the NDMP tape server is running version 2 or does not support NDMP_CONFIG interfaces, determine the autochanger handle from the device list.

Application information variables

In addition to the standard application information variables, HIST=y and UPDATE=y, that are required for all vendors. Table 114 on page 799 lists Auspex-specific variables that may be required. Type these variables in the Application Information attribute of the Client resource, as described in "Configuring an NDMP Client resource" on page 758.

Table 114 Auspex Variables

Variable	Required	Values and Meaning
DIRECT= y/n	Yes for DAR	This value must be set to **y** if you are performing a DAR recovery. If you are not performing a DAR recovery, set the value to **n**. The default is **n**. "Performing DAR backups" on page 767 provides additional information about DAR.
SNAPSHOT= y	No	Enables or disables snapshots. The default is **n**.

Procom NetFORCE

These sections provide information about autochanger configuration for the Procom host. The *EMC NetWorker Software Compatibility Guide* provides a complete list of supported versions.

How to find the device pathname for a Procom NetFORCE host

To determine the device pathname, log in as root and type the following command:

```
status dm
```

NetFORCE responds with a list of autochanger and media device names. Use the appropriate device name when configuring the autochanger with the **jbconfig** program, as described in "Configuring library support for NDMP" on page 751.

Note: If the NDMP tape server is running version 2 or does not support NDMP_CONFIG interfaces, you must determine the autochanger handle. On Procom NetFORCE, the SCSI device name format is *isp1*tSSL[L], where isp1 is the autochanger handle. The Fibre Channel device format is *ffx1*tSSL[L], where *ffx1* is the autochanger handle.

Postconfiguration requirements for Procom NetFORCE host

After using the **jbconfig** program to configure an autochanger for a Procom NetFORCE host:

1. Reset the autochanger by typing the following command.

   ```
   /usr/sbin/nsrjb -HE
   ```

2. To read previously-labeled tapes back into the database, inventory the autochanger.

   ```
   /usr/sbin/nsrjb -IE
   ```

Mirapoint

These sections provide information about requirements and autochanger configuration for the Mirapoint host. The *EMC NetWorker Software Compatibility Guide* provides a complete list of supported versions.

Requirements

These requirements apply to the Mirapoint system.

- After restoring a full backup, reboot the Mirapoint system. The system does not need to be rebooted after an incremental recovery.
- The application information variable MIRA_OPTIONS= (fromimagefull=) must be set. The *(fromimagefull=)* value allows the user to perform a full image backup and then perform message (file) based backups by using the date of the image in order to perform the selection.

Device pathname for a Mirapoint host

The device pathname is required to successfully run the **jbconfig** program. The device pathname is /dev/nrst*n*, where *n* starts at 0 and increases one number for each tape drive. This value is constant. Use the pathname to configure the autochanger with the **jbconfig** program, as described in "Configuring library support for NDMP" on page 751.

Once the autochanger is configured, certain attributes must be modified in the Autochanger resource window. Set the **Load Sleep** and **Eject Sleep** attributes to these values:

- Load Sleep to 90
- Eject Sleep to 30

Note: If the NDMP tape server is running version 2 or does not support NDMP_CONFIG interfaces, you must type the autochanger handle, /dev/ch0, when running **jbconfig**.

BlueArc

BlueArc systems provide network-attached, fixed storage solutions to share, manage, and protect data. This ensures continuous availability and simplified data management.

When configuring a BlueArc system, keep these points in mind:

- When performing NDMP recovery operations, a $__NDMP__ directory may be created at the root level of the file system to which the data is being recovered. This directory holds the file list that the Sun StorageTek EBS server passes for an index recovery, if that file list contains more than 1,024 filenames. Do not change the directory and its contents during an active recovery operation. If no recovery is in progress, the directory may be deleted.
- NDMP_BLUEARC_FH_NAMETYPE=UNIX, an application information variable (also known as an NDMP Environment Variable), must be specified in the Client resource.
- The backup command type, also known as the backup application name, is **dump**.
- While performing NDMP backup and recovery operations, a message similar to the following may appear:

  ```
  NDMP session—Unknown environment variable name
  ignored.
  ```

  You can ignore this message, as it does not impact the operation.

The BlueArc's documentation provides a complete information on configuring a BlueArc system for NDMP operations.

SnapImage software

The NetWorker SnapImage Module is a high-performance storage solution that backs up and recovers high-density file systems across a LAN or SAN. It is based on NDMP and provides full-image, block-level backups and recoveries, allowing large amounts of data to be backed up quickly, which reduces the backup window. The SnapImage software backs up high-density file systems faster than non-NDMP based file-based backup systems. It supports full backups as well as differential backups, which back up only the blocks that have changed since the last full backup. It can recover data at the block level, and also can recover specific files and directories at the volume level.

The SnapImage Module Installation and Administrator's Guide and Release Notes provide instructions on installing, configuring, and using the SnapImage software.

Filers that act as tape servers only

This section provides information about filers that act as tape servers only.

DinoStor TapeServer

These sections discuss the DinoStor TapeServer and the autochanger configurations.

Configuring the DinoStor TapeServer

The DinoStor software provides a web-based interface for administering and controlling the TapeServer settings.

Ensure that the following are set when configuring the server:

- On the Password page of the Configure tab, specify the password for administering the DinoStor TapeServer, not the NDMP password.
- When configuring the DinoStor TapeServer, set the NDMP version to 3 and the port number to 10000 on the NDMP page of the Configure tab.
- When connecting an autochanger to a DinoStor TapeServer, shut down the DinoStor TapeServer and power down. Attach the new devices and reboot the system. You can reboot or shut down the system from the web interface.
- Only SCSI tape devices are supported.
- DDS is not supported, since the DinoStor TapeServer is not fibre-equipped.
- GigE and 10/100 Base-T are supported.

How to specify device handles for a DinoStor TapeServer

When configuring the autochanger with the Sun StorageTek EBS Administration window, the DinoStor TapeServer has one hostname with multiple device handles. Use the hostname to configure the autochanger with the Sun StorageTek EBS Administration window, as described in "Configuring library support for NDMP" on page 751.

To specify the device handles:

1. Access the **DinoStor TapeServer interface**.
2. Click the **Configure** page.

3. Click the **SCSI** tab.

When running **jbconfig**, specify the DinoStor hostname as the device that controls the robotics and the device handles as the additional devices.

How to perform a remote backup with a DinoStor TapeServer

To perform a remote backup with a DinoStor TapeServer:

1. Set up the configuration with an NDMP host, a Sun StorageTek EBS server, and the DinoStor TapeServer with a library attached.
2. Configure the Sun StorageTek EBS server.
   a. For the **Name** attribute of the **Client** resource, specify the hostname of the NDMP server that contains the data to be backed up.
   b. For the **Storage Node** attribute of the **Client** resource, verify that the hostname of the DinoStor server is listed first. The client directs the data to the first storage node in the list that has an enabled device and is able to receive the data.
   c. Complete the rest of the required **Client** resource attributes and configure the Sun StorageTek EBS server as you would for any other NDMP backup.

"Configuring an NDMP Client resource" on page 758 provides instructions on configuring the Client resource.

SnapImage or NDMP TapeServer Module

These sections discuss the SnapImage/NDMP TapeServer software configurations and requirements.

Finding the device pathname for a SnapImage Module device

The required steps for determining the device pathname are different for Solaris and HP-UX computers. Complete the instructions for the appropriate platform. Use the pathname to configure the autochanger with the Sun StorageTek EBS Console Administration window as described in "Configuring library support for NDMP" on page 751.

How to identify the device pathname for a Solaris computer

To identify the device pathname:

1. Log in as root on the host.
2. Determine the media device pathname by typing the following command.

```
ls -l /dev/rmt/*
```

The computer responds with the list of device filenames.

```
lrwxrwxrwx   1 root     root          03 Nov 17 12:03
/dev/rmt/0mbn ->
../../devices/pci@1f.0/scsi@2/st@1,0:mbn
```

where:

- */st@1, 0* is the SCSI ID field for the media device.
- /dev/rmt/0mbn is the media device pathname.

If the NDMP tape server is running version 2 or does not support NDMP_CONFIG interfaces, determine the autochanger handle by typing the following command:

```
ls -l /dev/rsjb*
```

The computer responds with a list of *rsjb* device handles.

```
lrwxrwxrwx   1 root     other        45 Jan 11 13:35
/dev/rsjb2 ->
../devices/pci@1f.0/scsi@2,1/sjb@1,0:8r
```

where:

- :8r identifies the autochanger robotics.
- */sjb@1, 0* is the SCSI ID field for the rsjb device.
- */dev/rsjb2* is the autochanger handle.

How to identify the device pathname for an HP-UX computer

To display the device pathname, log in as root on the host, use one of the following methods:

- Use the **ioscan** command by typing the following command:

  ```
  ioscan -fnC tape
  ```

- Use the SAM utility.

  a. Start the utility.

  b. Select **Peripheral Device**s.

  c. Select **Tape Drives**.

  d. Highlight the appropriate tape drive and select **Show Device Files f**rom the **Action** menu.

Note: If the NDMP tape server is running version 2 or does not support NDMP_CONFIG interfaces, use the **ioscan** command or SAM utility to obtain the autochanger handle.

How to perform a remote backup with the SnapImage/NDMP TapeServer software

To perform a remote backup:

1. Install the Sun StorageTek EBS software on the server designated as the Sun StorageTek EBS server.
2. Install the SnapImage Module/NDMP TapeServer software on the server designated as a tape server.
   a. If you are backing up the data to the Sun StorageTek EBS server, install the SnapImage Module/NDMP TapeServer software on that Sun StorageTek EBS server.
   b. If you are backing up the data to a different NDMP server, install the SnapImage Module/NDMP TapeServer software on the host to which the data is being backed up.

   The *EMC NetWorker SnapImage Module Installation and Administrator's Guide*, and its Release Notes provide instructions on installing the SnapImage Module/NDMP TapeServer software.

   Note: Do not use the configuration instructions in the SnapImage Module guide. Instead, use the configuration instructions provided in this chapter for the NDMP TapeServer software.

3. To allow other servers to access the SnapImage server:
   a. Log in as root on the host.
   b. Run the **ndmp_passwd** command, located in /usr/ndmphome/bin.
      – Use this syntax:

        ndmp_passwd -u *user*

        where *user* is the root or administrator name of the SnapImage server.
      – When prompted, type the password.
4. Configure the Sun StorageTek EBS server.

a. For the **Name** attribute of the **Client** resource, specify the hostname of the NDMP server that contains the data to be backed up.

b. For the **Storage Node** attribute of the **Client** resource, verify that the destination server for the data being backed up is listed first.

   The client directs the data to the first storage node in the list that has an enabled device and is able to receive the data.

c. Complete the rest of the required **Client** resource attributes and configure the Sun StorageTek EBS server as you would for any other NDMP backup.

"Configuring an NDMP Client resource" on page 758 provides instructions on configuring the Client resource.

22

SNMP Module

This chapter covers these topics:

- SNMP traps 808
- Configuring NetWorker SNMP notifications 808
- Configuring SNMP management software 811

SNMP traps

The Sun StorageTek EBS Simple Network Management Protocol (SNMP) Module allows Sun StorageTek EBS servers to send notification messages to SNMP management agents.

SNMP-enabled network management software must be configured to accept traps from the Sun StorageTek EBS server. For detailed information about SNMP management operations refer to your network management documentation.

The NetWorker SNMP Module uses traps to communicate Sun StorageTek EBS event notifications to SNMP management stations. A trap is an unsolicited notification sent from the SNMP agent (such as the Sun StorageTek EBS server) to the SNMP event manager.

The types of traps that the Sun StorageTek EBS server sends are determined when the NetWorker SNMP notification is configured within the Sun StorageTek EBS server. Typical traps include warnings, critical errors, and other messages from the Sun StorageTek EBS server. "Configuring NetWorker SNMP notifications" on page 808 provides instructions on configuring the NetWorker SNMP notification.

Configuring NetWorker SNMP notifications

Sun StorageTek EBS software provides notifications to a variety of resources about Sun StorageTek EBS 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 *Sun StorageTek EBS Installation Guide* provides information on enabling and licensing the Sun StorageTek EBS software.

Command line options for nsrtrap

The NetWorker SNMP Module uses the **nsrtrap** program to communicate SNMP traps from the Sun StorageTek EBS server to the SNMP management software. Table 115 lists the command line options that can be included in the Action attribute when the SNMP notification is configured.

Table 115 Command-line options for nsrtrap

Option	Description
-c *community*	Specifies the SNMP community that is authorized to receive traps from the Sun StorageTek EBS 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 Sun StorageTek EBS 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 *trap_type*	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 Sun StorageTek EBS 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 Sun StorageTek EBS notification.
-s *specific_type*	A generic setting that can be used to identify the type of trap the Sun StorageTek EBS 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 Sun StorageTek EBS 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 Sun StorageTek EBS 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 Sun StorageTek EBS 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 Sun StorageTek EBS 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 809 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 810 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:

   /usr/sbin/nsr/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 809 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 Sun StorageTek EBS 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.

Sun StorageTek EBS 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 Sun StorageTek EBS 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 Sun StorageTek EBS servers, an icon of each Sun StorageTek EBS 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 Sun StorageTek EBS 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 Sun StorageTek EBS 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 809 provides more information about setting the **-s** *specific-type* option.

23

Support for Microsoft Automated System Recovery

This chapter covers these topics:

- Microsoft Automated System Recovery 814
- Sun StorageTek EBS support for ASR 814
- ASR Limitations and special considerations..... 816
- ASR backup..... 817
- Creating an ASR disk 819
- Using the ASR disk to recover a Sun StorageTek EBS client 823

Microsoft Automated System Recovery

Microsoft Automated System Recovery (ASR) is a Windows XP and Windows Server 2003 feature that enables backup and recovery application 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.

Note: The ASR feature in Microsoft Windows Vista and Microsoft Windows Server 2008 is not supported in this Sun StorageTek EBS release.

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.

Sun StorageTek EBS support for ASR

Sun StorageTek EBS software supports ASR on Sun StorageTek EBS clients that run Windows XP Professional or Windows Server 2003. ASR is not supported for Sun StorageTek EBS servers or storage nodes.

ASR is *not* supported for Sun StorageTek EBS clients configured with a shared disk. This applies to stand-alone clients, as well as clients operating in an MSCS cluster environment. The *EMC NetWorker Disaster Recovery Guide* provides disaster recovery information about these configurations.

Sun StorageTek EBS ASR save set

The Sun StorageTek EBS 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,

The ASR save set can be backed up by itself, or as a component of the save set All. Like the Sun StorageTek EBS SYSTEM save sets, ASR is backed up automatically at level full.

Usage overview

Because ASR is the recommended Sun StorageTek EBS disaster recovery method for Windows XP Professional and Windows Server 2003 Sun StorageTek EBS clients, perform ASR backups on a regular basis.

You can perform an ASR backup by using either of these methods:

- As part of a regularly scheduled group backup, configured by using the Sun StorageTek EBS Administrator program.
- In a manual backup by using the NetWorker User program.

You can create an ASR disk by using either of these methods:

- By restoring the ASR save set on the Sun StorageTek EBS client computer.
- By performing a directed recovery of the ASR save set to a different Sun StorageTek EBS client computer.

In the event of a disaster on the Sun StorageTek EBS client, use the ASR disk to perform an ASR recovery. This must be done locally, on the failed Sun StorageTek EBS client computer.

Note: If a current ASR disk or ASR save set is not available for the failed Sun StorageTek EBS client computer, you must use the legacy Sun StorageTek EBS disaster recovery method. The EMC NetWorker Disaster Recovery Guide provides more information .

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 822 provides details.

ASR Limitations and special considerations

This section describes limitations and special considerations that apply to ASR backup and recovery.



IMPORTANT

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 819 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 Sun StorageTek EBS client" on page 823 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 Sun StorageTek EBS recovery operation. Chapter 13, "Recovering 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.

Note: To perform an ASR recovery, you must have an official Microsoft Windows installation CD for the version of Windows you are recovering. The EMC NetWorker Disaster Recovery Guide and the OEM documentation provide information If you have only an OEM recovery CD that includes the Windows installation files.

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.

ASR backup

Include ASR in a scheduled group backup for each Windows XP Professional and Windows Server 2003 Sun StorageTek EBS client. When you use the ASR disk to recover a Sun StorageTek EBS client computer, any data and configuration changes made since the last ASR backup will be lost.

Note: The Automated System Recovery (ASR) feature in Microsoft Vista and Microsoft Windows Server 2008 is not supported with the 7.4 or later releases of the software.

How to include the ASR or VSS ASR DISK save set in a scheduled backup

Backing up save set All is the preferred method of performing scheduled ASR backups. This method ensures that the ASR or VSS ASR DISK save set includes information for all of the client's disk volumes. If individual save sets are selected, only the selected save sets are restored during an ASR recovery.

To include the ASR or VSS ASR DISK save set in a client's scheduled backup:

- In the Save Set attribute for the Client resource, enter All.
- At a minimum, select the following save sets:
  - The drive letter for %SystemDrive% (for example, C:\)
  - If VSS is not licensed and enabled, select:
    - ASR:
    - SYSTEM STATE:
    - SYSTEM DB:
    - SYSTEM FILES:
  - If VSS is licensed and enabled, select:
    - VSS ASR DISK:
    - VSS SYSTEM BOOT:
    - VSS SYSTEM SERVICES:
    - VSS SYSTEM FILESET:
    - VSS USER DATA:

Chapter 2, "Backing Up Data" provides more information about specifying save sets for a scheduled backup.

How to include ASR or VSS ASR DISK save set in a manual backup

Backup of data by selecting My Computer is the preferred method of performing manual ASR backups. This method ensures that the ASR or VSS ASR DISK save set includes information for all of the client's disk volumes. By selecting individual save sets, only the selected save sets are restored during an ASR recovery.

To include the ASR or VSS ASR DISK save set in a manual backup:

- In the NetWorker User program, click the Backup button and select My Computer to save all data.

- At a minimum, select the following save sets:
  - If VSS is not licensed and enabled, select:
    - ASR:
    - SYSTEM STATE:
    - SYSTEM DB:
    - SYSTEM FILES:
  - If VSS is licensed and enabled, select:
    - VSS ASR DISK:
    - VSS SYSTEM BOOT:
    - VSS SYSTEM SERVICES:
    - VSS SYSTEM FILESET:

Chapter 2, "Backing Up Data" provides more information about manual backups.

Creating an ASR disk

The information in this section refers to using the NetWorker User program 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 Sun StorageTek EBS 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 Sun StorageTek EBS 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. "ASR backup" on page 817 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 Sun StorageTek EBS 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 821 provides details.

To create an ASR disk locally:

1. Log in with administrator privileges to the Sun StorageTek EBS 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 Sun StorageTek EBS 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 Sun StorageTek EBS 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 822.

Creating an ASR disk by using directed recovery

If the Sun StorageTek EBS 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 Sun StorageTek EBS administration, and in cases where you cannot access the %SystemRoot%\Repair\nsr directory on the Sun StorageTek EBS client because the system is damaged.

To create an ASR disk by using directed recovery:

1. Log in with administrator privileges to a Sun StorageTek EBS client computer.
2. In the NetWorker **User** program, click **Recover**.
3. In the **Source Client** dialog box, select the Sun StorageTek EBS 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 Sun StorageTek EBS 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 Sun StorageTek EBS 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 822.

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 Sun StorageTek EBS 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 *EMC NetWorker Disaster Recovery Guide*.

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 Sun StorageTek EBS backup save sets to recover.
- The Sun StorageTek EBS 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 819 provides more information.

  If a current ASR disk or ASR save set is not available for the failed computer, use the legacy Sun StorageTek EBS disaster recovery method. The *EMC NetWorker Disaster Recovery Guide* provides details.

- The Microsoft Windows XP Professional or Windows Server 2003 installation CD for the computer you are recovering.

  Note: Recovery CDs provided by computer manufacturer are not supported for ASR recovery. "OEM recovery CDs are not supported" on page 817 provides more information about this restriction.

- The latest Sun StorageTek EBS backup for the computer you are recovering.

Note: If you routinely move Sun StorageTek EBS 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 or UNIX man pages.

Performing an ASR recovery

To perform an ASR recovery on a Sun StorageTek EBS 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.

   Note: 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 Sun StorageTek EBS 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 819 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** or UNIX man pages.

   The VSS components and certain components cannot be correctly restored during ASR recovery. "Components that require special handling after an ASR recovery" on page 825 provides information on how to ensure that all necessary components are properly recovered.

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 Sun StorageTek EBS 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 Sun StorageTek EBS 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 and VSS SYSTEM Save Sets" has more information about VSS writers.

Appendix D, "Additional Features of the Microsoft Windows Server" provides information about how the Sun StorageTek EBS software handles the Windows system state and system-protected files.

To verify the Sun StorageTek EBS client recovery:

1. Reboot the Sun StorageTek EBS client computer and verify that the Sun StorageTek EBS Remote Exec and Sun StorageTek EBS Power Monitor services have started.
2. Use the Windows Event Viewer to examine the event logs for errors. In particular, check for these:
   - 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.

24

Volume Shadow Copy Service

This chapter covers these topics:

- Overview of VSS 830
- VSS and the backup process 831
- Controlling VSS from Sun StorageTek EBS software 833

Overview of VSS

If the NetWorker Module for Microsoft Applications is installed on the client computer, information in this chapter may be superseded by information in the NetWorker Module for Microsoft Applications documentation. The *EMC NetWorker Module for Microsoft Applications Administration Guide provides* more information about the NetWorker Module for Microsoft Applications.

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 Sun StorageTek EBS 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 Sun StorageTek EBS software releases 7.2 and later, if a VSS client license is present, Sun StorageTek EBS software uses VSS technology to create snapshot backups of volumes and exact copies of files, including all open files. Microsoft Vista or Microsoft Server 2008 do not require a VSS client license. 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 no longer 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 backups with the Sun StorageTek EBS PowerSnap Module software. The EMC PowerSnap Module documentation provides more information.

Note: VSS backups are not deduplicated.

VSS and the backup process

In VSS terms, Sun StorageTek EBS 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 a particular piece of 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. At this time, Microsoft supports snapshots at the volume level only.

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 45 on page 832 provides a graphical representation of the VSS backup process.

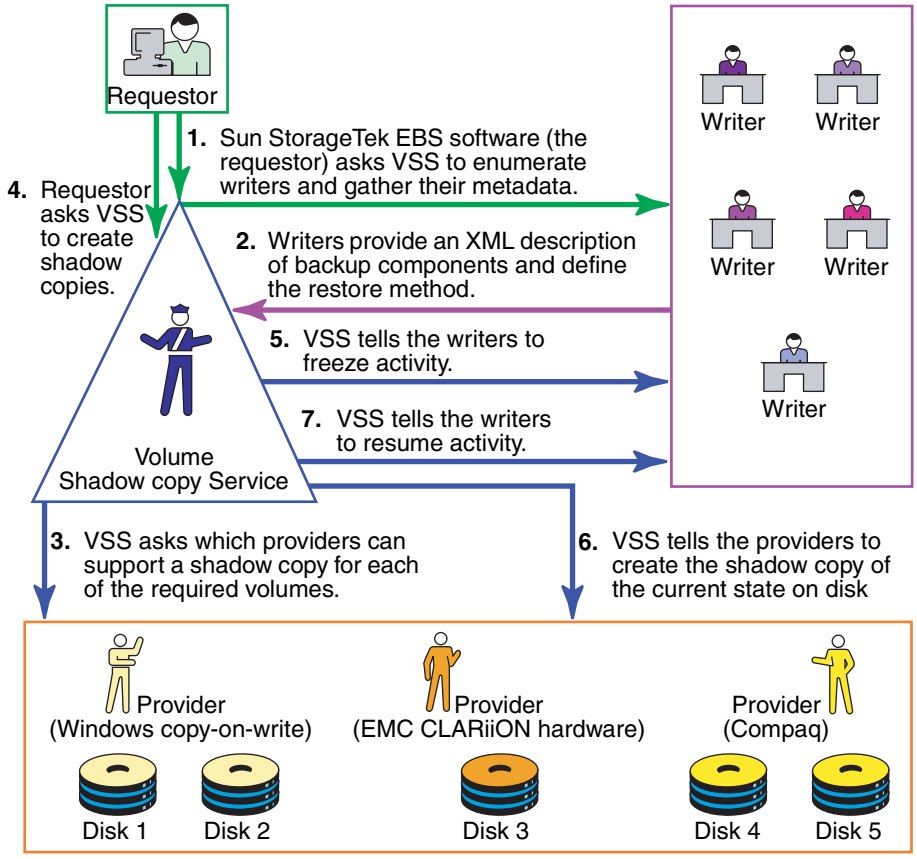


Figure 45 VSS backup process

Sun StorageTek EBS 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 physical 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.

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, Sun StorageTek EBS software maintains a list of supported writers in the NSRLA database of the client machine. When backing up data, the software confirms that a VSS client license exists, and then it 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, Sun StorageTek EBS 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, Sun StorageTek EBS 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 Sun StorageTek EBS software

By default, Sun StorageTek EBS software release 7.2 and later uses VSS technology to back up a client that is also running release 7.2 or later, and is licensed to use VSS. For VSS SYSTEM save sets, this means Sun StorageTek EBS software uses VSS for most save sets and

writers. Appendix A, "SYSTEM and VSS SYSTEM 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 Sun StorageTek EBS software uses VSS. For example, if you need to disable a writer or completely disable VSS. You can control VSS from the Administration window, the NetWorker User program, or the command prompt.

Note: With Windows Vista and Windows Server 2008, VSS is the only available backup method and cannot be disabled.

These sections provide more information:

- "Controlling VSS from the Administration window" on page 834 provides details on how to control VSS from the Administration window.
- "Controlling VSS from the Sun StorageTek EBS client" on page 836 provides details on how to control VSS from the NetWorker User program.
- "Control VSS from the command-prompt" on page 838 provides details on how to control VSS from the command prompt.
- "Globally disabling VSS" on page 838 provides details on how to disable VSS globally.

Note: The documentation that accompanied the VSS client license describes how to control VSS when backing up a client that is running Sun StorageTek EBS software release 7.2 or later, where the Sun StorageTek EBS *server* is running a prior release of Sun StorageTek EBS software.

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 116 on page 840, then click **OK**.
   - Separate multiple commands with a semicolon (;).
   - If the **Save Operations** attribute is left blank, Sun StorageTek EBS software backs up data by using VSS if a VSS client license is present.

Notes:

1. The **Save Operations** attribute does not support Sun StorageTek EBS Module save sets. If a Sun StorageTek EBS Module save set name is entered in the window, the backup fails.
2. 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.
3. When typing a writer name in the **Save Operations** attribute, 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, 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:
     - Sun StorageTek EBS 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.
4. Use the **Save Operations** attribute only for clients running Sun StorageTek EBS software release 7.2 or later. If anything is entered in this attribute for a client that is running an earlier Sun StorageTek EBS software release, the backup will fail.

Controlling VSS from the Sun StorageTek EBS client

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 Sun StorageTek EBS 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 116 on page 840, then click **OK**. Separate multiple commands with a semicolon (;).

   If the **Local Save Operations** dialog box is left empty, Sun StorageTek EBS software backs up data by using VSS if a VSS client license is present:

   - Local Save Operations does not support Sun StorageTek EBS Module save sets. If a Sun StorageTek EBS 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:
     - Sun StorageTek EBS 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 Sun StorageTek EBS client or the Console server by using the **-o** option and the Save Operations commands in Table 116 on page 840, but only while performing a **save, savefs,** or **nsrarchive** 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* or the UNIX man pages provides more information about the **save**, **savefs**, and **nsrarchive** commands and command options.

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.

Note: With Windows Vista and Windows Server 2008 and later Windows releases, VSS is the only available backup method and cannot be disabled.

To disable VSS:

1. Log in as root or as Windows Administrator on the Sun StorageTek EBS server.

   To disable VSS for all Sun StorageTek EBS 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

Table 116 on page 840 lists the commands and syntax used to control VSS, along with a description of what each command does. These commands may be entered in any of these locations:

- Save Operations attribute of the Administration window on the Console server.
- Local Save Operations dialog box of the NetWorker User program on the Sun StorageTek EBS client.

- Command prompt on the Sun StorageTek EBS client or the Console server.

Table 116 VSS commands

Task	Entry	Result
Enable VSS.	This attribute should be left empty.	Leaving the attribute empty will result in Sun StorageTek EBS software automatically using VSS (if VSS is licensed).
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.
Disable VSS for a particular drive.	VSS:*drive*:\=off For example: VSS:c:\=off	The specified drive is backed up by using the legacy method.
Disable an individual writer.	VSS:*writer*=off where *writer* is the name of the writer to disable. For example: **`VSS:WINS Writer=off`**	The application data served by *writer* is not saved, unless a Sun StorageTek EBS Module exists for that application and is installed and configured on the system. **Note:** When a writer is disabled, Sun StorageTek EBS software still processes the writer so its files are skipped during the file system backup.
Instruct Sun StorageTek EBS software to back up the file system by using VSS only.	VSS:*root drive path*=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.

Handling MSDE, SQL Server, and Microsoft Exchange databases

This release of Sun StorageTek EBS software does not support VSS backups of Microsoft Exchange Server databases or SQL Server databases, but *does* support VSS backups of MSDE databases through the MSDE writer. The MSDE writer was designed to handle backups of both MSDE databases and SQL Server databases. If you perform a backup operation that includes the MSDE writer, and both MSDE and SQL Server databases exist on the same system, Sun StorageTek EBS software excludes the SQL Server databases from the backup operation and backs up only the MSDE databases. If only SQL Server

databases exist under the MSDE writer, Sun StorageTek EBS software excludes the entire writer from the backup operation.

If you used a previous release of Sun StorageTek EBS software to perform a VSS backup of Exchange Server databases or SQL Server databases, you can use this release to recover them. Since they were previously backed up by using VSS, they will be recovered by using VSS.

Authoritative restores of the Active Directory Application Mode (ADAM) and DFS Replication writers

Authoritative restores of the ADAM and DFS Replication writers must be performed from the command line. Restores from the NetWorker User GUI will be nonauthoritative.

To perform an authoritative restore of the ADAM or DFS Replication writers, use the -U option for the recover command. To recover the ADAM writer, type the following command:

```
recover -s server -U -N "VSS USER DATA:\ADAM (Address
Book) Writer"
```

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"
```

To restore the DFS Replication writer on Windows 2008 systems, type the following command:

```
recover -s server -U -N "VSS SYSTEM BOOT:\DFS Replication
service writer"
```

Note: You cannot select individual components within the writer for recovery.

25

DiskXtender Data Manager File System Support

This chapter covers these topics:

- Supported configurations 844
- DiskXtender Data Manager file system overview 844
- Backup of DXDM file systems 846
- Recovery of DXDM file systems 849

Supported configurations

These configurations are supported for the backup and recovery of DiskXtender Data Manager (DXDM) file systems:

- DXDM software and Sun StorageTek EBS 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 Sun StorageTek EBS server, client, or storage node software.

Note: 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

Sun StorageTek EBS 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 Sun StorageTek EBS 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 Sun StorageTek EBS 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 46 on page 846, 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 46 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.

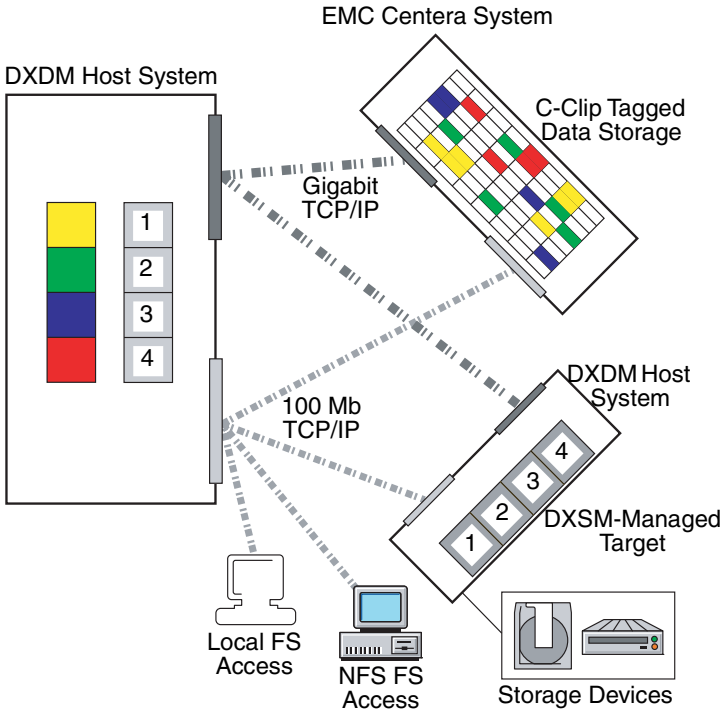


Figure 46 Prototypical DXDM installation

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 117 shows the types of

files and metadata that can exist in a DXDM file system and notes the types that are included in a Sun StorageTek EBS backup.

Table 117 DXDM file system data types included in a Sun StorageTek EBS 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[a]

a. The DMAPI metadata is included by default. It can safely be excluded. "Excluding the DMAPI attributes file" on page 848 provides more information.

File data which has not changed since being migrated is protected by the storage target. Sun StorageTek EBS 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, Sun StorageTek EBS 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 848 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 4, "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 6, "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 849 provides details. "Viewing existing DMAPI processes" on page 848 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 **prtdmsession** command:

```
prtdmsession
```

The **prtdmsession** command lists each existing DMAPI session.

How to remove a defunct DXULDMLIB process

Note: 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 848 provides details.
3. Type the **prtdmsession** 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 **deldmsession** command:

```
deldmsession -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 13, "Recovering 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 851 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 851 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 an earlier version of a file" on page 429 provides information on this procedure.
2. Complete the synchronization process. "File system synchronization" on page 851 provides details.
3. To expedite the synchronization of a particular file, complete the procedure described in "Manually synchronizing a file" on page 852.

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 Sun StorageTek EBS 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.

Note: 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 852 provides more information. "Automatic synchronization" on page 852 describes that synchronization is an automatic process which does not require administrative intervention.

"Manually synchronizing a file" on page 852 describes how make a file's data available more quickly.

If Sun StorageTek EBS **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. Sun StorageTek EBS **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 Sun StorageTek EBS **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 DMAPI processes" on page 848 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.

26

Support for Avamar Deduplication

This chapter covers these topics:

- Sun StorageTek EBS deduplication node 856
- Deduplication backups 856
- Creating a deduplication node and a replication node 860
- Cloning and deduplication 867
- Backup-to-tape for deduplication clients 868
- Backup-to-tape for deduplication clients 868
- Recovering deduplicated data 868
- Viewing deduplication events in NMC 869
- Disaster recovery 870
- Deduplication save sets 870
- Refreshing deduplication node information 872
- Tips for troubleshooting deduplication nodes 873

Sun StorageTek EBS deduplication node

A Sun StorageTek EBS deduplication node is an EMC Avamar® server that stores deduplicated backup data. The initial backup to a deduplication node should be a full backup. During subsequent backups, the Avamar infrastructure identifies redundant data segments at the source and backs up only unique segments, not entire files that contain changes. This reduces the time required to perform backups, reduces the network bandwidth and reduces the storage space used for backups.

Avamar server installation is separate from Sun StorageTek EBS installation, and is performed by EMC Customer Support. The Avamar server must be configured as a Sun StorageTek EBS deduplication node, and the Avamar server must be available when:

- A deduplication client resource is created.
- The Avamar server receives backup data from Sun StorageTek EBS deduplication clients.

Deduplication backups

All typical client operations, such as adding, deleting, and drag-and-drop, work the same for deduplication clients as they do for other types of clients.

A deduplication save set is treated as a regular NetWorker save set, although it actually has two parts: metadata (hash info) and the backed-up client data. Only the metadata is stored on a NetWorker storage node. The backed-up client data is stored on a deduplication node (Avamar server). The metadata must be cloned in the usual manner, while a copy of the backed-up data can be replicated on another deduplication node.



IMPORTANT

For disaster recovery, you must replicate the client data to another deduplication node and you must clone the metadata. Both the metadata and the client data are required to recover backed up client data.

Also consider the following points:

- Bootstrap and index data is not deduplicated.

- Both metadata and client data are required to recover a backed up client.
- Windows SYSTEM save sets do not support deduplication in NetWorker versions prior to 7.5 Service Pack 1.
- 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.
- Do not place both regular and deduplication clients in the same group.
- If deduplication is used for a client, directives cannot be applied.
- The metadata for the deduplication client must be cloned. Ensure that cloning is enabled on the group.
- The parallelism value for the client must not be greater than 4.

Replication of deduplication backups

Replication is configured by EMC Customer Support. If a replication node has been configured for a deduplication client, be aware that replication does not take place immediately after a backup finishes. There is a time lag before replication can happen, and all of the data must be available on the replication node before a recovery operation starts. Replication must be scheduled or started manually. Schedule the deduplication client's replication to start (or manually start replications) at a time that ensures that a backup operation can be replicated.

Schedules for deduplication clients

Deduplication backups must be scheduled to avoid the deduplication node's read-only periods, when such cron jobs as checkpoint and garbage-collection are run. The Avamar server documentation provides instructions.

With deduplication, backups require significantly less time, once the initial full backup has run.

Backup levels for deduplication clients

When a deduplication client is backed up for the first time after deduplication has been enabled, a full backup of the client data must

be performed. Subsequent backups of the client then take advantage of deduplication to back up only those data segments that have changed since the previous backup. The exception to this is that full deduplication backups always include a complete backup of the index.

Note: If the deduplication node for an existing deduplication client is changed or if another type of client is upgraded to a deduplication client, a new full backup of the client data must be performed, regardless whether the most recent backup performed was a full backup.

The initial full backup of a deduplication client takes longer than a regular backup of the same client data due to the initial overhead involved in deduplicating the data. All subsequent deduplication backups benefit from the fact that now only the data segments (not files) that have changed since the previous backup will be backed up. The choice of subsequent backup levels, however, depends on what is more important to the user: backup performance or recovery performance. Deduplication clients should use backup schedules that support that priority.

Initial full backup followed by daily (Level 1) incremental backups

Faster backups, longer recoveries: This schedule benefits an environment in which the speed of the backup is most important. The daily incremental (Level 1) deduplication backups will take significantly less time than daily regular incremental backups and even less time than daily full deduplication backups.

After one week of backups, this schedule results in slightly longer recovery times, due to the latency involved in restoring seven backup images (the initial full backup, followed by applying each incremental backup to the full backup).

Initial full backup followed by daily full backups

Longer backups, faster recoveries: This schedule benefits an environment in which the speed of data recovery is most important. After the initial full backup, daily *full* deduplication backups take only slightly longer than regular daily *incremental* backups.

After one week of backups, this schedule results in much shorter recovery times, since only a single backup image must be recovered.

Table 118 provides a a rough comparison of backup schedules. All estimates are relative and approximate. Recovery time estimates are for saveset recovery operations.

Table 118 Backup schedule comparison

Backup type	Full initial backup	Subsequent daily, full backups	Subsequent daily, incremental backups (with a weekly full)	Data recovery time after one week
Non-deduplication	x hours	x hours	1/10 x hours	1.6 to 2.0x, if the full backup and all incremental backups were restored.
Deduplication, daily full	1.5 x x hours	1/5 x hours	--	Less than or equal to x (only one backup need be restored).
Deduplication, daily incremental	1.5 x x hours	--	1/20 x hours	1.5 to 2.0x, if the full backup and all incremental backups were restored.

Chapter 4, "Backup Groups and Schedules" has more information about backup groups and schedules.

Retention policies

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 policy. 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 Sun StorageTek EBS 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.

Directives

If deduplication is used for a client, directives cannot be applied.

Creating a deduplication node and a replication node

Deduplication nodes and replication nodes must be created on Avamar servers. Once that has been done, you create access to them from the Sun StorageTek EBS side. The tasks in this section are grouped as follows:

- "Avamar setup" on page 860
- "NetWorker setup" on page 862



IMPORTANT

To avoid ambiguity in name resolution, always specify the fully qualified domain name (FQDN) to identify the deduplication node and the replication node, if used.

Avamar setup

This section includes the tasks that must be performed on the Avamar server. Complete the tasks in this section before completing the tasks in "NetWorker setup" on page 862.

Note: The tasks in this section assume that an Avamar 3.7.2-57 or later server is installed and configured at your site according to the instructions found in the *Avamar System Installation Manual* for your platform. The Avamar-specific tasks in this document are up-to-date at the time of writing, however, for the latest instructions, refer to the NetWorker De-Dupe Node Design and Implementation technical note available in Powerlink.

Task 1: Install the NetWorker Linux client on the Avamar server

1. Open a command shell and log onto the Avamar server utility node as the root user.
2. Obtain the NetWorker Linux Client software (release 7.6 or later) and familiarize yourself with the NetWorker *Installation Guide* and the NetWorker *Release Notes*.
3. Change to the directory containing the NetWorker software.
4. Install the NetWorker Linux client software in the default location, for example:

   `rpm -ivh --nodeps` *lgtoclnt-7.x-y.i686.rpm*

5. Start the NetWorker daemons:

   ```
   /etc/init.d/networker start
   ```

6. Verify that the NetWorker daemons have started:

   ```
   ps -ef | grep nsr
   ```

   The **nsrexecd** daemon should be running

Task 2: Modify Avamar Server mcserver.xml Preferences File

Next, change the Avamar server mcserver.xml "allow_duplicate_client_names" setting from false to true.

To modify the Avamar server mcserver.xml preferences file:

1. Open a command shell.
2. Log onto the Avamar server utility node as the user admin.

   Ensure that you are not still logged in as root.

3. Load the admin OpenSSH key:ssh-agent bash

   ```
   ssh-add ~admin/.ssh/admin_key
   ```

   You are prompted to enter a passphrase.

4. Enter the admin user account passphrase.
5. Stop the administrator server (mcs):

   ```
   dpnctl stop mcs
   ```

   Wait for this command to complete.

6. Open the following file in a Unix text editor:

   /usr/local/avamar/var/mc/server_data/prefs/mcserver.xml

7. Change the "allow_duplicate_client_names" setting from false to **true**.
8. Save your changes.
9. Restart the administrator server (mcs):

   ```
   dpnctl start mcs
   ```

   Wait for this command to complete.

Task 3: Add a NetWorker Domain to the Avamar Server

Finally, create a new NetWorker domain on the Avamar server.

To create a new NetWorker domain on the Avamar server:

1. Start the Avamar Administrator Graphical Management Console.
2. Select **Navigation** > **Administration**. The Administration window appears.
3. Click the **Account Management** tab.
4. In the tree, select the **server icon** (root domain).
5. Select **Actions** > **Account Management** > **New Domain**. The New Domain dialog box appears.
6. Enter the following information:

Table 119

Field	Description
New Domain Name	NetWorker
Contact	Optional contact name
Phone	Optional contact telephone number
Email	Optional contact email address
Location	Optional contact location

7. Click **OK**. A confirmation message appears.
8. Click **OK**. The confirmation message closes, the New Domain dialog box closes and the new domain appears in the Account Management tree.

NetWorker setup

Complete the tasks in this section after you have completed the tasks in "Avamar setup" on page 860.

Task 1: Create an AFTD on the NetWorker server or Storage Node

When using Avamar de-duplication with NetWorker, an Advanced File Type Device (AFTD) is required to back up the metadata of the deduplication save sets. The AFTD can either be defined as a storage node for the de-duplication client, or as the device for the pool that includes the de-duplication client.

- "File type device or AFTD on UNIX and Linux servers and remote storage nodes" on page 342 describes how to set up an AFTD on UNIX or Linux.
- "File type device or AFTD on local disks of Microsoft Windows servers and storage nodes" on page 344 describes how to set up an AFTD on Windows.

Considerations for setting up an AFTD

Before creating an AFTD, be aware of the following considerations:

- Use a local file system (local to the NetWorker storage node, client, or server) for the AFTD. The use of CIFS or NFS devices are supported, however these may have performance limitations.
- File type devices are basic disk devices that have similar restrictions to a tape device. Use an AFTD devices to maximize performance and functionality.
- Create AFTDs in a separate and dedicated directory, preferably with their own filesystem or mount point.
- Do NOT use the root directory, boot partition, or system directory for an AFTD because this could cause performance and system stability issues.
- Each directory to be used for an AFTD should have a minimum of 60 GB dedicated space. This should provide space for 6 Windows clients, backing up daily with 30 days of retention. This includes approximately 300MB for Windows system state information and a total of 1 million files of NetWorker metadata.

Task 2: Create a NetWorker deduplication node

To create a Sun StorageTek EBS deduplication node:

1. In the Sun StorageTek EBS server's **Administration interface**, click the **Devices** button.
2. Right-click **Deduplication Nodes** in the navigation tree, and select **New**. The **Create Deduplication Node** window appears, with the **General** tab shown.
3. In the **Identity** area:
   a. Type the fully qualified domain name (FQDN) of the deduplication node (an Avamar server) in the **Name** field.
   b. (Optional) Type a general comment to help identify the node, such as *Dedupe node1*.

4. In the **Access** area, type the remote username and password for the Avamar server that is the deduplication node.

   Note: Any subsequent changes to the remote password on the Avamar server must be updated in the corresponding NetWorker deduplication node resource. The *Avamar System Administration Guide provides* details about changing the password on the Avamar server.

5. If the **Configuration** area shows a value in the **Replication node** field, then a replication node that this deduplication node can use has already been created on the Avamar server. Use the name of the replication node in "Task 4: Optional, create a replication node for a deduplication node if required" on page 865 if you require a replication node to use for backup failover for the deduplication node you are creating.
6. Click **OK**. When the progress windows disappear, the new deduplication node is listed in the **Deduplication Nodes** table.

Task 3: Configure a deduplication client

After creating a deduplication node, configure a Sun StorageTek EBS deduplication client as follows (or start by editing an existing client and start at step 4).

To configure a deduplication client:

1. In the server's **Administration interface**, click the **Configuration** button.
2. Select **Clients** in the navigation tree. The **Clients** table appears. It lists any deduplication clients that have already been created along with any non-deduplication clients.
3. Right-click **Clients** in the navigation tree. or right-click any client in the **Clients** table, and select **New**. The **General** tab appears in the **Create Client** window.
4. On the **General** tab, type the information as for any client, with the following exceptions:
   - The **Name** of the client must be a computer qualified to be a Sun StorageTek EBS deduplication client: it must run an operating system that supports deduplication.
   - In the **Backup** area:
     – Ensure that the deduplication client has been assigned to a group that contains only deduplication clients.

- Do not mix regular and deduplication clients within a group. "Task 2: Set up a group for backup clients" on page 78 provides general information about creating backup groups.
- A recommended backup schedule for deduplication clients depends on whether your priority is faster backups or faster, less complicated recoveries. "Task 1: Set up a schedule for backups" on page 77 provides general information about setting backup schedules. "Backup levels for deduplication clients" on page 857 provides information about backup levels and schedules for deduplication clients.

5. Click the **Apps & Modules** tab:
   a. In the **Deduplication** area, select the **Deduplication backup** attribute to enable this client for deduplication backups.
   b. Select the name of the NetWorker deduplication node to which this client's backup data will be sent (the node created in "Task 2: Create a NetWorker deduplication node" on page 863). This step links this client with its own deduplication node. Do not select the name of a replication node here.



IMPORTANT

If the deduplication node for this client's backup data changes, the next backup done must be a level 0 (full) backup.

6. Click the **Globals (1 of 2)** tab and ensure that the **Parallelism** attribute is less than or equal to 4.
7. Complete the remaining configuration information as for any regular client.
8. Click **OK**.

Task 4: Optional, create a replication node for a deduplication node if required

If the Sun StorageTek EBS deduplication node requires a replication node for automatic failover of data backups, the replication node must first be configured by EMC Customer Support. The Sun StorageTek EBS software does not track which clients are being replicated. To guarantee that backups from the replication node can be recovered with the Sun StorageTek EBS software, create access to

the replication node from the Sun StorageTek EBS side before the backup is performed.

After the replication node has been configured on the Avamar server, give the replication node access to the Sun StorageTek EBS deduplication node:

1. Right-click **Deduplication Nodes** in the navigation tree on the Devices screen.
2. Right-click the newly created (primary) deduplication node and select **Properties**.
3. Select the **General** tab.
4. Make a note of the following:
   - The fully qualified domain name (FQDN) of the deduplication node. This value is displayed in the **Name** attribute.
   - The FQDN of the replication node. This value is in the **Replication node** attribute.
5. Close the primary NetWorker deduplication node resource.
6. To make the replication node usable, create a new deduplication node for the replication node that was noted in step 4 :
   - Follow the steps in "Task 2: Create a NetWorker deduplication node" on page 863.
   - Ensure that you type the FQDN name of the replication host (Avamar server) in the **Name** field.
7. In the Avamar application, open the Avamar **Administration** interface for the primary deduplication node:
   a. Select the **Services Administration** tab.
   b. Right-click on the **Replication Cron job** attribute and select **View Properties**.
   c. Ensure that the **Destination Directory:/REPLICATE/** attribute matches the exact value of the deduplication node that was noted in step 4 . If the value does not match, make the change and select **Apply**.

Note: For supported Avamar server releases prior to 4.1, the **Destination Directory:/REPLICATE/** attribute is known as the **System Name** attribute.

d. Ensure that the **Destination** attribute matches the exact value of the replication node that was noted in step 4 . If the value does not match, make the change and select **Apply**.



IMPORTANT

Changes to replication node configuration are not done often. When such changes are made, they must be made by EMC Customer Support. After such a change has been made, the user must then update the Sun StorageTek EBS deduplication node to propagate changes to the Sun StorageTek EBS software.

Cloning and deduplication

EMC Avamar deduplication technology decreases the amount of time, network bandwidth, and disk capacity required to back up client data.

The cloning of deduplication backups is somewhat different from the cloning of other Sun StorageTek EBS backups. Only the metadata (hash information) is stored on a Sun StorageTek EBS storage node. This metadata is not deduplicated. However, it can be cloned in the usual manner. "Cloning a save set automatically" on page 295 provides information on how to set up cloning operations.



IMPORTANT

For disaster recovery, you must replicate the client data to another deduplication node and you must clone the metadata. Both the metadata and the client data are required to recover backed up client data.

The backed-up data from a deduplication client, however, is stored on a deduplication node and cannot be cloned to a Sun StorageTek EBS storage node. This backed-up data can be replicated on another deduplication node if such a replication host has been configured. The NetWorker software does not initiate replication. A replication host (an Avamar server) must be configured by EMC Customer Support before a deduplication backup can be replicated. "Task 4: Optional, create a replication node for a deduplication node if required" on page 865 provides more information.

It is also possible to output the backup data of deduplication nodes to tape volumes. "Backup-to-tape for deduplication clients" on page 868 provides more information.

Backup-to-tape for deduplication clients

There is an alternate method to using a replication node to ensure that additional copies of deduplication data are available for recovery. That is to create a second instance of the client to be backed up, but do not configure the second instance as a deduplication client. The second client instance is treated as a regular NetWorker client and its data is backed up to tape.

Example

Clients *mars*, *venus*, and *jupiter* have been configured as deduplication clients on the Apps & Modules tab of the Create Client resource. On the General tab of the Create Client resource, they were assigned to a backup group named "Dedupe backups". This group is scheduled for a daily Level full backup.

To get a monthly tape backup of these clients, create another instance of the *mars*, *venus*, and *jupiter* clients, but do not select the Deduplication backup checkbox on the Apps & Modules tab of the Create Client resource. On the General tab of the Create Client resource, assign *mars*, *venus*, and *jupiter* to a backup group named "Tape backups". Schedule this group for a monthly full backup on one day of the month. Skip every other day of the month.

Recovering deduplicated data

The process for recovering data from a deduplication node is basically the same as that for recovering from a storage node. However, there is an underlying difference in where the information is kept. Only the metadata, which is hash information, is stored (or cloned or staged) on a NetWorker storage node. The backed-up data from a deduplication client is stored on a deduplication node that is an Avamar server.

Both the deduplication node and the volume on the recovered side must be online during the recovery of deduplicated data.

Recovery from a replication node

If recovery from the primary deduplication node fails, autorecovery will use a replication node only if the following are true:

- A replication node has been configured on the Avamar server by EMC Customer Support. A replication node name can be visible from Sun StorageTek EBS even if the node has not been configured. However, it is not usable unless all of the requirements in this section have been met.
- The replication node has also been created as a Sun StorageTek EBS deduplication node. "Creating a deduplication node and a replication node" on page 860 provides more information.
- That replication node was specified in the **Replication node** field of the deduplication client resource and existed at the time the backup was run. "Task 4: Optional, create a replication node for a deduplication node if required" on page 865 provides more information.

If a recovery operation fails to establish a session with the primary deduplication node, it then attempts to recover from the replication node.

Ensure that the system name mentioned in the replication cron job properties on the Avamar server matches the system name in the **Replication node** field of the deduplication node's **Properties**. If the system names do not match, recovery from the replication node will fail with login errors. The Avamar documentation provides information about replication cron jobs.

Viewing deduplication events in NMC

Avamar events can be viewed in the Sun StorageTek EBS Management Console by clicking the Events button in the Console window.

The following types of events can be viewed:

- **hfscheck** — validation of the data stored on the Avamar server.
- Disk health/hardware — status of the disks on the Avamar server (limited to "bad" or "failed" disk status).
- Capacity — monitor and report capacity utilization of the deduplication node and notify with an alert when the threshold is reached.
- Replication events that occur on the Avamar server.

To enable these Avamar events for viewing in NMC, the following one-time configuration in Avamar Administrator is required:

Note: You must be logged into the root domain with the role of Administrator in order to perform this procedure.

1. Start Avamar Administrator.
2. Select **Navigation > Administration** or click the **Administration** launcher button. The Administration window appears.
3. Click the **Event Management** tab and then, click the **Unacknowledged Events** tab.
4. Select one or more of the event types listed above.
5. Select **Actions > Event Management > Acknowledge Unacknowledged Events**.

The events are then removed from the Unacknowledged Events list.

Disaster recovery

Only the metadata is stored on a NetWorker storage node. The backed-up client data is stored on a deduplication node (Avamar server). The metadata can be cloned in the usual manner, while a copy of the backed-up data can be replicated on another deduplication node.



IMPORTANT

For disaster recovery, you must replicate the client data to another deduplication node and you must clone the metadata. Both the metadata and the client data are required to recover backed up client data.

The **scanner** program can recover only the metadata (hash data) for deduplication backups. To recover backed-up replicated client data, see "Recovery from a replication node" on page 868.

If the backup-to-tape method described in "Backup-to-tape for deduplication clients" on page 868 is used regularly, it provides a tape backup of the deduplicated data up to a particular point in time.

Deduplication save sets

This section details how to delete and query deduplication save sets.

Note: Windows SYSTEM save sets do not support deduplication in Sun StorageTek EBS versions prior to 7.5 Service Pack 1.

Deleting deduplication save sets

When index entries for deduplication save sets are deleted, or when an index entry for a volume that contains deduplication save sets is deleted, the Sun StorageTek EBS software performs two tasks:

1. First, the Sun StorageTek EBS software updates the Deduplication node resource for the deduplication node that was used to create the save sets, to indicate that the deduplicated data should be deleted from the deduplication node.
2. Once the Deduplication node resource has been updated with this information, the information in the media database related to the deduplication save set is deleted.

As a result of the process, if the resource for the deduplication node that was used to create the backup has been deleted, then the deletion of the deduplication save set will fail.

As a separate activity, the Sun StorageTek EBS server periodically checks the Deduplication node resource to determine if there is any deduplication data to be deleted; if there is, the Sun StorageTek EBS server instructs the Avamar server to delete the data. Because the deletion of index entries for deduplication save sets and the deletion of deduplication data on the deduplication node are separate activities, space on the deduplication node will not be freed immediately when the index entries for the deduplication save set are deleted.

Querying deduplication save sets using mminfo

Use the **mminfo** program to query information about save sets created using deduplication. The output of **mminfo** identifies save sets that were created with deduplication by using the **dedupe** flag.

To limit the output of **mminfo** to only those save sets created using deduplication, use the **-q dedupe** option. For example, to query deduplication save sets for the local host, use the following command:

```
mminfo -S -q dedupe
```

Be aware that this command does not list the following:

- Empty save sets
- Save sets in which nothing was backed up as the result of skip directives

The previous save sets are treated as regular save sets, not as deduplicated save sets.

Refreshing deduplication node information

Deduplication nodes appear in the Deduplication Nodes table. Information in this table comes from two sources:

- Information in the Name and Replication Node columns comes from the Sun StorageTek EBS server.
- Information in the other columns comes from the deduplication node. The Avamar server documentation provides details about these columns.

To update the deduplication node information:

1. Right-click **Deduplication Nodes** in the navigation tree (or any node or multiple nodes in the Deduplication Nodes table).
2. Select **Refresh**. Selecting **Refresh** from the navigation tree causes the following warning message to display:

   ```
   "Refreshing data for the deduplication nodes may take
   considerable time. Would you like to start the
   operation? Yes or No."
   ```



IMPORTANT

A Refresh done from the navigation tree refreshes each node. If many deduplication nodes are in the table, this can take a long time. If you do not need updated information from all of the deduplication nodes, do the following:
1. In the **Deduplication Nodes** table, select only the node or nodes to be updated.
2. Right-click and select **Refresh** from the table.

Tips for troubleshooting deduplication nodes

Tips for troubleshooting deduplication nodes and clients include:

- Have a dedicated pool for deduplication clients. That pool must use an AFTD (defined on the Sun StorageTek EBS server or storage node) to store the metadata of the deduplication save sets. Do not include deduplication clients in the same group as non-deduplication clients.
- In order to take full advantage of deduplication function, once a deduplication node for a given client has been selected for an initial full backup, continue to use that same deduplication node for the client's backups. Changing the node requires that an initial Level 0 (full) backup be run immediately, negating the time-saving benefit of deduplication.

27

Support for HomeBase

This chapter covers these topics:

- Introduction to HomeBase Support (Bare Metal Recovery) 876
- BMR backups 877
- Recovering with BMR 878

Introduction to HomeBase Support (Bare Metal Recovery)

This section provides information on bare metal recovery (BMR).

BMR Recovery

A BMR recovery enables the restoration or migration of a server from one hardware type to another. This type of recovery begins with the HomeBase Server and the profile data that is gathered and includes the file system and application data restore from the Sun StorageTek EBS server through the Console server.

NetWorker BMR HomeBase Server

A Sun StorageTek EBS BMR server is an EMC HomeBase Server that manages and restores operating system and server state data, collected as profiles. The HomeBase Server is the collection point for profiles produced by HomeBase Agents. HomeBase Agents are automatically installed with the Sun StorageTek EBS client software.

The profile data includes:

- hardware configurations
- operating system levels
- system tuning
- network connections
- security
- storage layouts

The HomeBase Agent captures this information on a scheduled basis with a NetWorker save set, and sends the profile results to a HomeBase Server for management and analysis.

Installation Notes

HomeBase Server installation is separate from Sun StorageTek EBS server installation.

The HomeBase Server must be configured and available when:

- A BMR client resource is created.

- BMR backup data (profiles) are generated from Sun StorageTek EBS clients.

The EMC HomeBase product documentation provides details about HomeBase Server installation, profile management, and recovery specifics.

The NetWorker *Installation Guide* for information provides information on how to install and enable the BMR feature.

BMR backups

Sun StorageTek EBS clients can be configured to take advantage of EMC HomeBase BMR technology to ensure that a client can be restored or migrated from one hardware type to another. When a backup is run on a client, a HomeBase profile is automatically generated and saved to a specified HomeBase Server. This process is enabled with a HomeBase Agent license.

Configuring a client for BMR support

To configure a Sun StorageTek EBS client for BMR support, perform the following tasks:

Task 1: Enable BMR server support

Connection with a HomeBase Server is enabled during the set up of the Sun StorageTek EBS server. This connection enables the delivery of profile data from the Sun StorageTek EBS client to the HomeBase Server.

To enable BMR support:

1. From the **Administration** window, click **Configuration**.
2. Select the Sun StorageTek EBS server name.
3. From the **File** menu, select **Properties**.
4. In the **Properties** dialog box, click the **Configuration** tab.
5. Type the IP address or hostname for the HomeBase Server in the BMR server field.
6. Click **OK**.

Note: A HomeBase Agent license batch code is required to enable the client to generate profiles and send them to the HomeBase Server. Contact your EMC HomeBase representative for licensing information.

Task 2: Configure a client for BMR backup

After enabling communication with a HomeBase Server, configure a Sun StorageTek EBS client as follows (or start by editing an existing client).

To configure a BMR backup for a client:

1. From the **Administration** window, click **Configuration**.
2. Click **Clients**.
3. Right-click the client to be enabled, then select **Properties**. The **Properties** dialog box appears.
4. Click the **Globals (2 of 2)** tab.
5. Click the **BMR** check box.
6. Type the following in the BMR options field:

   -L *license_batch_code*

   Additional options can be specified in this field that define how the profile is generated. The EMC HomeBase product documentation for HomeBase Agent command line options provides more information.

   Note: If there is a change to the HomeBase Server or license batch, the new license must be updated.

7. Click **OK**.

The HomeBase profile, used during a BMR recovery, is generated when the client's scheduled save set is run. The profile is then stored in the HomeBase Server.

Recovering with BMR

This section provides an outline of the steps required to perform a BMR recovery or server migration. The EMC HomeBase product documentation provides complete details and instructions pertaining to the HomeBase Server and HomeBase Agents.

A BMR recovery is performed from the HomeBase Server and Console in the following order:

1. On the new machine, deploy an operating system that matches the operating system of the profiled client that is being restored.

   Windows and Linux operating systems can be deployed from the HomeBase Server web pages. UNIX operating systems are deployed manually with package matching and validation performed through the HomeBase Server web pages.

2. Install a Sun StorageTek EBS client on the machine that will serve as the recovery server. The machine on which the client data is recovered does not need to match the hardware of the original client.

   Note: The Console server and the HomeBase Server must be able to communicate with the recovery server.

3. The HomeBase Server will communicate with the client through with a HomeBase Agent license batch. Communication can be established from the client command line using the **hba licensing announce** command.The *EMC HomeBase Agent Installation and Administration Guide* provides details, or enter **hba -h** for command usage.

4. From the HomeBase Server web pages, restore the base profile onto the new machine. The EMC HomeBase Agent Recovery Guide provides platform specific information.

   Note: The base profile will configure the new client to match the IP and network configuration of the client that was profiled. To avoid IP conflicts on the network, make sure that the original client is not available.

5. Reboot the client.

6. From the client, recover the client data. "Choosing a recovery method" on page 424 provides more information.

7. Ensure that the **Options> Exclusion Options> Exclusion File List> Perform special handling for Bare Metal Recovery** checkbox is selected. This automatically excludes files that the HomeBase Server has recovered with the base profile restore.

Note: The exclude.NETWORKER file is located in the HomeBase Agent exclude subdirectory and can be edited to exclude additional data. However, this should be done with extreme care.

8. From the HomeBase Server web pages or the client command line, restore the extended profile. The *EMC HomeBase Agent Installation and Administration Guide* for details provides more information, or enter **hba -h** for command usage.

   The extended profile must exist on the client. It can be downloaded from the HomeBase Server web pages or restored through the data recovery in step 6 .

9. Reboot the recovered client.

Excluding files for BMR recovery

A HomeBase profile restore is one of the first steps in a BMR recovery. A Sun StorageTek EBS client data restore may restore files that were previously restored by HomeBase. An exclusion file is provided to ensure that there is no overlap between a HomeBase profile restore and the Sun StorageTek EBS client data restore. The EMC HomeBase Recovery Guides provide complete information about excluding files from recovery.

Before performing a data restore in a BMR recovery, enable the use of the default exclude file.

To enable the exclude file:

1. In the NetWorker **User** program, select **Exclusion Options** from the **Options** menu.
2. In the **Exclusion File List** dialog box, select the **Perform special handling for Bare Metal Recovery** check box.
3. To add a custom exclude file, enter the full path in the field provided and click **Add**.



CAUTION

Custom exclude lists should be used with caution and tested before recovery to a production system is performed.

4. Click **OK**.

28 Licensing

This chapter covers these topics:

- Product licenses 882
- Evaluating the NetWorker software 882
- Permanently licensing the NetWorker software 884
- NetWorker License Manager 884
- Licensing Sun StorageTek EBS software components 886
- License Conformance Summary 893
- How to obtain NetWorker license information 895
- Querying the local server 896

Product licenses

NetWorker software and added features, such as modules, can be downloaded and evaluated for free from Powerlink, the EMC website, or the media kit. The software can be evaluated for 30 days without an enabler code or license. After that, you can obtain evaluation enablers to extend the evaluation period for an additional 45 days.

To permanently use the NetWorker software, you must license the software on the Powerlink Licensing site and apply the supplied licensing key on the NetWorker server. The license key includes permanent enabler codes and corresponding authorization (auth) codes.

Base enabler

Each installation of NetWorker server software must be licensed with a base enabler. Only the server's enabler is called a *base* enabler. This enabler "turns on" the software and allows the use of a particular bundle of features, such as a specified number of clients and devices. All license keys are entered and stored on the NetWorker server, which enforces the licensing.

Base enablers come in different editions that enable varying degrees of functionality. Add-on enablers allow a broader scope of features.



IMPORTANT

You cannot delete the base enabler code that enables the basic Sun StorageTek EBS software. The base enabler can only be upgraded or downgraded.

Evaluating the NetWorker software

When you first install the NetWorker software, you can evaluate the software with all the modules and features for 45 days free without entering any enabler codes or licenses.

The licensing chapter in *The NetWorker Installation Guide* provides detailed instructions on how to evaluate the software.

How to apply an evaluation enabler code

To apply an evaluation enabler on the **NetWorker** server:

1. Start the **NetWorker Management Console** software.
2. Open the **Administration** window:
   a. In the **Console** window, click **Enterprise**.
   b. In the left pane, select a NetWorker server in the **Enterprise** list.
   c. In the right pane, select the application.
   d. From the **Enterprise** menu, click **Launch Application**. The **Administration** window is launched as a separate application.
3. From the **Administration** window, click **Configuration**.
4. In the left pane, select **Registration**.
5. From the **File** menu, select **New**.
6. In the **Enabler Code** attribute, type the enabler code.
7. (Optional) In the **Comment** attribute, type a description of the license.
8. Click **OK**.

Deleting an enabler code

You can delete a license enabler code at any time, whether or not it has been permanently authorized. The license is *not* removed when the software is uninstalled.

Note: You cannot delete the enabler code that enables the basic Sun StorageTek EBS software, called the "base enabler code". A base enabler can only be upgraded or downgraded.

To delete an enabler code:

1. From the **Administration** window, click **Configuration**.
2. Click **Registrations**.
3. Right-click the license to delete, then select **Delete**.
4. Click **Yes** to confirm the deletion.

Evaluation enabler alert

An alert message is generated 30 days before a NetWorker evaluation enabler code expires. The alert remains until the NetWorker evaluation enabler is authorized or deleted.

To view the license alert one of two ways:

- From the **Sun StorageTek EBS Administration** window, click **Monitoring** and select the **Alert** tab.
- From the **Console** window, click **Events**.

Color-coded icons appear with the alert message:

- A yellow icon indicates that the enabler will expire within 30 days. The message appears daily until 10 days prior to the evaluation enabler code expiration.
- A red icon indicates that the enabler will expire within 10 days. The message appears daily up to and including the day of the enabler code expiration.

By the end of the evaluation period, you must permanently license the software to continue using modules or features that you have evaluated. The Licensing Process Guide provides more information.

Permanently licensing the NetWorker software

To permanently use the NetWorker software to back up and recover data, you must license the software. This licensing process is the same for all editions of the NetWorker software, and for individual modules and features.

The licensing chapter in the *NetWorker Installation Guide* provides detailed instructions on how to permanently license the software.

NetWorker License Manager

The NetWorker License Manager software provides centralized license management, which enables you to maintain all of an enterprise's NetWorker licenses from a single computer.

With the NetWorker License Manager, you can move NetWorker software from one computer to another, or change the IP address on an existing NetWorker server without having to reauthorize the

software. The NetWorker License Manager can be installed 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 enter an enabler code if you are using LLM:

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, enter 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 Sun StorageTek EBS 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, enter the hostname of the appropriate server and click **OK**.

Licensing Sun StorageTek EBS software components

This section provides licensing information on the following Sun StorageTek EBS features:

- "Archiving licensing" on page 886
- "Bare Metal Recovery (BMR) licensing" on page 887
- "Client connection licenses" on page 887
- "Cluster licensing" on page 887
- "DDS licensing" on page 888
- "HomeBase server/agent" on page 889
- "NDMP licensing" on page 889
- "NetWorker Modules" on page 889
- "Sun StorageTek EBS cloud backup option" on page 890
- "Sun StorageTek EBS storage node licensing" on page 890
- "Virtual environments simplified licensing" on page 890
- "Virtual Edition Client Connection license" on page 891

Archiving licensing

The archive feature must be purchased and licensed separately from other Sun StorageTek EBS software components.

Before licensing and using the Sun StorageTek EBS archive feature, ensure that you have a device, either stand-alone device or in an autochanger or silo, connected to a Sun StorageTek EBS server or storage node. If you are cloning archives, you must have at least two devices available.

Bare Metal Recovery (BMR) licensing

Connection with a HomeBase server is enabled during the setup of the Sun StorageTek EBS server. This connection enables the delivery of profile data from the Sun StorageTek EBS client to the HomeBase server.

Note: A HomeBase agent license batch code is required to enable the client to generate profiles and send them to the HomeBase server. Contact your EMC HomeBase representative for licensing information.

Client connection licenses

Every computer to be backed up in a NetWorker datazone requires a client connection license, even the NetWorker server. The client connection license may be one of the licenses that is supplied with the base enabler or purchased separately. An NDMP data server requires a special type of client connection license.

Note: EMC ClientPak® enablers are no longer required. Client licensing is now based solely on the client connection enablers.

Cluster licensing

This section discusses Sun StorageTek EBS cluster licensing.

Note: Cluster client connection enablers are no longer required. Cluster client licensing is now based solely on the client connection enablers. However, existing cluster client connection licenses will still be honored. Each physical node of the cluster requires a client license. Virtual clients of a cluster do not require their own license.

Enabling the highly available Sun StorageTek EBS server

A client connection license is required for each physical node in the cluster on which you intend to run the highly available Sun StorageTek EBS server.

In addition, a storage node that is configured locally, that is, a storage node that uses a hostname that matches the physical hostname of the node that is running the Sun StorageTek EBS virtual server, does *not* require a separate storage node enabler.

For example, on a node in a cluster with the physical hostname Node A, a remote device defined as rd=NodeA:/dev/tape1 would not require a separate storage node enabler while the Sun StorageTek EBS virtual server was running on Node A. "Configuration option 3" on page 697 provides an example of this configuration.

Enabling cluster and virtual clients

Sun StorageTek EBS client licensing differentiates between stand-alone computers and computers that participate in a cluster. The licenses are bound to physical nodes. Therefore, once a client connection license is allocated, any virtual clients that are running on that physical node can be backed up.

Backing up application modules on cluster clients

To back up a NetWorker Module:

- From the *virtual* client, an application module license is required for the virtual client.
- From the *physical* host, an application module license is required for the physical client computer.

For example, physical nodes A and B are clustered and run a database. Virtual node C runs on this cluster and provides access to the database. If node C is the only way the database is backed up, one database application module license is required. If the physical node also backs up the database, a second database application module is required. Of course, client connection licenses are required for both of the physical nodes.

DDS licensing

To enable DDS, one license is required for each drive that is to be shared. Once a drive is licensed as shared, any number of storage

nodes can share it. DDS licensing is independent of library and storage node licensing.

HomeBase server/agent

The HomeBase server/agent have a separate licensing scheme and require separate enablers.

More information on HomeBase licensing is provided in the following documentation:

- EMC HomeBase Server Installation and Configuration Guide
- EMC HomeBase Agent Installation and Configuration Guide

NDMP licensing

The Sun StorageTek EBS NDMP interface is an optional feature that must be licensed separately from the Sun StorageTek EBS base product according to a tiered licensing structure.

The following apply to NDMP licensing:

- One NDMP Client Connection license is required for each NDMP host.
- For NAS arrays, where each host in the array shares one license, you can specify the user-defined array name in the **NDMP array name** attribute of the Client resource.
- When the NetWorker NDMP interface is not used, a storage node license is required to use the Sun StorageTek EBS storage node software with NDMP.
- Additionally, a file type device is supported for NDMP operations, which can also be used as a bootstrap device. A DiskBackup license is required for the file type device.

NetWorker Modules

NetWorker Modules are licensed on the basis of one enabler per database type host.

For example, to back up the Oracle database on two hosts, two NetWorker Module for Oracle enablers are required, even if the two hosts are backed up by the same server.

However, if multiple Oracle instances are running on a single NetWorker client, only one NetWorker Module for Oracle enabler is required.

Sun StorageTek EBS cloud backup option

A NetWorker cloud backup option license enables you to add an unlimited number of cloud backup devices.

Sun StorageTek EBS storage node licensing

The following apply to storage node licensing:

- Each storage node requires a storage node license, in addition to its client connection license.
- A storage node which allows backup only of local data is licensed separately as a dedicated storage node.
- Devices such as files or jukeboxes require device enablers, regardless of whether they are on the server, storage node, or dedicated storage node.

Note: If the Sun StorageTek EBS storage node software is disabled (if the Sun StorageTek EBS server is disabled or the storage node enabler is expired), you will not be able to recover backed-up data by using the storage node. To recover that data, the remote volume must be moved to a local drive connected to the server, and the recovery performed from there.

Virtual environments simplified licensing

Sun StorageTek EBS introduces a simplified licensing model for virtualized environments. The *EMC Software Compatibility Guide* provides a detailed list of supported server virtualization environments.

Two new attributes have been added to the General tab of the Client resource to identify the client as a virtual client:

- **Virtual client**. Set the attribute to **Yes** by selecting the Virtual Client attribute checkbox if the client is a virtual client.
- **Physical host**. If the client is a virtual client, set the attribute to the hostname of the primary/initial physical machine by hosting the virtual client.

Virtual Edition Client Connection license

A new license type, Virtual Edition Client Connection, works with all server virtualization environments supported by the Sun StorageTek EBS software.

The Physical host identified in the **Virtual client hostname** attribute will consume one Virtual Edition Client Connection license, regardless of how many virtual clients are running on that host.

- If a virtual machine is licensed on one physical host and then migrates to another physical host, the new physical host requires its own Virtual Client Connection license.
- The hostname does not need to be fully-qualified, and must be less than 64 bytes.
- All clients sharing the same physical host must use an identical name. Do not mix name formats such as short, FQDN or IP address.

To free a Virtual Client Connection license that has been assigned to a physical host, all references to the physical host must be removed, by doing one of the following:

- Changing the virtual client and/or physical hostname attribute in the Client resources for all virtual clients that reference the physical host.
- Deleting all Client resources for virtual clients that reference the physical host.

Sun StorageTek EBS Modules in virtual environments

When Sun StorageTek EBS Module software is running natively on one or more virtual machines, one module license per module type will be consumed per physical host, regardless of the number of virtual clients associated with that physical host. The physical host itself will also require a Virtual Edition Client Connection license. Every physical machine that might host virtual clients with modules should be licensed in this way.

One license is required for each application type (Microsoft SQL, Exchange, and SharePoint, Oracle, and SAP) used within all of the virtual machines on a single physical server. There are no changes to model codes for NetWorker Modules, so use existing codes and license enablers.

Virtual Tape Library licensing

A Virtual Tape Library Frame license must be purchased for each physical hardware frame that supports VTLs. If the Virtual Jukebox attribute is set to Yes during configuration, but a Virtual Tape Library license does not exist, the configuration will succeed but subsequent library operations will fail.

Applying the Virtual Client license to an existing virtual machine (VM) after upgrading from a previous release

The VCC license is not automatically applied to an existing VM after upgrading to NetWorker 7.6. As a result, the VM uses one standard client license for the preexisting virtual client instead of using the virtual client license.

To take advantage of the Virtual Edition Client Connection licensing policies and free up the use of the standard client license:

1. Select the **Virtual Client** attribute checkbox for this Client resource.
2. Specify the physical host.

Enabling virtual client licensing for existing Client resources

After upgrading from a release prior to Sun StorageTek EBS release 7.4 Service Pack 3, to enable the use of the virtual client license:

1. Edit the existing client by right-clicking the client in the in the **Configuration** screen of the **Administration** window and selecting **Properties**.
2. Select the **Virtual client** option.
3. For the **Physical host** attribute, enter the hostname of the physical machine hosting the virtual client.

Note: The hostname does not need to be fully-qualified, and must be less than 64 bytes.

Licensing for VMware Consolidated Backup

When VMware Consolidate Backup (VCB) is being used, each proxy backup host will consume one standard Client Connection license, regardless of how many virtual clients are configured to backup via the proxy backup host.

License Conformance Summary

In the course of using Sun StorageTek EBS software to protect enterprise data, users add clients and modules as needed. It can be helpful to see a summary report of the current license information for any given Sun StorageTek EBS server. A summary report enables users to determine several things.

For example:

- Which or how many products or features have been enabled
- Which or how many licenses have been purchased and authorized
- Which or how many of those licenses have been used and, conversely, how many are still available to use
- How many additional licenses might be needed
- Whether the number of licenses conforms to the number of licenses in use

From Console, a License Conformance Summary report that shows relevant information about the Sun StorageTek EBS related product licenses purchased or under evaluation can be run. The information is presented in a tabular form that enables you to see installed products, licenses that are in conformance, and which additional licenses, if any, must be obtained in order to close conformance gaps.

In addition to product and license information, the summary displays certain details about the environment and the software version that is run by the Sun StorageTek EBS server.

Accessing the License Conformance Summary

To display the License Conformance Summary:

1. In the main **Console** window, click the **Enterprise** button on the taskbar.
2. Highlight a host in the navigation tree; in the right pane, right-click Sun StorageTek EBS, and select **Launch Application**. The **Sun StorageTek EBS Administration** window appears.
3. Click the **Configuration** button on the taskbar.
4. Select **Configuration** > **License Conformance Summary** on the toolbar. The **License Conformance Summary** appears.

5. When finished, click **OK** or **Cancel** to exit the summary.

Note: The License Conformance Summary can also be displayed by right-clicking **Registrations** in the server's navigation tree, and selecting **License Conformance Summary**.

License Conformance Summary details

The License Conformance Summary provides the following information:

- Server and environment information:
  - **Sun StorageTek EBS Server**: Name of the Sun StorageTek EBS server.
  - **Version**: Sun StorageTek EBS software release and build number.
  - **Full Conformance**: Possible values are **Yes**, **No**, or, if in evaluation mode, **Eval**.
  - **Base License**: Possible values include:
    - The Sun StorageTek EBS edition that is installed
    - **Eval** indicates evaluation mode.
    - **Yes** or **No indicates** whether an authorization code has been applied to the license.
    - **Disabled** indications that the server has been disabled.
  - **Operating System**: Operating system installed on the Sun StorageTek EBS server.
  - **Report Date**: Date and time when the summary was created.
- License related information is displayed in the following columns:
  - **License**: Type of license, for example, Storage Node, Client, module name.
  - **Number Used**: How many licenses of this type have been used.
  - **Number of Licenses**: How many licenses of this type are installed and not expired on the server. While the license is in evaluation mode, the number displayed is the maximum number possible for the license type.

- **% Conformance**: Degree to which the number of licenses purchased is less than or equal to the number used. For DiskBackup Option and Virtual Tape Library, this might be blank, because more than one device can be created for each license of these license types.
- **Notes**: Additional information, specific to the license type, provided by the system (for example, the capacity of a DiskBackup license).

◆ A checksum **(five groups of generated alphanumeric characters) is listed at the bottom of the summary if the summary contents are valid. Users may ignore this line**.

When viewing the License Conformance Summary, consider the following:

◆ DiskBackup Option and Virtual Tape Library (VTL) Frames are the only license types for which the **Number Used** can exceed the **Number of Licenses,** because more than one device can be created for each license of this type.

◆ The License Conformance Summary is not available for use with the Sun StorageTek EBS License Manager software.

◆ The License Conformance Summary supports the standard Console table functions, such as **Print** and **Export**.

◆ A time-stamped snapshot License Conformance Summary is sent to the */nsr/logs* directory at the start of each quarter.

◆ The Sun StorageTek EBS server updates license information only once daily, at noon. Changes made after noon will be reflected in the next day's update.

◆ **License Conformance** is a new attribute for the server nsr resource.

◆ **License Conformance Checksum** is a new attribute used by the system to maintain the integrity of the summary.

How to obtain NetWorker license information

The **nsrlic** command is installed as part of the server installation and is not available on machines that only have the client software installed. There are a number of ways that a customer can obtain license information from a server or servers.

Querying the local server

To query the local server, that is, the server where **nsrlic** is stored, type **nsrlic** at the command line.

Example UNIX**: /usr/sbin/nsrlic**

Windows**: networker_install_dir\bin\nsrlic**

A report is produced with various quantities and servers indicated. The following is an example report output.

Example 52 Report

```
STANDARD CLIENT LICENSES
Available: 12
Used: 5
Loaned to Virtual: 0
Remaining: 7
Connected Clients

VIRTUAL CLIENT LICENSES
Available: 0
Borrowed from Server: 0
Used: 0
Remaining: 0
Connected Clients

NDMP CLIENT LICENSES
Available: 0
Used: 0
Remaining: 0
Connected Clients

SERVER/CLUSTER CLIENT TYPES
AIX: 0
Digital UNIX: 0
HP UX: 0
HP MPE: 0
Linux: 2
NetWare: 0
Network Appliance: 0
SGI: 0
Solaris: 2
SunOS: 0
UnixWare: 0
Windows NT Server: 1
```

```
   WORKSTATION CLIENT TYPES
   DOS: 0
   Macintosh: 0
   OS/2: 0
   Windows 3.1x: 0
   Windows 95: 0
   Windows NT Workstation: 0
   UX/4800: 0
   Others: 0

   Defined Clients          PRE-5.0 CLIENT TYPES
   APPLICATION LICENSES

NetWorker Module for Informix, Unix Client/1
   Available: 1
   Used: 0
   Remaining: 1

NetWorker Module for Databases and Applications,
Windows/Linux Client/1
   Available: 1
   Used: 1
   Remaining: 0
```

Determining the number of available client licenses

To determine the number of available client licenses, look at the "nsrlic: Remaining" line.

Example In the previous listing, where it indicates "sv=12", this means this particular server has 12 licenses available. Also, you can look below to the STANDARD CLIENT LICENSES section and locate: "Remaining: 12." This is another indication of the number of available servers.

Querying a server

To query a specific server, type the **-s** <*server*> option at the command line.

Example `/usr/sbin/nsrlic -s bacoor`

where *bacoor* is the server name being queried.

Querying a server for all information

To query a specific server for all information, type the **-v** <*server*> option at the command line.

Example `/usr/sbin/nsrlic -v bacoor`

where *-v* is for a verbose query for bacoor.

Note: This query may be helpful in troubleshooting license issues.

A verbose report is produced with various quantities indicated.

Example 53 Verbose report

```
STANDARD CLIENT LICENSES
Available: 12
Used: 5
Loaned to Virtual: 0
Remaining: 7
Connected Clients: bu-cosmos, bu-orantx64ja
bu-star1.lss.emc.com,
bu-sunscreen, kip;

VIRTUAL CLIENT LICENSES
Available: 0
Borrowed from Server: 0
Used: 0
Remaining: 0
Connected Clients

NDMP CLIENT LICENSES
Available: 0
Used: 0
Remaining: 0
Connected Clients

SERVER/CLUSTER CLIENT TYPES
AIX: 0
Digital UNIX: 0
HP UX: 0
HP MPE: 0
Linux: 2
NetWare: 0
Network Appliance: 0
SGI: 0
Solaris: 2
SunOS: 0
UnixWare: 0
Windows NT Server: 1
```

```
   WORKSTATION CLIENT TYPES
   DOS: 0
   Macintosh: 0
   OS/2: 0
   Windows 3.1x: 0
   Windows 95: 0
   Windows NT Workstation: 0
   UX/4800: 0
   Others: 0

   Defined Clients     PRE-5.0 CLIENT TYPES
   APPLICATION LICENSES

NetWorker Module for Informix, Unix Client/1
   Available: 1
   Used: 0
   Remaining: 1
   Connected Clients

NetWorker Module for Databases and Applications,
Windows/Linux Client/1
   Available: 1
   Used: 1
   Remaining: 0
   Connected Clients: bu-orantx64ja;
```

29

Troubleshooting

This chapter covers these topics:

- Before contacting technical support 902
- Viewing log files 904
- Sun StorageTek EBS functionality issues 909
- Devices and autochangers 925
- Sun StorageTek EBS locale and code set support 936
- Resource database notes 936
- Network and server communication errors 938
- UNIX communication issues 941
- Sun StorageTek EBS archiving and retrieval 950
- Storage nodes 952
- Console error messages and corrective actions 954
- Console log files 957
- Console troubleshooting notes and tips 959

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 Powerlink website for technical assistance. Provide this information:

- The software version of the Sun StorageTek EBS 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, IRIX, 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 SUNWebsn
    ```
  - For AIX, enter the **lslpp -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 Sun StorageTek EBS command was successful before you made any changes and, if so, the changes you made.

Determining the version of Sun StorageTek EBS software running on a client

To determine the version of the Sun StorageTek EBS software running on a client, use either the Sun StorageTek EBS client interface or the **nsradmin** command.

Determining the software version by using the client interface

To determine the software version by using the client interface:

1. Open the Sun StorageTek EBS client program:
   - On Windows, open the NetWorker **User** program.
   - On UNIX, open the **nwrecover** program.

   Chapter 1, "Overview" provides more information about opening these programs.
2. From the **Help** menu, select **About Sun StorageTek EBS User**. The Sun StorageTek EBS version number appears in the **About** dialog box.
3. Click **OK** to close the dialog box.

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 Sun StorageTek EBS software running on each client is displayed.

Displaying diagnostic mode attributes

Sun StorageTek EBS 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

For Sun StorageTek EBS release 7.4 and later, the log files viewed by using the **nsr_render_log** (UNIX/Linux) or **nsr_render_log.exe** (Microsoft Windows) program include:

- Daemon log file (daemon.raw)
- gstd log file (gstd.raw)
- Sun StorageTek EBS User log file (networkr.raw; Microsoft Windows only)

The **nsr_render_log** program renders internationalized Sun StorageTek EBS 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 Sun StorageTek EBS 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.

Rendering log files in the current locale at runtime

You can also instruct the Sun StorageTek EBS 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 Sun StorageTek EBS 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 Sun StorageTek EBS 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 Sun StorageTek EBS software to render log files into the current locale at runtime:

1. Log in as root or as Windows administrator on the Sun StorageTek EBS client.

   Note: You must have security administrator privileges to view security audit 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 Sun StorageTek EBS log file direction on Microsoft Windows, type the following:

   ```
   update runtime rendered log: "C:\\Program
   Files\\Legato\\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\Legato\nsr\logs\daemon.raw"
```

Note: The **nsr_render_log** program is located in the bin directory of the Sun StorageTek EBS 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 "C:\Program
Files\Legato\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, "C:\Program
Files\Legato\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 Sun StorageTek EBS 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 Sun StorageTek EBS

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 "C:\Program
Files\Legato\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 54 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.

```
# jobquery
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.

Sun StorageTek EBS functionality issues

This section describes workarounds for Sun StorageTek EBS issues.

Backup and recovery

This section covers backup and recovery operations.

Checking the Sun StorageTek EBS services

If you have trouble starting Sun StorageTek EBS 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/nsrmmd -n 1
```

```
12234 ?        S  0:00 /usr/sbin/nsr/nsrmmd -n 2
12410 pts/8    S  0:00 grep nsr
```

If daemons are not present, start the Sun StorageTek EBS daemons.

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.

"Unable to connect to host" error in the Client Backup Configuration wizard

In the **Specify the Backup Configuration Type** page, when you click **Next**, the following error is displayed:

```
Unable to connect to host: Please check Security setting
and daemon logs on the Networker client and Console
server for more details
```

Check for one of the following when you receive this error:

1. Verify whether the SSL key in the NSR Peer Information under nsrladb matches between the NMC Server and Client. This occurs when nsrladb is deleted during installation.

   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: Once 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.

Unsynchronized client clock errors

This situation affects computers with Sun StorageTek EBS client software release 5.0 or later. If the setting for the system clock on a Sun StorageTek EBS client differs from that of the Sun StorageTek EBS server by more than one minute, the following message appears during a scheduled backup:

```
Warning: unsynchronized client clock detected
```

If the gap in synchronization is less than two minutes, this message appears, but the schedule is still honored. To avoid receiving this message, keep the server and client clocks synchronized.

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.

Clone ID timestamp does not reflect the time the cone was created

To guarantee that cloned save sets created on different storage nodes do not have the same timestamp, the Sun StorageTek EBS 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 549 provides more information.
- Shut down and restart the Sun StorageTek EBS services. "Stopping and starting a Sun StorageTek EBS server, client, or storage node" on page 70 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 by save set selection" on page 436 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, **nsrmmd** 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 Sun StorageTek EBS 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" volid
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 Sun StorageTek EBS server.

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 Sun StorageTek EBS 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 Sun StorageTek EBS binaries.

Illegal characters in configurations

When naming label templates, directives, groups, policies, and schedules, these characters are not allowed:

```
/ \ * [ ] ( ) $ ! ^ ' " ? ; ` ~ < > & | { }
```

Hostname aliases

Savegroups fail when aliases are improperly defined. Under certain conditions, such as improperly configured DNS servers or hosts files, the Sun StorageTek EBS 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 Sun StorageTek EBS 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.

Note: 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.

Failed recover operation directly after new installation

If you attempt to start the **nwrecover** program immediately after installing Sun StorageTek EBS software for the first time, this error message appears:

```
nwrecover: Program not found.
```

To avoid the problem, perform a Sun StorageTek EBS backup on the client.

Recovering files from an interrupted backup

If you terminate a backup by stopping the Sun StorageTek EBS services, you cannot recover the files because the media database is not updated when the services stop. Consequently, the Sun StorageTek EBS server does not know on which volumes the requested files reside.

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 Sun StorageTek EBS 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 88 provides more information.

Renamed clients cannot recover old backups

The Sun StorageTek EBS 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. "Directing recoveries to another client" on page 442 provides information about performing directed recoveries.

Client file index errors

These issues are related to client file indexes:

- "Missing client file indexes" on page 918
- "Check failure of client file indexes" on page 918
- "No notification of client file index size growth" on page 919

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 Sun StorageTek EBS 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, recover the index using the procedure outlined in "Recovering a client file index from a specific time and date" on page 453.

No notification of client file index size growth

The Sun StorageTek EBS 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 636 provides information on how to manage the Sun StorageTek EBS 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 recovered are listed.
- Messages similar to this might appear:

```
Recover: ***Canceled***
Recover: Unable to read checksum from save stream
Recover: error recovering /APPSTART.ANI
Didn't recover requested file /APPSTART.ANI
```

The messages indicate that a recovery was not stopped cleanly.

RPC error

If Sun StorageTek EBS has trouble backing up a directory path, a message similar to this appears, which notes the path:

```
* jupiter:/ save: xdr of win32 attributes failed for
'/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 and Windows 2000 have 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.

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 Sun StorageTek EBS 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 85 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 Sun StorageTek EBS 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 Sun StorageTek EBS server.

The information to recover server indexes is stored in the media database on the Sun StorageTek EBS server.

Note: The **savegrp** no longer performs a server index backup for servers not in an active group.

Copy violation

If Sun StorageTek EBS software is installed on multiple servers and the same Sun StorageTek EBS 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 Sun StorageTek EBS software from the extra servers.
3. Restart the Sun StorageTek EBS services on the server where the backups are to go.

Xview errors

If this error message appears when you attempt to start Sun StorageTek EBS client graphical interfaces (for example, **nwrecover**), it means that the client is not authorized to display Sun StorageTek EBS graphical interfaces:

```
Xlib: connection to "mars:0.0" refused by server
Xlib: Client is not authorized to connect to Server
Xview error: Cannot open display on window server:
mars:0.0 (Server package)
```

To correct the situation, configure the client to display the graphical interfaces:

1. From the client machine, invoke the **xhost** command:

   xhost *server_name*

2. Remotely log in to the Sun StorageTek EBS server and issue the **setenv** command:

```
setenv DISPLAY client_name:0.0
```

For command shells other than **csh** type:

```
DISPLAY=client_name:0.0
export DISPLAY
```

Converting sparse files to fully allocated files

The Sun StorageTek EBS 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 workaround 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.

Note: To perform this workaround, 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 172 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 Sun StorageTek EBS 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 a host has multiple network interface cards (NICs), all NICs must be configured and must be have resolvable names, even if one or more NICs are not being used. Failure to have all NICs be resolvable may cause problems with host connectivity to the Sun StorageTek EBS server.

Libraries entering ready state

When starting Sun StorageTek EBS or after configuring a library, it may take a short amount of time for the library to enter the Ready state within Sun StorageTek EBS. 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 Sun StorageTek EBS 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 Sun StorageTek EBS** check box is selected in the Sun StorageTek EBS 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 Sun StorageTek EBS Console, select **Setup**.
2. Select **Setup>System Options**.
3. Unselect the **RPC ping via UDP when connecting to Sun StorageTek EBS** checkbox.

Devices and autochangers

This section explains how to resolve problems with devices and autochangers.

Note: 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 903 provides information about displaying hidden attributes.

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

Note: Do not change time related attributes unless advised to do so by a Technical Support representative.

Maintenance commands

Sun StorageTek EBS 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 Sun StorageTek EBS 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 Sun StorageTek EBS server.
3. Reset the autochanger by typing this command:

   ```
   nsrjb -Hv
   ```

4. Perform an inventory by typing this command:

   ```
   nsrjb -Iv
   ```

The Sun StorageTek EBS server can use the autochanger after the inventory operation completes.

Chapter 3, "Configuring Libraries and Devices" provides complete information on the use of 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 Sun StorageTek EBS 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 Sun StorageTek EBS 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 Sun StorageTek EBS 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.
- Sun StorageTek EBS filemarks consume space on the tape.
  - The Sun StorageTek EBS 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 Sun StorageTek EBS 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:

- Automedia 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 Sun StorageTek EBS 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 .
2. If the serial number of the arm is reported, follow the procedure at "Scanning for libraries and devices" on page 141 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 Sun StorageTek EBS 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 Sun StorageTek EBS backups. The Sun StorageTek EBS 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 Sun StorageTek EBS services.
2. Shut down and restart the autochanger with the TZ89 drives.
3. When the autochanger is back online, restart Sun StorageTek EBS services. This resets Sun StorageTek EBS 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 926 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.

Managing optical drives with Solaris 9

With Solaris 9, the Volume Management daemon (**vold**) is changed so that it automatically attempts to manage all removable media devices. Because of this change, the Volume Management daemon may interfere with Sun StorageTek EBS operations related to optical drives.

To resolve this possible conflict, disable the **vold** daemon *or* modify the daemon configuration file.

Disable the vold daemon

To disable the volume management (**vold)** daemon:

1. Log in as **root** on the Sun StorageTek EBS storage node, and remove or rename the /etc/rc2.d/*volmgt script.

2. Enter the **/etc/init.d/volmgt stop** command.

Modify the daemon configuration file

To modify the daemon configuration file:

1. Log in as **root** on the Sun StorageTek EBS storage node, and open the daemon configuration file, /etc/vold.conf, in a text editor.
2. Comment out this line in the **Devices to Use** section.

   ```
   use rmscsi drive /dev/rdsk/c*s2 dev_rmscsi.so rmscsi%d
   ```

   After commenting out this line, the **Devices to Use** section of the configuration file looks similar to this:

   ```
   # Devices to use
   # use rmscsi drive /dev/rdsk/c*s2 dev_rmscsi.so
rmscsi%d
   ```

3. Save the configuration file.
4. Reinitialize the **Volume Management** daemon with the new configuration file settings. One way to do this is to send a hang-up signal to the daemon, for example:

   ```
   ps -ef | grep vold
   kill -HUP vold_pid
   ```

   where *vold_pid* is the process ID of the volume management daemon, **vold**.

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 Sun StorageTek EBS 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 Sun StorageTek EBS 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 Sun StorageTek EBS server or storage node. "Editing a client" on page 650 provides information about editing a Client resource.

SCSI reserve/release with dynamic drive sharing

When the Sun StorageTek EBS 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 Sun StorageTek EBS 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 Sun StorageTek EBS 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 Sun StorageTek EBS 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.

Tru64

SCSI reserve/release is only available on version 5.1B and later.

Device ordering issues

"Device ordering" on page 361 provides information about issues related to drive ordering, including how to determine if drive reordering has happened and procedures to correct the problem.

Sun StorageTek EBS locale and code set support

Sun StorageTek EBS 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.

Sun StorageTek EBS software might not function normally if the locale is changed. The previously existing indexes can become invalid.

Issues with fonts when using X servers on Linux

A number of issues exist in Linux X font support for non-english locales that may cause problems when displaying the **nwrecover** program. These problems may be resolved by switching to a different locale for the same language (for example, by using ja_JP.eucjp instead of ja_JP.utf8).

Resource database notes

With Sun StorageTek EBS release 6.2 and later, resource information resides in this structure:

```
<Sun StorageTek EBS_install_path>/res/nsrdb/00
                .
<Sun StorageTek EBS_install_path>/res/nsrdb/09
```

Each resource is stored 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 Sun StorageTek EBS client (**nsrexecd**) continues to use the <Sun StorageTek EBS_install_path>/res/nsrla.res file.

Viewing resources

You can view all Sun StorageTek EBS resources through the Administration window.

Although you can view the contents of the new Sun StorageTek EBS 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** *Sun StorageTek EBS_install_path*/**res**/**nsrdb**

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 Sun StorageTek EBS resource database files can be caused by a power outage, operating system crash, or manual editing of the database. If the Sun StorageTek EBS 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 904 provides information about viewing log files.

The Sun StorageTek EBS 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 Powerlink website for technical assistance.

Network and server communication errors

This section provides general, UNIX and Windows network and communication issues that you may encounter in a Sun StorageTek EBS 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 Sun StorageTek EBS servers, add the Sun StorageTek EBS 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 Sun StorageTek EBS client should contain these names for a Sun StorageTek EBS 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 Sun StorageTek EBS Setup program in maintenance mode and edit the Allowed Servers list. The *Sun StorageTek EBS Installation Guide provides details*.

Unapproved server error during client setup

If you add a Windows client to a UNIX Sun StorageTek EBS 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 Sun StorageTek EBS 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
- Member of the Administrator group

- Members of a group with Change Security Settings rights

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 Sun StorageTek EBS server may fail if multiple Sun StorageTek EBS 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 Sun StorageTek EBS 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 Sun StorageTek EBS server.

To enable the host to authenticate to the Sun StorageTek EBS server:

1. On the Sun StorageTek EBS host that cannot authenticate, stop the Sun StorageTek EBS client service. "Stopping and starting a Sun StorageTek EBS server, client, or storage node" on page 70 provides more information.
2. Delete the **nsrladb** database, which is located in *Sun StorageTek EBS_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 637 provides more information about reducing the media management database.

Firewall issues

If network communications are not functioning correctly, it may be related to incorrectly configured firewall rules on one or more of the hosts involved. Determine if a firewall is enabled and make sure that

it is correctly configured to allow communications with Sun StorageTek EBS programs. Appendix B, "Firewall Support," has information about firewall configuration.

1. Type this command:

   nwinstcreate -n *long_host_name*

   where *long_host_name* is the fully qualified host name, for example, *jupiter.accounting.company.com*. The *EMC NetWorker Command Reference Guide* or the UNIX man page provide more information about **nwinstcreate**.

2. Restart the Sun StorageTek EBS client service. "Stopping and starting a Sun StorageTek EBS server, client, or storage node" on page 70 provides more information.

   The host can now authenticate by using the long name.

UNIX communication issues

This section covers communication issues on UNIX networks.

Troubleshoot IP errors

If any of these error conditions occur, there may be an IP naming problem:

- RPC errors
- Unknown host messages
- Failure contacting the portmapper
- Connection failures or timeouts
- Program unexpectedly exiting
- Refused connections
- Remote command (the rcmd() function) to active client failures
- Name-to-address translation failures
- Program not registered messages
- Sun StorageTek EBS services failing to start
- Sun StorageTek EBS services failing to remain active
- Invalid path messages

To troubleshoot IP errors:

1. Document the steps you take and the results, especially error messages, in case you need to contact EMC Technical Support. This enables you to email or fax the steps and error message to EMC.
2. Set up a host table for each Sun StorageTek EBS client and Sun StorageTek EBS server. "Setting up a host table" on page 942 provides details.
3. Disable other name servers to simplify testing. "Disabling the DNS server" on page 943 provides details.
4. Use **ping** to establish basic connectivity. "Using ping to verify network connections" on page 943 provides details.
5. Use **rpcinfo** to verify that sessions can be established and that portmapping is correct. "Using rpcinfo to verify that sessions can be established" on page 944 provides details.

Setting up a host table

Troubleshoot IP problems by using only the host table. Troubleshooting by using only the host table does not mean you cannot use a name service, for example, DNS, with Sun StorageTek EBS software. Run tests by using only the host table to determine whether you have the software installed correctly. After you know the software works with the host table, enable the name server.

To set up a host table:

1. On the Sun StorageTek EBS client, list the client and the Sun StorageTek EBS servers to which it connects, for example:

```
127.0.0.1 localhost loopback
123.456.789.111 client client.domain.com
123.456.789.222 server server.domain.com
```

2. On the Sun StorageTek EBS server, list the Sun StorageTek EBS server and all of its clients, for example:

```
127.0.0.1 localhost loopback
123.456.789.111 server server.domain.com
123.456.789.222 client client.domain.com
```

3. Use the guidelines in "Using ping to verify network connections" on page 943 to ensure the highest success rate for parsing a host table within any operating system.

Recommendations for host table configuration:

- Do not use blank lines in the body of the host table.
- The end of the host table should always contain a blank line.
- The first un-remarked entry should always be the loopback line in the exact order and format shown in step 1 and step 2 .
- The last character of each unremarked line should be a space, not a carriage return.

On UNIX platforms, the host table resides in the /etc/hosts file.

Disabling the DNS server

To simplify the troubleshooting of hostname resolution problems, disable services like DNS, DHCP, and Windows Internet Naming Service (WINS).

Disable DNS just for the initial setup of the Sun StorageTek EBS clients and the Sun StorageTek EBS server to be tested. Only disable the ability of a client to obtain IP naming information from a DNS server.

To disable the DNS server on most UNIX platforms, rename the file /etc/resolv.conf and reboot. Alternatively, set up the IP name search order so that the host table is searched before DNS.

Note: If the hostname cannot be resolved, the Sun StorageTek EBS server caches this information for five minutes before attempting the resolution again. If the hostname is successfully resolved, the information is cached for 30 minutes.

Setting the IP name search order

To set the IP name search order:

1. Edit the /etc/nsswitch.conf file and verify that the /etc/resolv.conf file exists.
2. Set the host file to be first in the search order, with DNS second and NIS last. For example:

```
hosts: files [NOTFOUND=continue] DNS
[NOTFOUND=continue] nis
```

Using ping to verify network connections

After you create the host table, test it using **ping**.

On the Sun StorageTek EBS client, **ping**:

- Client short name (hostname) from the client
- Client long name (hostname plus domain information) from the client
- Client IP address from the client
- Server short name from the client
- Server long name from the client
- Server IP address from the client

This example shows how to **ping** the client short name and client long name from a Sun StorageTek EBS client called *mars* in the *jupiter* domain:

```
ping mars
ping mars.jupiter.com
```

On the Sun StorageTek EBS server, **ping**:

- Server short name from the server
- Server long name from the server
- Server IP address from the server
- Client short name from the server
- Client long name from the server
- Client IP address from the server

Using rpcinfo to verify that sessions can be established

If **ping** is successful and backup problems still exist, test with **rpcinfo**. Use **rpcinfo** to test the operation of the portmapper. Use the same tests with **rpcinfo** as with **ping**.

To use **rpcinfo**, the host whose hostname you enter must have a portmapper running. In most cases, EMC portmappers are compatible with third-party portmappers. If you are using a product that provides its own portmapper, do not load the third-party portmapper until you have verified that the Sun StorageTek EBS software works with the rest of the environment.

The **rpcbind** service must be running. The **rpcinfo** utility is part of the operating system.

The syntax for using **rpcinfo** to display ports using TCP is:

```
rpcinfo -p hostname
```

Substitute the long name and short name for the variable *hostname*.

View other **rpcinfo** options by typing **rpcinfo** at the command-prompt. The UNIX man page provides more information. Repeat **rpcinfo** using all the locations and all the iterations listed in this document for **ping**.

When **rpcinfo** runs successfully, the output is a list of port numbers and names. For example:

```
rpcinfo for mars
program vers proto   port
100000    2   tcp    111  portmapper
100000    2   udp    111  portmapper
390103    2   tcp    760
390103    2   udp    764
390109    2   udp    764
390107    4   tcp    819
390107    5   tcp    819
```

Binding to server errors

Sun StorageTek EBS 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 Sun StorageTek EBS services are not running and a Sun StorageTek EBS 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 Sun StorageTek EBS services **nsrd**, **nsrexecd**, **nsrindexd**, **nsrmmd**, and **nsrmmdbd** might not be running.

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 120 lists the location for the boot-time file. The **nsrexecd** man page provides information about **nsrexecd**.

Table 120 Boot-time file locations

Operating system	Boot-time file
AIX	/etc/rc.nsr
HPUX	/sbin/init.d/networker
IRIX (SGI)	/etc/rc2.d/S95networker
SCO	/etc/rc2.d/S95networker
Solaris	/etc/init.d/networker
SunOS 4.1.x	/etc/rc.local
Ultrix	/etc/rc./local
Others	/etc/rc2.d/S95networker

Microsoft Windows communication issues

This section covers communication issues on Windows networks.

Certified protocols

This release of Sun StorageTek EBS 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 Sun StorageTek EBS software at this time.

Resolving connection problems

This section addresses DNS and local host file related connection problems and solutions.

During installation, this error message might appear after an apparently successful client-server connection.

hostname `is not a valid hostname.`

This type of error message indicates possible problems with hostname resolution. This section notes items to check when you encounter connection problems.

IP configuration

If you experience IP configuration problems:

- If you are using DNS, ensure the IP address of the Windows computer is correct.
- If you are not using DNS, ensure the host table is set up correctly.
  - The host table is in the hosts file located in %SystemRoot%\system32\drivers\etc
  - Ensure that the hostname is set correctly.
  - Check the host table on both the server and client computers. Check that all hostnames, fully qualified domain names, and IP addresses are listed.
  - Both server and client hostnames should be in the host table with the IP address, short names, and long names in this order:

```
Server A Host Table
127.0.0.1localhostloopback
xxx.xxx.xxx.xxxserver_aserver_a.domain.com
xxx.xxx.xxx.xxxclient_a client_a.domain.com
Client A Host Table
127.0.0.1localhostloopback
xxx.xxx.xxx.xxxclient_aclient_a.domain.com
xxx.xxx.xxx.xxxserver_a server_a.domain.com
```

To ensure that the computer returns the correct hostname:

1. In the **Control Panel**, access the dialog box for DNS configuration.
2. Enter the hostname in the **Hostname** text box.
3. Edit the domain name in the **Domain Name** text box. The hostname and the domain name together make up the fully qualified domain name.
4. To ensure that TCP/IP is set up correctly, check the hostname and fully qualified domain name by typing this at the command-prompt:

```
ipconfig /all
ping hostname
```

Domain Name Server

If you experience DNS problems, ensure that:

- The DNS server address is correct.
- The DNS server is running.
- The default gateway is set up correctly.
- DHCP is set up correctly. If you are using DHCP, the Sun StorageTek EBS server must have a reserved IP address.

Connectivity

If you experience connectivity problems, ensure that:

- The Windows computer can **ping** its destination using both the destination's hostname and its fully qualified domain name.
- The Windows computer can **ping** itself using both its hostname and its fully qualified domain name.

Sun StorageTek EBS software

If you experience Sun StorageTek EBS software problems, ensure that:

- You are using the **-s server** argument when starting a Sun StorageTek EBS utility from the command prompt.
- There is a space in the path for the **-s server** argument or the path to the Sun StorageTek EBS program.
- The Sun StorageTek EBS server to which you are connecting is running.

Tips

These tips can help prevent problems using the Sun StorageTek EBS software:

- Use the host table instead of DNS to test if DNS is working properly.
- Eliminate DHCP and use a static IP address instead.
- Use the 8.3 file naming convention for directory names in a path. Do not include spaces in pathnames.

Changing the Sun StorageTek EBS server address

When the TCP/IP address changes on the Sun StorageTek EBS server, the Sun StorageTek EBS 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 Sun StorageTek EBS software stops working.

If you are using DHCP, use a static IP address for the Sun StorageTek EBS server.

New.Net and Sun StorageTek EBS 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 Sun StorageTek EBS software. This causes many Sun StorageTek EBS programs, including **save.exe**, to fail on exit. This is a New.Net problem that Sun StorageTek EBS 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.

Note: Do not manually delete the newdotnet.dll file. Doing so renders the system unusable.

Troubleshooting port configurations

The following error messages might appear if there are problems with port configurations:

- Cannot bind socket to service port in configured range on system *hostname*

  This indicates that the configured service port range is not large enough.

  To resolve this problem:

  1. Increase the service port range configured for the specified hostname.
  2. Set the firewall rules to allow inbound and outbound TCP/UDP packets within the port range specified in step 1 .

- `Cannot bind socket to connection port in range on system` *hostname*

This indicates that the configured connection port range is not large enough. The current backup or recovery session has been stopped.

To resolve this problem:

1. Increase the connection port range for the specified hostname.
2. Set the firewall rules to allow inbound and outbound TCP/UDP packets within the port range specified in step 1 .

Conflicting portmappers

Portmappers used with Hummingbird Maestro and Exceed software earlier than version 5.1 conflict with the Sun StorageTek EBS portmapper. This conflict occurs because Maestro and Exceed software install a service named **HCL Inetd**. You cannot successfully install the Sun StorageTek EBS software while this service is running. You must disable the **HCL Inetd** service through the Services window in the Control Panel before installing the Sun StorageTek EBS software (the Maestro and Exceed software continue to run).

If you installed both the Sun StorageTek EBS software and the Hummingbird software and you have difficulties with portmappers, contact Hummingbird Technical Support.

Sun StorageTek EBS 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 Sun StorageTek EBS 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 Sun StorageTek EBS administrator privileges for root@*client_system* in the Administrator attribute in the Server resource. However, be aware that Sun StorageTek EBS 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 904 provides information about viewing log files.

Filenames with carriage return not allowed

If you archive and attempt to groom a file that has a carriage return as part of the name, the grooming operation will fail with a "could not start" error. Filenames that include carriage returns are not allowed.

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 Sun StorageTek EBS services.
2. Remove the environment setting for NSR_MMDCONTROL.
3. Restart the Sun StorageTek EBS 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 timeout attribute for storage node operations" on page 128 provides more information.

Console error messages and corrective actions

Table 121 provides a list of Console error messages or symptoms and corrective actions to take.

Table 121 Error messages or symptoms (1 of 4)

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 Navigator: JavaScript (enabled by default) has been disabled.	In Internet Explorer: Lower the web browser's security setting or enable Active Scripting. In Navigator: Enable JavaScript.
Authorization code not accepted.	Sun StorageTek EBS 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.	• Ensure that the Console server is running. "Console troubleshooting notes and tips" on page 959 provides details. If it is not, close all application windows and check the gstd log file for errors. "Viewing log files" on page 904 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.

Table 121 Error messages or symptoms (2 of 4)

Error message or symptom	Possible cause	Corrective action
Connection refused: no further information. or Problem contacting server *server_name:*	Console server is in the process of crashing or has already crashed.	Check to see if the Console server is running. • 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 904 provides information about viewing log files.
	Console server has been started within the previous few minutes.	Wait a couple of minutes and retry.
Database fetch operation failed.	Console database is corrupt.	Recover the database. "Recovering the Console database" on page 456 provides details.
Display problem: In Internet Explorer: The page cannot be displayed. In Navigator: Netscape search results page appears.	Console server is not running.	Restart the Console server.
	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 957 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.	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 "*pathname*" already exists.	• 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.	• 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

Table 121 Error messages or symptoms (3 of 4)

Error message or symptom	Possible cause	Corrective action
Could not contact License Manager on <*hostname*>. - or - Program not registered.	License Manager hostname has not been assigned or License Manager is not running or installed. **Note:** Use of the License Manager is optional. The License Manager is only available from EMC.	If you are using the License Manager and a hostname has not been assigned: 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. **Note:** The *NetWorker License Manager Installation and Administrator's Guide* provides details.
	Sun StorageTek EBS client was stopped, but the License Manager was not stopped, and then the Sun StorageTek EBS client was restarted. Although both services are now running, Sun StorageTek EBS client must be started *before* License Manager is started. If the services are not started in the correct order, an error condition occurs.	1. Stop the Sun StorageTek EBS software. 2. Stop License Manager if it is running. 3. Restart License Manager. 4. Restart the Sun StorageTek EBS software.
License allocation failed.	Temporary license for Sun StorageTek EBS software is expired.	Enter enabler codes and register the product.
License managed event indicates that license is expiring/expired even though it has been authorized.	License has been authorized within the last 24 hours.	None needed. To remove the managed event from the display, dismiss the event or it 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.

Table 121 Error messages or symptoms (4 of 4)

Error message or symptom	Possible cause	Corrective action
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 Sun StorageTek EBS Console".	Java Web Start preferences are set to something that is incompatible with the rest of the environment. (For example, a proxy server has been set up that stops Java Web Start from downloading the Console client software from the Console web server.) This error message may also occur if the Console is being launched on a localized operating system and the Java Web Start cache path contains non-English characters.	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. (In the proxy server example, go to the General tab of the Preferences dialog box and select None, for Proxies.) If the Java Web Start cache path contains non-English characters, change the path to contain no non-English characters.
gstd.log file error: internal error: could not end transaction	When 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
- gstd
- db output
- dbstop 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/SUNWebsg/logs.

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.

By default, the gst log file is located in /opt/SUNWebsg/logs.

"Viewing log files" on page 904 provides information about viewing log files.

Table 122 lists the variables that control the gstd.log file.

Table 122 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 575 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 Sun StorageTek EBS servers), open a separate instance of the Console server to run each operation.

- Are the following processes running?
  - GST server (**gstd**)
  - Database server **(dbsrv9**)
  - GST web server (**gsttclsn**)

  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/SUNWebsg/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:

  ```
  SUNWebsg root6140   1   0   12:54:10 ?0:03
  /opt/SUNWebsg/sybasa/bin/dbsrv9
  ```
- 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:

  ```
  SUNWebsg root6140   1      0   12:54:10 ?0:03
  /opt/SUNWebsg/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 Sun StorageTek EBS Console

After the Sun StorageTek EBS 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 Sun StorageTek EBS console may fail to launch.

Workaround

Remove the Sun StorageTek EBS Management Console Application and Language Pack (if applicable) from the Java Application Cache Viewer:

1. Run the **Java Application Cache Viewer**. From the command line, use the **javaws** command to launch the application.
2. Select the **Sun StorageTek EBS Management Console Application and Language Pack**.
3. Click **Remove Selected Application**.
4. Clear the temporary files used by Java applications:
   a. In the **Java Application Cache Viewer**, select **Preferences** from the **Edit** menu.
   b. 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.

A

SYSTEM and VSS SYSTEM Save Sets

This appendix covers these topics:

- SYSTEM save sets 964
- VSS SYSTEM save sets 970

SYSTEM save sets

The SYSTEM save sets discussed in this section work with these operating systems:

- Windows 2000
- 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, Sun StorageTek EBS software backs up all Windows system state components except the SFP component. "Components of the SYSTEM FILES save set" on page 965 provides information about how Sun StorageTek EBS software treats the SFP component.

Certain basic system state components are present with every Windows Server 2003, Windows XP Professional, or Windows 2000 installation. These components are present on every system, and are always part of the SYSTEM STATE save set. Table 123 lists these basic components and references to their backup and recovery procedures.

Table 123 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 1006 provides details.
Registry	"Windows registry" on page 1007 provides details.
Performance Counters	None

Optional components of the SYSTEM STATE save set

In Windows 2000, 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 124.

Table 124 SYSTEM STATE save set optional components

SYSTEM STATE optional component	Special backup and restore considerations
Active Directory (AD)	"Active Directory" on page 998 provides details.
Certificate Server	None
Cluster Server	Chapter 19, "Cluster Support" provides details.
File Replication Service (FRS, also called SYSVOL)	None

Components of the SYSTEM FILES save set

In Windows 2000, 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, ntdetect.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 Sun StorageTek EBS SYSTEM STATE save set, backing up or recovering the system state would be extremely resource-intensive. Therefore, Sun StorageTek EBS 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), Windows XP Professional, or Windows 2000 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*, Sun StorageTek EBS software backs up *all* of the system-protected files.

Components of the SYSTEM DB save set

The Sun StorageTek EBS SYSTEM DB save set is used to back up the Windows Server 2003 32 and 64 bit, Windows XP Professional, or Windows 2000 system databases that are installed and started.

Table 125 on page 967 lists the system databases that are present by default and references to their backup and recovery procedures.

Table 125 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 104 and "Restoring Windows Content Index Server on Windows 2000 or later" on page 470 provide details.
Disk Quota Database	"Granting full permissions for backup of Disk Quota database" on page 108 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.

Optional component of the SYSTEM DB save set

If the Terminal Services Licensing (TSL) database is installed and the corresponding TSL service is started, the database is also backed up as part of the SYSTEM DB save set.

"Terminal server licensing" on page 1008 provides information about TSL backup and recovery by using Sun StorageTek EBS software.

Windows databases not included in the SYSTEM DB save set

Table 126 lists the additional Windows Server 2003, Windows XP Professional, and Windows 2000 databases that Sun StorageTek EBS 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 126 Windows databases not in SYSTEM DB save set(1 of 2)

Database	Backup and restore procedures
Distributed File System (DFS)	Appendix C, "Backing Up and Restoring a Microsoft DFS" provides details

Table 126 Windows databases not in SYSTEM DB save set(2 of 2)

Database	Backup and restore procedures
Encrypting File System (EFS)	"Encrypting file system" on page 1004 provides details.
Event logs	"Event logs" on page 1006 provides details.
Sparse files	"Sparse files" on page 1008 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:

- Content source crawls scheduled with the Microsoft Windows 2000 Scheduled Tasks utility. You must re-create these on the restored server by using the schedule settings on the Properties page for each content source.
- 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.

Components of the ASR save set

The ASR (Automated System Recovery) save set contains all the information necessary to return the failed computer to its condition at the time of the last ASR backup, including:

- Automated reinstallation of Windows
- Restoration of the system configuration
- Recovery of one or more disk volumes

The ASR save set can be backed up by itself, or as a component of save set *All*. Like other SYSTEM save sets, ASR is automatically backed up at level full.

Chapter 23, "Support for Microsoft Automated System Recovery" provides more information about Sun StorageTek EBS software support for ASR.

VSS SYSTEM save sets

The VSS SYSTEM save sets discussed in this section work with Windows Server 2003, where VSS is licensed and enabled, and Windows Server 2008.

The VSS SYSTEM SAVE SET includes the VSS System Boot, VSS System Fileset, and VSS System Services, along with three non-system components, VSS User Data, VSS Other and VSS ASR Disk (Windows Server 2003 only).

Components of the VSS SYSTEM BOOT save set

The VSS SYSTEM BOOT save set includes all elements of the Windows system state, except the VSS System Writer writer and the SFP component. It also includes the ASR System Files legacy component. All other components are dynamically generated, and therefore may differ each time Sun StorageTek EBS software runs.

"Components of the VSS SYSTEM FILESET save set" on page 970 provides information about how Sun StorageTek EBS software treats the SFP component.

Components of the VSS SYSTEM FILESET save set

The VSS SYSTEM FILESET save set includes the VSS System Writer writer, the SFP component, and the IA-64 EFI FAT partition legacy component. This save set component is recovered by using the legacy method of recovery.

Most of the elements of the Windows system state are part of the VSS SYTEM BOOT save set. However, the SFP component typically consists of more than 200 MB of data in more than 1,500 files. If the SFP component were part of the Sun StorageTek EBS VSS SYSTEM BOOT save set, backing up or recovering the system state would be extremely resource-intensive. To relieve the strain on resources, Sun StorageTek EBS software backs up the SFP component as part of the VSS SYSTEM FILESET save set.

"Components of the SYSTEM FILES save set" on page 965 provides more information about the SFP component.

System File Protection (SFP)

For a full backup operation, specification of the VSS SYSTEM FILESET 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 VSS SYSTEM FILESET 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.

Note: 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 Disaster Recovery Guide* provides more information.

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 corresponding to the system component installation date. They also 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 Sun StorageTek EBS backup of save set *All* or the VSS SYSTEM FILESET save set detects new SFP files, Sun StorageTek EBS software checks for the more recent of the file creation and file modification dates. If the more recent date is after the *as of time* value, Sun StorageTek EBS software backs up *all* of the system-protected files.

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 Sun StorageTek EBS software runs.

The FSRM Disk Quota component can only be backed up with a VSS enabled backup, which is requires a VSS license.

Windows Server Cluster writers

Backup and recovery of Windows Server cluster is supported for Windows Server 2003 and Windows Server 2008. There are some differences in the Cluster Writer save set names between Windows Server 2003 and Windows Server 2008.

In Windows Server 2003, the Sun StorageTek EBS 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.

In Windows Server 2008, the Sun StorageTek EBS software uses the following specifications:

- Cluster writer name: Cluster Database
- Backup and recovery save set: Under the VSS SYSTEM SERVICES system save set

Components of the VSS USER DATA save set

The VSS USER DATA save set includes application data that is not system-critical. Components of this save set are dynamically generated and may differ each time Sun StorageTek EBS software runs.

Components of the VSS OTHER save set

The VSS OTHER save set includes application data that is not system-critical. Components of this save set are dynamically generated and may differ each time Sun StorageTek EBS software runs.

Components of the VSS ASR DISK save set

The VSS ASR DISK save set includes the ASR legacy component. This is a legacy save set and is recovered by using the legacy method of recovery.

Note: The VSS ASR DISK save set is not supported by NetWorker for Windows Vista or Windows Server 2008.

B

Firewall Support

This appendix covers these topics:

- Overview 976
- Configuring the port ranges in Sun StorageTek EBS software .. 976
- Ports needed in the service ports range for Sun StorageTek EBS 7.3.x and later 979
- Ports needed in the service ports range for Sun StorageTek EBS 7.2.x. 982
- Ports used by Sun StorageTek EBS add-on products 983
- Example of a firewall configuration 987
- Access control to local host resources 988
- Firewall port enhancement 989

Overview

Many differences exist in the number of ports used by Sun StorageTek EBS release 7.3 and later, from the earlier releases. Also if you operate in a mixed environment, calculating the ports required by the Sun StorageTek EBS software is more complicated with the later releases. Consequently, this appendix is divided into sections that specifically describe each Sun StorageTek EBS release. Separate sections describe mixed environments and Sun StorageTek EBS add-ons.

This appendix also provides an example of firewall configuration, access control to local host resources, and firewall port enhancement.

Configuring the port ranges in Sun StorageTek EBS software

Sun StorageTek EBS software and some add-on products use ports from a range that is stored in the nsrla database. The port ranges are stored in the "NSR system port ranges" resource and there are several ways to view or update their values. There are two sets of port ranges:

- Connection port range
- Service port range

Connection ports range

TCP client side ports are only selected from this range, which should be kept as wide as possible. There are no security concerns in keeping this range as wide as possible although Sun StorageTek EBS performance problems or random malfunctions may occur if the range is too narrow. In release 7.2 the default value for this range is: 10001-30000. In Sun StorageTek EBS 7.3, the default value, which is 0-0, indicates that the Sun StorageTek EBS software will allow the operating system to select the port for TCP clients. You can only do this when using Sun StorageTek EBS 7.3 and later.

Service ports range

The Sun StorageTek EBS TCP/UDP server processes use ports in this range. The default value for the service ports range is: 7937-9936,

although this range can be narrowed down (as discussed in this document).

When configuring a firewall, you must open the port range described by the "Service ports" range in the firewall, for TCP connections. The Sun StorageTek EBS software does use some ports from this range for UDP services, but these are not essential. UDP service ports must be opened in the firewall for TCP connections, although there is no need to allow UDP connections through the firewall. The ports required by various daemons are selected at random from the service ports range.



IMPORTANT

The nsrexecd command, when used on every type of Sun StorageTek EBS host, will always try to listen on 7937 and 7938. These two ports are used regardless of what you set the service ports range to, provided that another process is not already listening on those ports when the Sun StorageTek EBS software starts. The Sun StorageTek EBS software requires that either port 7938 be available, or for rpcbind (portmapper) to be running and available through the firewall. Otherwise the Sun StorageTek EBS software will not function correctly. Ports 7937 and 7938 are included in the number of ports used by nsrexecd. The service port range must therefore not include these ports.
Additionally, Port 514 is used as a fallback connection if communication with nsrexecd cannot be established. To avoid potentially slow performance with the connection, ensure that port 514 is not blocked.

Anytime you change the service or connection port range used by a Sun StorageTek EBS host, you must restart all Sun StorageTek EBS processes on that computer for those changes to take effect. You must do this no matter which method you use.

In order to change the port ranges on a system, you must have update access to the "NSR system port ranges" resource in nsrexecd, that is, you must be in the "administrator" list for that resource. “Access control to local host resources” on page 988 provides more information about access control on this resource.

Using the nsrports command to update port ranges

The **nsrports** program is used to view and update the port ranges. To view port ranges, enter:

```
$nsrports -s shadow.emc.com
```

In this example, the port ranges for Sun StorageTek EBS host shadow.emc.com are displayed:

```
Service ports: 7937-9936
Connection ports: 0-0
```

To update the service ports range, enter:

```
$nsrports -s shadow.emc.com -S 7937-8937
```

As seen in this example, the service port range on shadow.emc.com was changed from 7937-9936 to 7937-8937. You can enter multiple ranges for those cases where you can not have a a solid range. This is illustrated in this example

```
nsrports -s shadow.emc.com -S 7937-7938 10000-11000 11500
```

In this case the service ports range for host: shadow.emc.com is set to use ports 7937-7938, 10000-11000, and 11500.

Using the Console to update the port ranges

You can only use the Console to update the port ranges if you are using the Console server that ships with the current Sun StorageTek EBS release.

To use the Console to update the port ranges on a host:

1. Go to the **Configuration** screen.
2. Select the **Local Hosts** entry for the computer on which you want to update the port ranges. You select the entry from the left side of the screen in tree view.
3. Right-click the **Local Hosts** entry.
4. Select **Configure Port Ranges...**.
5. Change the value of the desired range. For multiple ranges, type one range per line.
6. Click **OK** when complete.

Using nsradmin to update the port ranges

To view the current port ranges on a system, type:

```
$nsradmin -p nsrexec -s shadow.emc.com

NetWorker administration program.
Use the "help" command for help, "visual" for full-screen
mode.
nsradmin> p type: nsr system port ranges
   type: NSR system port ranges;
   service ports: 7937-9936;
   connection ports: 0-0;
   administrator: *@localhost;
nsradmin>
```

To update the service ports range, type:

```
$nsradmin -p nsrexec -s shadow.emc.com

NetWorker administration program.
Use the "help" command for help, "visual" for full-screen
mode.
nsradmin> . type: nsr system port ranges
Current query set
nsradmin> update service ports: 7937-7938, 9000-9100
                service ports: 7937-7938, 9000-9100;
Update? y
updated resource id
0.13.3.219.0.0.0.0.0.0.0.0.67.30.89.147.137.69.168.98(6)
nsradmin>
```

Multiple port ranges were created in the previous example. Notice that each range is separated by commas.

Ports needed in the service ports range for Sun StorageTek EBS 7.3.x and later

The Console is shipped and installed as part of Sun StorageTek EBS 7.3.*x*. Because of this, a new required component called the Console server is available. The Sun StorageTek EBS Console server must be installed to manage and monitor Sun StorageTek EBS servers. It is best to install it on just one computer in the network to take full advantage of the Console's consolidated reporting feature. As in previous versions of Console, the ports used by the Console server are configured separately, in separate configuration files, from the ports that the Sun StorageTek EBS software uses. As a result, the ports that the Console server uses are not calculated by **nsrports**, in the normal range of ports. This also means that these ports can be addressed separately from the ports calculated for the rest of the Sun

StorageTek EBS software. "Sun StorageTek EBS Console" on page 983 provides additional information on configuring these ports.

These daemons were introduced in Sun StorageTek EBS 7.3 and are shipped with Sun StorageTek EBS release 7.3.*x*:

- **nsrmmgd**
- **nsrlcpd**
- **nsrjobd**

Each of these daemons use one port and two new ports are used by **nsrexecd**. This demonstrates why the configuration of Sun StorageTek EBS 7.3 is different from that performed for Sun StorageTek EBS 7.2.*x*.

The Sun StorageTek EBS server

These daemons only run on the server:

- **nsrd**
- **nsrmmdbd**
- **nsrindexd**
- **nsrmmgd**
- **nsrjobd**

Each of these daemons uses one port for TCP servers. The nsrexecd daemon also runs on the server and it uses 4 ports. Both the **nsrlcpd** and **nsrmmd** daemons run on either a storage node or on the server (because the server is acting as a storage node). These daemons take one port per instance. There is one **nsrlcpd** daemon for each jukebox managed by the Sun StorageTek EBS software, and this is started on the computer that manages that jukebox. This computer can be either the Sun StorageTek EBS server or storage node.

There is one **nsrmmd** daemon for each Sun StorageTek EBS configured device, unless that is an adv_file type device. The adv_file type device requires two nsrmmds be started. The **nsrmmd** daemon starts on the computer that manages the device. When spanning from one medium to another, an extra **nsrmmd** is started to mount the new tape. This **nsrmmd** is known as a 'helper mmd' and each instance requires one port.

Helper **mmds** take dynamic ports, so that even though the ports are not in use when you start the Sun StorageTek EBS software, they will

be in use when you span media. Since multiple devices can span simultaneously, it is best to assume that each device requires two ports.

Unfortunately, the processes **nsrd** and **nsrmmgd** processes each require one port for UDP services. These ports do not really need to be accessed from outside the firewall, but they are taken from the service ports range configured using **nsrports**. Therefore, they need to be added to the formula to calculate the number of ports needed in the service ports range.

The formula for how many ports the Sun StorageTek EBS server requires is therefore:

11 + 2 * #*devices* + #*jukeboxes*

where:

- #*devices* is the number of devices configured on the Sun StorageTek EBS server only.
- #*jukeboxes* is the number of jukeboxes configured on the Sun StorageTek EBS server only.

Sun StorageTek EBS storage node

The only storage node daemons that are run are:

- **nsrexecd**
- **nsrmmd**
- **nsrlcpd**

As described previously, **nsrexecd** uses four ports, **nsrmmd** uses two ports per device, and nsrlcpd uses one port per jukebox. From this, the number of ports required by a Sun StorageTek EBS storage node is:

4 + 2 * #*devices* + #*jukeboxes*

where:

- #*devices* is the number of devices configured on the Sun StorageTek EBS storage node only.
- #*jukeboxes* is the number of jukeboxes configured on the Sun StorageTek EBS storage node only.

Sun StorageTek EBS client

The only daemon running on the Sun StorageTek EBS 7.3 client is **nsrexecd** and it uses four ports. Therefore the port usage formula for a Sun StorageTek EBS client is 4.

Ports needed in the service ports range for Sun StorageTek EBS 7.2.x.

In Sun StorageTek EBS release 7.2.*x*, the dynamic ports used by nsrexec become redundant.

Sun StorageTek EBS server

These daemons only run on the server:

- **nsrd**
- **nsrindexd**
- **nsrmmdbd**

Each of these daemons uses one port although, **nsrexecd** also runs on the Sun StorageTek EBS server and it uses two ports. The **nsrmmd** process can run on either the server or storage node. It runs on the computer responsible for managing the device. One **nsrmmd** process is started per device (except for adv_file type device which has two **mmds** started). During spanning, one additional **nsrmmd** helper, which uses a dynamic port, is started. Since spanning can simultaneously take place on all devices, **mmd** daemons essentially use two ports per device configured on the host in question.

Unfortunately, the **nsrd** process uses one port for a UDP service. This port does not really need to be accessed from outside the firewall, but the port is taken from the service ports range configured by using **nsrports**. Therefore, this port needs to be added to the formula to calculate the number of ports needed in the service ports range.

The formula for how many ports the Sun StorageTek EBS server requires is therefore:

6 + 2 * *#devices*

The #devices only includes those devices that are configured on the Sun StorageTek EBS server.

Sun StorageTek EBS storage node

Both **nsrexecd** and **nsrmmd** are the only daemons that run on the storage node. The **nsrexecd** (like on the server) daemon uses two ports, while the **nsrmmd** daemon uses two ports for each device configured on the host.

The Sun StorageTek EBS 7.2.*x* port usage formula is:

2 + 2 * #*devices*

Note: The #*devices* only includes the devices configured on the Sun StorageTek EBS storage node in question.

Sun StorageTek EBS client

The **nsrexecd** daemon, which uses two ports, is the only one running on the Sun StorageTek EBS client. Therefore, the formula for port usage on the Sun StorageTek EBS client is 2.

Ports used by Sun StorageTek EBS add-on products

This section details the ports used for various Sun StorageTek EBS add-on products. Some add-on products use ports from the range configured using **nsrports**, others have their own configuration files.

If a Sun StorageTek EBS add-on product uses ports from outside the range configured by using **nsrports**, they do not need to be included in the formula used to calculate the range entered using **nsrports**. However, it is still important that the ports used by that add-on, ***not*** overlap the range used by the Sun StorageTek EBS software on the host where the add-on is installed. If the add-on product ports do overlap the Sun StorageTek EBS port range, neither product will be able to use the ports they require.

Sun StorageTek EBS Console

The Sun StorageTek EBS Console is used to administer the Sun StorageTek EBS server and can also administer some of the Sun StorageTek EBS Client resources.

The Console server requires these three ports:

- One port is used for the web server. This is required to provide a way to download the java application code which acts as the Console front end. This port is selected during the installation process.
- The second port is used for RPC calls from the Console Java client to the Console server. This port is selected during the installation process.
- The last port is used for database queries. The default number for this port is 2638.

The ports used for the Console are **not** taken from the range configured by using **nsrports**.

How to enable Console client/server connections through a firewall

To enable clients to connect to the Console server through a firewall, configure the firewall for communication over the HTTP, Client Service, and TDS (Tabular Data Stream protocol) port numbers:

1. Determine the two port numbers used for connecting to the Console server.

   Note: To identify the specified ports, refer to the GST server package in the install.log file. Look for the following expressions as part of larger expressions:

   - http_svc_port -t int -v *http_service_port*
   - clnt_svc_port -t int -v *client_service_port*

   where *http_service_port* and *client_service_port* are port numbers. By default the HTTP service port is set to 9000 and the client service port is set to 9001.

2. Use the firewall software to accept inbound and outbound (TCP packets from the HTTP and client service ports determined in step 1.
3. Use the firewall software to accept inbound and outbound TDS packets from the Database Communications port (port 2638).
4. Repeat step 2 and step 3 for any other firewalls used with this connection.

Sun StorageTek EBS Console and the software distribution feature

If the software distribution feature is being used to remotely distribute Sun StorageTek EBS software and the Sun StorageTek EBS

Console server is running on a different host than the Sun StorageTek EBS server, and a firewall is in use between the Sun StorageTek EBS server and the Sun StorageTek EBS Console server, an additional port will be required. This port will be taken from the taken from the range configured for the Sun StorageTek EBS software.

AlphaStor

If there is a firewall between AlphaStor and Sun StorageTek EBS software, the following ports should be opened:

- With Sun StorageTek EBS release 7.2:
  - Port 44444 should be opened between the AlphaStor software and the Sun StorageTek EBS server.

    This port uses long-running connections. The firewall should not be configured to close the port due to inactivity to prevent media from being unmounted during backup.
  - Port 41025 should be opened between the AlphaStor software and the Sun StorageTek EBS server and/or Sun StorageTek EBS client for command line operations.
- With Sun StorageTek EBS release 7.3 and newer:
  - Port 44475 should be opened between the AlphaStor software and the Sun StorageTek EBS server.

Additionally, if there is a firewall between the AlphaStor software and the host running the AlphaStor Graphical User Interface, port 44470 should be opened in the firewall.

Ports required for communication between AlphaStor software and storage libraries

If there is a firewall between the AlphaStor software and a host with an AlphaStor-controlled library connected to it, the following ports should be opened:

- Port 3500 should be opened between the AlphaStor software and the host controlling the library's robotics.
- Port 3000 should be opened between the AlphaStor software and the host access the drives.

Client Backup Configuration Wizard

The Client Backup Configuration wizard speeds the process by which a Sun StorageTek EBS client is configured, on the Sun StorageTek EBS server, for being backed up. Prior to Sun StorageTek EBS release 7.3, the client backup configuration wizard was shipped as a separate add-on product. With Sun StorageTek EBS release 7.3 it became, like the Console, a part of the main Sun StorageTek EBS product.

The Sun StorageTek EBS Client Backup Configuration wizard uses one port for each open user interface on the Sun StorageTek EBS client. This is a dynamic port and that is closed when the wizard is closed.

These ports are selected from the port range configured by using **nsrports**.

NDMP

The NDMP extension uses 1 transient port (port number 10,000) whenever an NDMP filer is being backed up. When backing up an NDMP Server to a Sun StorageTek EBS storage node that uses DSA, DSA is dynamically spawned for each backup and recover. As it is spawned, DSA listens on a port that falls within the Sun StorageTek EBS service port range and closes the connection once the backup or recover is finished. DSA listens on an unused transient port from within the Sun StorageTek EBS service port.

Sun StorageTek EBS modules

SnapImage uses 1 transient port (port number 10,000), specified in /etc/services much like other standard services, although this can be changed manually during installation.

The data connection used between the Sun StorageTek EBS storage node and the SnapImage data server (if they are remote), is within the Sun StorageTek EBS service port range. Another NDMP service (NDMP tape server) uses the system reserved port (0-1024 range) to accept the connection. The connecting client (NDMP data server) uses the ephemeral ports generated by the system.

Example of a firewall configuration

The following is an example of a firewall configuration.

ABC Company has a Sun StorageTek EBS server behind a bi-directional firewall that controls both incoming and outgoing traffic. The Sun StorageTek EBS server has a locally attached media library with six devices. This will require 24 service ports: 11 + 2 * [number of devices] + [number of libraries]. Therefore, the server must be configured to use 24 service ports and the firewall must allow incoming traffic from outside the firewall to the server on these 24 ports.

The company also has three storage nodes outside the firewall. Each storage node has one library with six devices. Each storage node will require 17 service ports: 4 + 2 * [number of devices] + [number of libraries]. The firewall will need to be configured to allow outgoing traffic from the server through the firewall to the storage nodes. However, each storage node can use the same service ports. Therefore, the firewall only needs to allow outgoing traffic to the storage nodes' IP addresses through a total of 17 ports. Furthermore, these can be the same ports that the server uses for its service ports.

The company also has several clients outside the firewall. Each of these clients requires four service ports. The firewall will have to allow outgoing traffic from the server to the clients on the clients' service ports. However, all client can use the same four service ports. Furthermore, the ports can be the same ports that the storage nodes use for their service ports.

Therefore, the simplest configuration for ABC Company to allow all necessary communication through their firewall would result in a total of 24 ports opened for incoming communication to the server, with 17 of those ports also opened for outgoing communication from the server to the clients and storage nodes.

Multiple firewalls

The ABC Company can include multiple firewalls, with one or more storage nodes and some clients behind each firewall. In this case, each firewall should be configured the same.

Additional ports that may be required

If ABC Company also wishes to have the Sun StorageTek EBS Console server installed inside the firewall, and allow users to launch the Console GUI from outside the firewall, four additional service ports should be opened through the firewall. The client backup configuration wizard will use a client-side service port, so they may decide to configure an additional service port on each client (for a total of five). NetWorker modules may also require additional ports.

Access control to local host resources

Most resources inherit their administrator list from the NSR resource. The NSR administrator list is only a copy of the list of users in the Administrators entry from the NSR usergroup resource. However, there are a few resources for which this is not true. These exceptions are the resources that you configure using the Local hosts list in the Console Configuration Window.

When selecting a host from the Local hosts list in the Console Configuration Window, you are administering resources on the host that you selected. You are not administering resources on the Sun StorageTek EBS server. It is important to remember this because all other resources in the Console configuration window reside on the Sun StorageTek EBS server. The host on which you configure a resource is the one that you select in the Local hosts tree entry. This selection is made from the left window pane in the Console configuration window.

The Administrator attribute for resources selected from a Local hosts entry is different from the Administrator attribute in the NSR resource. The NSR resource appears when you right click the root of the tree in the Console configuration window and then select Properties.

The user identified as *administrator@gstd_machine* is not automatically included in the Administrator attribute for the these resources. In this case, *gstd_machine* is the hostname of the device where the Console server is installed.

Firewall port enhancement

NetWorker 7.5 uses the sun portmapper on port 111 and its own portmapper (default port 7937 and 7938) for RPC lookups. If you have a firewall that blocks these ports, you may see delays up to four minutes per connection. You can reconfigure NetWorker to use its own portmapper on a different static port number you have opened through your firewall. Setting the portmapper causes it to avoid using the system portmapper.

To enable NetWorker portmapper:

1. Add the following lines to your services file:

```
nsrrpc <portnumber>/udp
nsrrpc <portnumber>/tcp
```

2. Restart the NetWorker services.

Configure the server, its storage nodes, and its clients in the same way.

C

Backing Up and Restoring a Microsoft DFS

This appendix covers these topics:

- Overview of a Microsoft Distributed File System (DFS) 992
- DFS backups and restores 993

Overview of a Microsoft Distributed File System (DFS)

Microsoft distributed file system (DFS) is a Windows file system feature that allows 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 versus registry-based DFS

Windows Server 2003, Windows XP Professional, and Windows 2000 support two types of DFS:

- Domain-based DFS was introduced with Windows 2000. The DFS topology information is stored in Active Directory (AD) on the domain controllers. 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 (also called *stand-alone DFS*) is the type of DFS supported in Windows NT 4.0. It is also supported in Windows 2000 and later operating systems for backward compatibility. The DFS topology information is stored in the Windows registry on the DFS host server.

DFS junctions

A DFS junction is a DFS root (namespace for files and DFS links) or a DFS link (connection to a shared file or folder). DFS junctions are file system objects, but are not files or directories. Therefore, the Sun StorageTek EBS 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.

Note: The DFS terminology Microsoft uses with Windows 2000 and later operating systems is different than that used with Windows NT 4.0. For example, in Windows NT 4.0 DFS junctions are called DFS components.

DFS backups and restores

If the NetWorker Module for Microsoft Applications is installed on the client computer, refer to the *EMC NetWorker Module for Microsoft Applications Administration Guide* for DFS information.

A complete Sun StorageTek EBS backup of a DFS requires backups of the following:

- DFS topology information
- DFS junctions
- DFS destination directories (shared directories connected to DFS links)

Note: The Sun StorageTek EBS 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.

DFS topology information

Domain-based DFS topology information is backed up as part of AD, which is a component of the SYSTEM STATE or VSS SYSTEM BOOT 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 SYSTEM STATE or VSS SYSTEM BOOT save set. Appendix A, "SYSTEM and VSS SYSTEM Save Sets" provides more information about the SYSTEM or VSS SYSTEM save sets.

DFS backup consideration

If VSS is enabled, a DFS backup or a VSS SYSTEM BOOT backup may fail under the following conditions:

- The namespace folder is not a shared folder target.

- Replication is configured for the namespace.
- Files exist in the namespace folder and have replicated to the member server.

To avoid this issue, create a namespace in the folder target or do not enable replication.

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.

To configure a scheduled backup for a DFS:

1. From the **Administration** window, include the following clients in the Sun StorageTek EBS group that will be 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 Sun StorageTek EBS 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 these save sets in the **Save Set** attribute for 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.
   - For registry-based DFS only, include the following SYSTEM save sets in the DFS host server's **Save Set** attribute:

     SYSTEM STATE:

SYSTEM FILES:

If VSS is licensed and enabled, include these VSS SYSTEM save sets:

VSS SYSTEM BOOT:

VSS SYSTEM FILESET:

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 VSS SYSTEM save sets:

   VSS SYSTEM BOOT:

   VSS SYSTEM FILESET:

Restoring a DFS

Note: 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 and VSS SYSTEM 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 13, "Recovering Data" provides recovery procedures.

2. On the DFS host server:
   - 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.

   - If necessary, restore any local DFS destination directories.
3. If necessary, restore the remote DFS destination directories.

D

Additional Features of the Microsoft Windows Server

This appendix covers these topics:

- NetWorker Module for Microsoft Applications 998
- Active Directory 998
- Encrypting file system 1004
- Event logs 1006
- Internet information server 1006
- Windows registry 1007
- Sparse files 1008
- Terminal server licensing 1008
- Windows Change Journal 1010
- Advanced Configuration and Power Interface 1014

NetWorker Module for Microsoft Applications

NetWorker Module for Microsoft Applications (NMM) was previously known as NetWorker VSS Client for Microsoft Windows Server 2003. 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 Office SharePoint Services.

If NMM is installed on the client computer, refer to the *NetWorker Module for Microsoft Applications Administration Guide* for documentation about that product.

There are also separate modules available for Microsoft Exchange Server and Microsoft SQL Server:

- **EMC NetWorker Module for Microsoft Exchange** to back up and recover Microsoft Exchange Server.
- **EMC NetWorker Module for Microsoft SQL Server** to back up and recover Microsoft SQL Server.

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 2003 and Windows 2000 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 Sun StorageTek EBS 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

To recover AD, Windows requires that the system state be recovered. Each backup of the system state includes an expiry token for the AD backup. During recovery, the operating system compares the expiry token to the tombstone lifetime of the domain. An AD recovery fails if the expiry token is older than the tombstone lifetime. (This is a Microsoft AD restriction. To recover AD data that is past the tombstone lifetime, reinstall the operating system on the domain controller *before* recovering.)

The tombstone lifetime is a configurable attribute of a Windows domain. By default, the tombstone lifetime is 60 days. The Microsoft Windows Server Resource Kit documentation on AD provides more information.

Preparing to recover the system state on a domain controller

To prepare to recover the system state on a domain controller:

1. If the domain controller is a DNS server, perform one of the following steps:
   - Edit the following files to include the name and IP address of the Sun StorageTek EBS server (the backup server for the computer) and the primary deduplication node where the system save sets were previously saved to:
     - %SystemRoot%\system32\drivers\etc\hosts
     - %SystemRoot%\system32\drivers\etc\lmhosts
   - Type the name and IP address of the replication node if applicable.

   Note: If the deduplication nodes are not added to the hosts files, the recovery will fail because the deduplication nodes cannot be resolved.

2. Before you start a recovery of the SYSTEM STATE or VSS SYSTEM BOOT save set on a Windows Server 2003 or Windows 2000 domain controller, boot the system in Directory Services Restore Mode. (During a disaster recovery, however, it is not necessary to boot in Directory Services Restore Mode. The *EMC NetWorker Disaster Recovery Guide provides* more information.

Note: If you recover the SYSTEM STATE or VSS SYSTEM BOOT save set on a functioning Windows Server 2003 or Windows 2000 domain controller without first booting the system in Directory Services Restore Mode, recovery of the AD files will fail.

To boot the domain controller in Directory Services Restore Mode, do one of the following:

- Select the operating system to boot, but press F8 before the system begins booting.
- Select the Directory Services Restore Mode as the boot option. The Microsoft Windows documentation provides information about modifying the boot.ini file for multiple-boot configurations.

Performing a nonauthoritative restore

A nonauthoritative restore is the preferred and default method for restoring Active Directory. Nonauthoritative restores rely on at least one other domain controller in the domain to receive updates. The procedure consists of restoring Active Directory on a domain controller and then replicating the most recent data from another domain controller.

To perform a nonauthoritative restore:

1. On the domain controller, close all programs.
2. From the Windows **Start** menu, click **Shut Down**.
3. Select **Restart** and then click **OK**. When the computer restarts, a list of startup choices appears. The startup choices and duration of the startup display are based on the settings in the operating systems section and the boot loader section of the boot.ini file. "Preparing to recover the system state on a domain controller" on page 999 provides more information.
4. Choose the Windows boot option for the domain controller.
5. Press **F8** to display a list of special boot options.
6. Select **Directory Services Restore Mode** (Windows Domain Controllers Only) from the list of special boot options. When you reboot in this mode, AD is taken offline.

7. Log in as administrator by using the password specified when you created the domain controller. Windows starts in **Safe Mode**.
8. Recover the save sets.
   - To initiate the operation from the command-prompt, recover the save sets in the order listed:
     – SYSTEM FILES or VSS SYSTEM FILESET
     – SYSTEM STATE or VSS SYSTEM BOOT

     The command used to recover each save set should look similar to the following:

     ```
     Sun StorageTek EBS_install_path\bin\recover.exe -iY
     [-s Sun StorageTek EBS_server_name] -N
     "saveset_name"
     ```

     Note: To avoid use of the Windows version of **recover.exe**, either include <Sun StorageTek EBS_install_path>\bin\recover.exe at the command-prompt, or ensure that <Sun StorageTek EBS_install_path>\bin is listed before %SystemRoot%\System32 in the $PATH environment variable.

   - To initiate the operation from the NetWorker User program's Recover window, mark both the SYSTEM STATE and SYSTEM FILES (or VSS SYSTEM BOOT and VSS SYSTEM FILESET) save sets in the **Recover** window, and click **Start**.
9. When the restore is finished, verify that the operation was successful.
10. Restart the domain controller computer in normal mode (not in Directory Services Restore Mode). The AD recovery is now complete.

Performing an authoritative restore

If the domain was configured with multiple domain controllers, perform an authoritative restore to replicate the recover AD data to all domain controllers in the domain. An authoritative restore recovers any Windows objects that were corrupted or deleted or modified. After the authoritative restore, the domain administrator can delete any newer objects that are not needed.

AD objects may have associated group policies, for example, organizational units, domains, and site objects. Group policies are

stored in the SYSVOL directory. A restore of the SYSVOL directory cannot be separated from an authoritative restore of AD.

How to perform an authoritative restore

To perform an authoritative restore:

1. Follow step 1 through step 9 in "Performing a nonauthoritative restore" on page 1000.
2. Run the Windows **ntdsutil** command-line utility.

   The **ntdsutil** prompt appears. The **ntdsutil** utility is a command interface similar to the Sun StorageTek EBS **recover** interface. For help in the **ntdsutil** utility, type **?** at the **ntdsutil** prompt.
3. At the **ntdsutil** prompt, type:

   ```
   NTDSUTIL: activate instance ntds (On Windows 2008 systems only)

   NTDSUTIL: authoritative restore
   ```
4. To perform an authoritative restore:
   - Of the entire database, type:

     ```
     NTDSUTIL: restore database (On Windows 2003 systems only)
     ```
   - 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 AD provides information.
5. Exit the **ntdsutil** utility by typing **quit** at each successive **ntdsutil** prompt until the command-prompt appears.
6. Restart the domain controller in normal mode and log in.

Checking authoritative restore results

To ensure that authoritative restore is working properly, perform the procedures in this section with Windows domain administrator privileges.

Prepare to check authoritative restore results

To check authoritative restore results:

1. Select two networked computers, for example, *node_a* and *node_b*.
2. On each computer, perform a fresh installation of the operating system. When you perform the installation, ensure that the hard disk on each computer is reformatted.
3. Set up a Windows domain with two domain controllers.
   a. Use the Windows **dcpromo** command to configure *node_a* as a domain controller.
   b. Add *node_b* to the domain.
   c. Use the **dcpromo** command to configure *node_b* as another domain controller for that domain.
4. Install the Sun StorageTek EBS client software on both nodes. If either node is to function as a Sun StorageTek EBS server, install the server software as well.
5. Back up either one of the domain controllers. For this example, back up *node_a*.
6. After the backup completes, test the authoritative restore that is about to be performed, as follows:
   a. Enter some changes (from the Active Directory Users and Computers snap-in) on *node_a*, the backed-up domain controller. For example, delete an existing user, add one or more new users, and modify an existing user's **Description** field. Record the settings before and after the change.
   b. Wait for the changes to be replicated to *node_b*. The changes are replicated when the settings on *node_a* appear on *node_b*.
7. For an authoritative restore, perform the steps in "How to perform an authoritative restore" on page 1002.
8. Ensure that the authoritative restore worked correctly, as described in "How to check the authoritative restore results" on page 1004.

How to check the authoritative restore results

To check the authoritative restore results:

1. Open the **Active Directory Users and Computers MMC** snap-in.
2. If the authoritative restore was successful, these settings appear on both node_a and node_b after the reboot and a complete replication cycle:
   - The modified user description now has its original description.
   - The deleted user is present again.
   - The users that were added after the backup are still present.

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:

- Sun StorageTek EBS 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.
- Sun StorageTek EBS 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.

Event logs

Event logs can be used for troubleshooting hardware problems as well as monitoring security conditions, and system and application software problems.

If no VSS client license exists, or if VSS is disabled, the Sun StorageTek EBS 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 licensed and enabled, event logs are backed as a component of the VSS SYSTEM SERVICES save set.

Note: 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) 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 Sun StorageTek EBS 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.

Note: If you perform a Sun StorageTek EBS 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

In Windows Server 2003, Windows XP Professional, and Windows 2000, 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 Sun StorageTek EBS software automatically backs up or recovers the registry as well as this save set.

The registry is always saved and restored at level full.

Backward compatibility with earlier Sun StorageTek EBS releases

The Windows NT registry save set is no longer supported in the Sun StorageTek EBS software. In Windows Server 2003, Windows XP Professional, and Windows 2000, the registry is part of the system state. The Sun StorageTek EBS software includes the registry as a component of the SYSTEM STATE or VSS SYSTEM BOOT save set. Specifying registry as a save set in a Client resource does *not* result in a backup of the Windows Server 2003, Windows XP Professional, or Windows 2000 registry. Backup and recovery of the All save set includes the SYSTEM STATE or VSS SYSTEM BOOT save set, which includes the registry.

Note: Attempting to restore a Windows NT 4.0 registry save set to a computer that is running Windows Server 2003, Windows XP Professional, or Windows 2000 will yield unpredictable results.

Files and folders not backed up in Windows 2000

Windows 2000 excludes some files and folders from backup. The files and folders that are excluded are listed in the following registry keys:

- HKEY_LOCAL_MACHINES\SYSTEM\CurrentControlSet\Control\BackupRestore\FilesNotToBackup
- HKEY_LOCAL_MACHINES\SYSTEM\CurrentControlSet\Control\BackupRestore\KeysNotToRestore

The Microsoft Knowledge Base article 233427 on the Microsoft website provides more information.

Sparse files

The NTFS sparse files feature enables a program to create huge files without actually committing disk space for every byte. The Sun StorageTek EBS software provides complete backup and recovery support for sparse files.

Terminal server licensing

Terminal Services provides access to Windows 2000 applications running on a server host. TSL (Terminal Server Licensing) is the client license management service for Terminal Services Windows 2000 Server edition.

Backing up TSL data

TSL is a component of the SYSTEM DB save set. To back up TSL, back up both of the following:

- SYSTEM DB save set
- Directory in which the license server is installed

This ensures that data in both the registry and the license server database is saved.

To help protect licensing data from accidental loss:

1. Back up the SYSTEM STATE save set so that the registry is backed up. If the TSL database changes and the SYSTEM STATE is *not* backed up before the next SYSTEM STATE and SYSTEM DB save set recovery.
2. Edit the Windows 2000 registry to reflect the TSL changes that were not backed up.
3. Back up the SYSTEM FILES save set whenever other SYSTEM save sets are backed up.

Recovering TSL data

If you restore the system state data and license server database to the same license server host, any unissued licenses are restored correctly as long as you have not replaced the operating system on the computer. If you do not restore to the same computer, any unissued licenses that are detected are not restored and an event appears in the system log of the Event Viewer. To restore the unissued licenses, contact Microsoft and request that the licenses be reissued.

You can recover a former state of TSL:

- After it is removed or lost from a computer.
- After the operating system is reinstalled.

Recovering the TSL former state

To recover a former state of TSL:

1. If TSL is *not* currently installed on the computer, reinstall it to the same location it occupied before.
2. Ensure that the TSL service is started.
   a. Register a new TSL installation with Microsoft.
   b. Activate the service by typing the number that Microsoft provides.
3. From the NetWorker **User** program, click **Recover** to open the **Source Client** dialog box.
4. Select the source client with the data to recover, and click **OK**.
5. Select the destination client for the recovered data, and click **OK**.

6. Mark the **SYSTEM FILES** and **SYSTEM STATE** save sets for recovery to a point in time that closely precedes the time TSL was moved or changed.
7. Click **Start**.
8. When the recovery operation is finished, reboot the computer.
9. Restore the **SYSTEM DB** save set. The TSL database is recovered as part of this save set.
10. Reboot the computer again to activate the Terminal Service licensing.

Windows Change Journal

Microsoft Windows Change Journal is a Windows Server 2003, Windows XP Professional, and Windows 2000 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 Sun StorageTek EBS 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 Sun StorageTek EBS 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.

Sun StorageTek EBS support for Change Journal

Windows does not have an administrative interface for enabling or disabling the Change Journal. That functionality is provided with the Sun StorageTek EBS Change Journal Manager.

The Sun StorageTek EBS Change Journal Manager is installed during the Sun StorageTek EBS software setup and can be run from **Start>Programs>Sun StorageTek EBS**.

The Sun StorageTek EBS Change Journal Manager allows you to:

- Enable or disable the Change Journal for each NTFS 5.0 volume.
- Enable or disable the Sun StorageTek EBS software's use of each volume's Change Journal.
- Set parameters that control the size of the Change Journal log file.

How Sun StorageTek EBS software uses the Change Journal

When configured to use the Change Journal, the Sun StorageTek EBS 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 190 provides information about level and incremental backups.

Note: The Sun StorageTek EBS 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 Sun StorageTek EBS **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 Sun StorageTek EBS **save** command's **-t** argument. The *EMC NetWorker Command Reference Guide* or the UNIX man page provides information about the Sun StorageTek EBS **save** command.

Sun StorageTek EBS save criteria when using the Change Journal

A Sun StorageTek EBS 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 Sun StorageTek EBS 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 127 shows the Sun StorageTek EBS save triggers when using the Change Journal.

Table 127 Sun StorageTek EBS 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 Sun StorageTek EBS software to use the Change Journal

The Sun StorageTek EBS Change Journal Manager enables one to view or edit the Change Journal configuration of each volume in the Sun StorageTek EBS server or client host computer.

How to configure Sun StorageTek EBS software to use the Change Journal

To view or edit the Change Journal configuration:

1. On the Sun StorageTek EBS server or client host computer, select **Start>Programs>Sun StorageTek EBS>Sun StorageTek EBS Change Journal Manager**.

2. In the **Sun StorageTek EBS 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 Sun StorageTek EBS software, ensure that the **Sun StorageTek EBS Uses Change Journal** checkbox is selected. To disable the Change Journal, ensure that this checkbox is clear.

   If you select **All NTFS Volume**s, 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 **Sun StorageTek EBS Change Journal Manager**.

Advanced Configuration and Power Interface

Sun StorageTek EBS software supports the following:

- Windows Server 2003
- Windows XP Professional
- Windows 2000 Microsoft Advanced Configuration
- Power Interface (ACPI), which is also called OnNow

Sun StorageTek EBS support for ACPI

Support for ACPI is provided by the *Sun StorageTek EBS Power Monitor* service. The executable for this service (<Sun StorageTek EBS_install_path>\bin\nsrpm.exe) is installed and configured for automatic startup during Sun StorageTek EBS setup.

For scheduled backups of Windows Server 2003, Windows XP Professional, and Windows 2000 Sun StorageTek EBS clients, each client must be able to respond when a Sun StorageTek EBS server contacts it for backup. Therefore, by default, the Sun StorageTek EBS Power Monitor service does not allow a Sun StorageTek EBS server or client host operating system and its network interface to enter the ACPI standby mode.

However, the Sun StorageTek EBS 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 Sun StorageTek EBS software

does not prevent the host's power-management policies from forcing the system to power down. The Sun StorageTek EBS 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 Sun StorageTek EBS 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 Sun StorageTek EBS client or server host in standby or hibernation mode *while* a Sun StorageTek EBS backup or restore operation (involving that host) is in progress. Doing so will yield unpredictable results.
- If a Sun StorageTek EBS server is powered down while Sun StorageTek EBS 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 Sun StorageTek EBS's tape processes will have incorrect positioning information.

E

UNIX and Linux Platform-Specific Notes

This appendix covers these topics:

- Solaris 1018
- Linux 1021
- HP-UX 1023
- AIX 1029
- IRIX 1030

Solaris

This section provides information specific to Sun StorageTek EBS software that runs on the Solaris platform.

Support for Solaris zones

The Sun StorageTek EBS software provides support for local zones for a client and a dedicated storage node only. You can install and back up a Sun StorageTek EBS client or storage node on a machine running in a local zone. "Dedicated storage nodes" on page 131 provides more information about storage node support in a local zone.

The NetWorker Console and the NetWorker server can only be run in a global zone.

"Using the All keyword for the Save Set attribute on Solaris 10" has information about the "All" keyword for the Save Set attribute in the Client resource with Solaris zones.

Using the All keyword for the Save Set attribute on Solaris 10

The All keyword for the Save Set attribute in the Client resource backs up all locally-mounted filesystems listed in the /etc/vstab file. Any file systems other than those found in the vfstab file are ignored.

Using the All keyword with Solaris zones

Because the All keyword backs up only locally-mounted file systems listed in the /etc/vstab file, it is possible that this configuration will result in no files being backed up on a Sun StorageTek EBS client in a local zone. This may occur for two reasons:

- The root directory of a local zone may be a directory off the global zone's root, and not a mount point for a device.
- The root directory of a local zone is a mount point for a device, but it is not listed in the vstab file.

ZFS file systems and the All keyword

ZFS filesystems are ignored when the **All** keyword is used for the Save Set attribute in the Client resource, because ZFS filesystems are not listed in the vstab file.

To back up the ZFS filesystems:

1. Explicitly set the file systems for backup in the client's save set list.
2. Set the ZFS file systems to legacy mount point behavior.

   For example:

   ```
   # zfs umount <zpool>
   # zfs set mountpoint=legacy <zpool>
   ```
3. Add an entry to the /etc/vfstab.

   For example:

   ```
   <zpool> - /mypool zfs - yes -
   ```
   where /mypool is the mount point

   ```
   # mount /mypool
   ```

Sun StorageTek EBS executables not found for Solaris client

On Solaris, Sun StorageTek EBS executables are installed by default in /usr/sbin/nsr. If you start a group backup on a Sun StorageTek EBS server that does not have /usr/sbin/nsr in the search path for root, the backup fails on a client that has its Sun StorageTek EBS executables in /usr/sbin/nsr. 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/nsr.
6. Click **OK**.

Alternate solution: Modify the search path for root on the Sun StorageTek EBS server to include /usr/sbin/nsr even if it does not exist locally.

Extended file attribute data included in Save Set File Size attribute

The save set file size shown in Sun StorageTek EBS 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 Sun StorageTek EBS.

Linux

This section provides information about Sun StorageTek EBS 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.

Configuration requirements for the inquire command

Depending on the specific OS requirements, and the configuration of the Sun StorageTek EBS server or storage node, device files may need to be created so that the **inquire** command can detect all devices.

For example, on a Sun StorageTek EBS 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

Note: 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 6, "Directives" provides information about directives.

HP-UX

This section provides information specific to Sun StorageTek EBS storage node software 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 Sun StorageTek EBS 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.

Note: 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 1024.
   - 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 1024.

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 *majornum*, 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 *majornum* is the number in the Character column.

Minor number

To determine the value for *minornum*, 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
```

This is an example of the output from this command (with the **-s** option):

```
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
```

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.

The "Hardware address not found" message

The HP-UX system configuration logger may generate the error message:

`bootpd` *`pid#`* `Hardware address not found` *`hardware-address`*

A similar message is written to the Sun StorageTek EBS /nsr/logs/messages file. If this error message appears:

1. Log in as root and edit the /etc/syslog.conf file to change every instance of daemon.notice to local7.notice.
2. Edit the **Log Default** notification in the **Notification** resource to change the value for the **Action** attribute:
   - Change daemon.notice to local 7.notice.
3. Apply and save the changes to the **Log Default** notification.
4. Enter this command:

   ```
   cat /etc/syslog.pid
   ```

5. Use the PID number obtained from the /etc/syslog.pid file to kill the designated PID number and cause the **syslogd** daemon to reread its configuration file:

   `kill -HPU` *`pid#`*

   The local7 facility is provided as an example.
6. View the syslog.h system header file to determine other possibilities for the facility.

- By default, LOG_LOCAL0 to LOG_LOCAL7 are reserved for local use.
- If they are not used by other local utilities on the system, the Sun StorageTek EBS server can use them to avoid the hardware address problems encountered with **bootpd**.

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 storage node

When the **nsrjb -u -S** command is used to unload a tape drive in an autochanger attached to an HP-UX 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.

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 Sun StorageTek EBS server will *not* back up the file system that the symbolic link points to.

Font warnings when exporting display to a Solaris computer

If the Sun StorageTek EBS motif-based graphical interface (**nwrecover**) is being run on an HP-UX computer and exporting the display to a Solaris computer, the motif application may fail to open and the following warnings may appear:

- Warning: Missing charsets in String to FontSet conversion
- Warning: Cannot convert string *string* to type FontSet

If this occurs, set the LC_CTYPE environment variable on the HP-UX computer to the value C.iso88591. For example, from an xterm window, telnet to the HP-UX computer and run the following command: **export LC_CTYPE=C.iso88591**

The **nwrecover** should work correctly when this environment variable is set. Add this environment variable setting to your .profile on the HP-UX computer.

AIX

This section provides information specific to Sun StorageTek EBS storage node 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.

IRIX

This section provides information specific to Sun StorageTek EBS storage node software running on the IRIX platform.

IRIX displays a “no space left in tape” message

If an IRIX Sun StorageTek EBS storage node is used, the following message may appear if the tape drive is not configured within the machine’s kernel:

```
BSF invalid argument no space left in tape
```

This is not a Sun StorageTek EBS error. Ensure that the device is supported and configured by applying the latest kernel patch from the Silicon Graphics.

Go to the SGI Services and Support website for the required patches and information about the current drives, robotics, and firmware revisions that IRIX supports.

F

MAC OS X Support

This appendix covers these topics:

- Support for Mac OS X..... 1032
- Mac OS X backup considerations 1032
- Recovering files and directories on Mac OS X..... 1035

Support for Mac OS X

This section describes Sun StorageTek EBS client support for the Mac OS X platform. Mac hosts can be set up as Sun StorageTek EBS clients by using the NetWorker client for Mac OS X. Mac hosts that are set up as Sun StorageTek EBS clients can be backed up and restored by using any supported Sun StorageTek EBS server on Solaris. Currently the Sun StorageTek EBS server and the Console server are not supported on Mac OS X.

Mac OS X metadata support

The Sun StorageTek EBS client on Mac OS X supports backup and recovery of all file system metadata including:

- Finder information
- Resource forks
- Extended attributes (Mac OS X 10.4)
- Access Control Lists (Mac OS X 10.4)

Supported file systems

The Sun StorageTek EBS client for Mac OS X software supports these file systems:

- HFS+ (including journaled)
- HFS
- UFS

Known limitations

For the NetWorker client for Mac OS X, there is no graphical user interface, however there is full support of the command line interface.

Mac OS X backup considerations

Use this section to help plan successful backups for Sun StorageTek EBS clients on the Mac OS X platform.

Scheduling a Sun StorageTek EBS client backup on Mac OS X

This section provides information on configuring backups for a Sun StorageTek EBS 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 81 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 230 provides more information about Mac OS directives.
3. Click **OK**.

"Scheduled backups" on page 74 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 Disaster Recovery Guide* provides more information about disaster recovery.

Note: The Sun StorageTek EBS 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 **savepnpc** 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 **savepnpc** in the Backup Command attribute when configuring the Mac OS X client as a Sun StorageTek EBS Client resource.

   "Using the savepnpc command with a customized backup program" on page 116 provides more information about enabling the Client resource to use the **savepnpc** command.

2. Create a custom **savepnpc** 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 **savepnpc** 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 55 Custom savepnpc script for Mac OS X

A Mac OS X Sun StorageTek EBS client that belongs to the Default group will have a /nsr/res/Default.res script with this content:

```
type: savepnpc;
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 **savepnpc** 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 Sun StorageTek EBS server defined in the /nsr/res/servers file that comes first in alphabetical order.

To specify an alternative Sun StorageTek EBS server, use the **save** command with the **-s** *Sun StorageTek EBS_server* option.

Recovering files and directories on Mac OS X

These sections provide information on recovering individual files and directories from a Sun StorageTek EBS client on Mac OS X:

- "Task 1: Browse backed-up Mac OS X data" on page 1035
- "Task 2: Recover individual files or directories" on page 1036

Each task in this section uses the Sun StorageTek EBS **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 Sun StorageTek EBS server by using this command:

   ```
   $ recover
   ```

   By default, the **recover** command contacts the Sun StorageTek EBS server defined in the /nsr/res/servers file that comes first in alphabetical order. To specify an alternative Sun StorageTek EBS server, use the **recover** command with the **-s** *Sun StorageTek EBS_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
   ```

Note: 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.

G

Direct SCSI Backup and Recover

This appendix covers these topics:

- Introduction to direct SCSI backup and recover 1038
- System requirements 1038
- Performing direct SCSI backup 1040
- Performing direct SCSI recover 1044
- Licensing 1047

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 to the Sun StorageTek EBS 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 Sun StorageTek EBS 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.

System requirements

The direct SCSI feature is only supported on Solaris storage nodes, and you must have EMC Solution Enabler version 5.5 or later installed.

The following hardware devices are supported:

- Raw device path of a SCSI device
- Symmetrix device

Note: 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.

Unsupported features

If performing backup and recover over a SCSI bus, the following features are not supported with Sun StorageTek EBS release 7.4:

- 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 Sun StorageTek EBS 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 Sun StorageTek EBS 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 Sun StorageTek EBS server.
- Runs **scsi asm** on the raw device path to move data from the SCSI device to a storage node (**nsrmmd**) by using SCSI commands.



IMPORTANT

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 Sun StorageTek EBS Client resource. "Task 6: Create a backup Client resource" on page 81 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 Sun StorageTek EBS server.

The Sun StorageTek EBS 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:

7. For the AppHost or the host that controls the data, perform either of the following steps:
   - Create a new Sun StorageTek EBS Client resource. "Task 6: Create a backup Client resource" on page 81 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**.
8. 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/rdsk*, type the following:

   ```
   /dev/rdsk/c1t2d0s2
   ```

9. On the **Apps and Modules** tab, select **SCSI** for the **Proxy backup type** attribute.
10. In the **Proxy Backup host** attribute, enter the name of the storage node on which the nsrscsi_save command will be run.
11. 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 Sun StorageTek EBS 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/rdsk/c1t2d0s2
  - For a device-set-name: {OracleDevices}

  Note: The braces are mandatory 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 Sun StorageTek EBS 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 Sun StorageTek EBS server.
- Runs **scsi asm** on the raw device path to move data from the storage node (nsrmmd) to a raw device via SCSI CDB commands.



IMPORTANT

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 Sun StorageTek EBS 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
oracledisksetscip2b081.legato.com.001
oracledisksetscip2b081.legato.com.001
oracledisksetscip2b081.legato.com.001
oracledisksetscip2b081.legato.com.001
```

```
scip2b081.legato.com.001
oracledisksetㄴ000187910217/0366 11/21/06 06:14:26 PM
3680744065
scip2b081.legato.com.001.RO
oracledisksetㄴ000187910217/0366 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.

   Note: 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 Sun StorageTek EBS 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.legato.com.001 scip2b081.legato.com 8839 MB
full /dev/rdsk/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/rdsk/c1t1d0s2).

   The Sun StorageTek EBS software recovers the contents of the save set to the destination location. The destination location is the original location by default.



IMPORTANT

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.

Glossary

This glossary provides definitions for terms used in this guide.

A

ACL (Access Control List) A list that specifies the permissions assigned to a specific file or directory.

active group A Sun StorageTek EBS backup group that has its Autostart attribute enabled.

Administrators group A 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.

agent A term used by Sun Microsystems to denote a cluster server. Also known as a package (HP-UX), and a virtual server (Microsoft).

annotation

1. A comment associated with an archive save set.
2. A comment associated with an event.

archive Backing up directories or files to an archive volume to free disk space. Archived data is not recyclable. See also "groom."

archive request A Sun StorageTek EBS resource used to schedule and manage archiving.

archive volume A volume used to store archive data. You cannot store archive data on a backup or cloned volume. See also *backup volume* or See "clone volume."

ASM (Application Specific Module) A 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 Sun StorageTek EBS directive used to compress files.

ATMOS EMC's cloud backup offering. ATMOS can be configured as an On Premise option (onsite) or as an Online option (offsite).

attribute A feature, such as an entry field, of a resource. See also "resource."

authentication The process by which a user or software process is determined to be trusted or not trusted.

authorization Privileges assigned to users of the Console and NetWorker server.

authorization code A unique code that unlocks the Sun StorageTek EBS software for permanent use.

autochanger See "library."

autochanger sharing See "library sharing."

Auto Media Management A feature that enables the storage device to automatically label, mount, and overwrite an unlabeled or recyclable volume.

B

backup An operation that saves data to a volume.

backup cycle The period of time from one level full backup to the next level full backup.

backup group See "group."

backup level See "level."

Backup Operators group A Microsoft Windows user group that can log in to a domain from a computer or a server, and back up and restore its data. Backup operators can also shut down servers or computers.

backup volume A volume used to store backup data. You cannot store backup data on an archive or cloned volume. See "archive volume." or See "clone volume."

base enabler code See “enabler code.”

bootstrap A save set that is essential for Sun StorageTek EBS disaster recovery. A bootstrap is composed of two components that reside on the Sun StorageTek EBS server: the media database, and the resource database. The bootstrap also provides information that is essential for recovering client file indexes. See “media database.” See “resource database.” See “client file index.”

browse policy A policy that determines how long backed up data is available for quick access. See “retention policy.”

C

canned report A preconfigured report that can be tailored by the user.

carousel See “library.”

client A computer, workstation, or fileserver whose data can be backed up or recovered. See “Client resource.”

client file index A database that tracks every database object, file, or file system that is backed up. The Sun StorageTek EBS server maintains a single client index file for each client. See “file system.”

client-initiated backup See “manual backup.”

Client resource 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. Multiple Client resources can be configured for one client computer. For example, you could create one Client resource to back up business data and another to back up operating system files. See also “client.”

clone A reliable copy of backed up data. Unlike volumes created with a simple copy command, clone volumes can be used in exactly the same way as the original backup volume. Single save sets or entire volumes can be cloned.

clone volume A volume used to store clone data. You cannot store clone data on either an archive or a backup volume. See “archive volume.” See “backup volume.” .

cloud For NetWorker, a cloud refers to an ATMOS On Premise option (onsite) or an ATMOS Online option (offsite) configuration of backup disks. NetWorker supports the notion of a cloud backup device. However, no specific knowledge or management of the backup disks behind a cloud is required. NetWorker only requires that the user specify credentials for a valid ATMOS account when setting up a cloud backup device.

cluster Two or more nodes that are connected and appear to network users as a single, highly available system.

command-prompt The shell prompt, where you enter commands.

CIF (Common Internet File System) Formerly known as Server Message Block (SMB). A message format used by DOS and Windows to share files, directories, and devices.

connection port A port that Sun StorageTek EBS processes use to perform backup and recovery sessions through a firewall.

Console Application Administrator A Console server user role whose members can configure features (except security) in the Console sever application.

Console Security Administrator A Console server user role whose members can add Console users and assign them to Console roles.

Console Server NetWorker servers and clients are managed from the NetWorker Console server. The Console server also provides reporting and monitoring capabilities for all NetWorker servers and clients.

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 A group of datazones managed by the Sun StorageTek EBS software.

D

daemon A process that runs in the background and performs a specified operation at predefined times or in response to certain events.

data retention policy See “retention policy.”

datawheel See “library.”

datazone A group of computers administered by a Sun StorageTek EBS server.

DDS See "DDS (Dynamic Drive Sharing)."

device
1. A storage unit that reads from and writes to backup volumes. A storage unit can be a tape device, optical drive, "autochanger", or file connected to the server or storage node. See "library." See "server." See "storage node."
2. When dynamic drive sharing (DDS) is enabled, refers to the access path to the physical drive.

Device Central The interface from which one can manage all Sun StorageTek EBS libraries.

DFS (Distributed File System) A Microsoft Windows add-on that allows you to create a logical directory of shared directories that span multiple machines across a network.

DFS component
1. A namespace for files and DFS links, called a DFS root.
2. A connection to a shared file or folder, called a DFS child node

DAR (Direct Access Restore) An NDMP data recovery operation that allows you to recover data in the middle of a tape set without having to parse the tape set sequentially. This reduces recovery time of large backups.

directed recovery A recovery method used to recover data that originated on one computer to another computer.

directive Instructions to take special actions on a given set of files for a specified client during a backup.

DMA (Data Management Application) Initiates a backup or recovery request through the NDMP connection.

document mode A mode that displays static reports (as charts or tables) in a format that resembles the Print Preview view displayed by a PDF file viewer.

domain controller A computer that manages user interactions with a domain, including logon, authentication, directory searches, and access to other shared resources.

drill-down reports Organizes basic report data in terms of granularity. For example, a user could start with group summary data, then "drill down" to data for a client within the group, and then to data for a selected save set.

drive When dynamic drive sharing (DDS) is enabled, refers to the physical backup object, such as a tape drive, disk, or file. See “device.”

DSA (Data Server Agent) An agent save or agent recover between a Sun StorageTek EBS server and a non-Sun StorageTek EBS client. An example of a DSA is an NDMP host that generates proprietary save data and sends that data to a Sun StorageTek EBS storage device to have a save set associated to it.

DSA save set Save sets of an NDMP client that are backed up to non-NDMP tape device.

DSP (Data Service Provider) Controls access to disk storage during an NDMP back up.

DDS (Dynamic Drive Sharing) A feature that allows Sun StorageTek EBS software to recognize shared drives.

E

enabler code A code that allow you to run Sun StorageTek EBS software.

enterprise Computers and folders organized into a tree-based visual representation.

event A notification of an application event. This notification contains information that may require user action.

exit code An indicator that specifies whether a backup or recovery session succeeded. An exit code of zero (0) indicates the session completed successfully. An exit code other than zero indicates that the session did not complete successfully.

expiration date The date when a volume changes from read/write to read-only.

expired save set A save set that has reached its browse time and is no longer browsable.

F

file index See "client file index."

file system
1. A file tree located on a specific disk partition or other mount point.
2. The entire set of all files.
3. A method of storing files.

firewall A system designed to prevent unauthorized access to or from a private network.

folder The graphical representation of a branching node in the Enterprise.

full backup See "level."

G

groom To remove the original files from a disk after an archive.

group A client or group of clients configured to back up files to the Sun StorageTek EBS server at a designated time of day.

GST (Generic Services Toolkit) Provides the underlying framework of the Console server.

H

heterogeneous networks Computers of different platforms that interact across the network.

highly available system A system that allows cluster-aware services to continue despite a hardware or software failure.

high-water mark The percentage of disk space that when filled, automatically starts the staging process.

host authentication A feature that provides encryption and verification services between NetWorker hosts. See also "user authentication."

host ID A serial number that uniquely identifies a computer.

I

inactivity timeout The number of minutes to wait before a client is considered to be unavailable for backup.

incremental See “level.”

individual user authentication The process by which Console administrators restrict or grant user access to Sun StorageTek EBS servers, based on Console usernames.

insertion time The time that the save set record was most recently introduced into the save set database.

Interactive mode A 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.

J

JAR (Java Archive) A file that contains compressed components needed for a Java applet or application.

Java A high-level programming language. The Java Virtual machine (JVM) enables the same, unmodified, Java program to run on most computer operating systems. See “JVM (Java Virtual Machine).”

Java Plug-in A JVM that can be used by a web browser to run Java applets.

jukebox See “library.”

JVM (Java Virtual Machine) An execution environment for interpreting the Java programming language. Each operating system runs a unique JVM to interpret Java code.

L

label A Sun StorageTek EBS assigned label that uniquely identifies a volume. Templates can be used to define label parameters.

LDAP (Lightweight Directory Access Protocol) A set of protocols for accessing information directories. LDAP is based on the X.500 standard, but is significantly simpler.

legacy method The use of special-case Microsoft APIs to back up and recover operating system components, services, and applications.

level A measurement that determines how much data is saved during a scheduled or ad hoc backup.

A full (f) backup backs up all files, regardless of whether they have changed. Levels one through nine [1-9] back up files that have changed since the last lower numbered backup level. An incremental (incr) backup backs up only files that have changed since the last backup.

library A mechanism that uses a robotic arm to move media among various components in a device. 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 tape drives in a library.

license enabler A code that is required to run a feature or product.

LLM (License Manager) An application that provides a central point for managing product licenses.

local cluster client A Sun StorageTek EBS 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.

locale settings Settings that specify the input and output formats for date and time, based on local language conventions.

logical cluster client See “virtual cluster client.”

logical device A specific Sun StorageTek EBS device used in the integration of Sun StorageTek EBS software with SmartMedia or AlphaStor. SmartMedia or AlphaStor assigns the logical device to an actual device.

low-water mark The percentage of disk space filled that, when reached, automatically stops the staging process.

LUS (EMC User SCSI) A proprietary device driver that sends arbitrary SCSI commands to an autochanger.

M

managed application A program that can be monitored and/or administered from the Console server.

manual backup A backup that a user performs from the client, also known as an unscheduled backup. The user specifies the files, file systems, and directories to back up.

media The physical storage medium, such as magnetic tape, optical disk, or file system to which backup data is written.

media database Indexed entries about the location and the life cycle status of all data and volumes that the Sun StorageTek EBS server manages. See "volume."

media index See "media database."

member A physical host in a cluster environment. Each member has its own IP address. A member can also be a "virtual server".

metadata VSS-defined information that is passed from the writer to the requestor. Metadata includes the writer name, a list of VSS components to back up, a list of components to exclude from the backup, and the methods to use for recovery. See "writer." See "requestor." See "VSS component."

mount To make a volume physically available so that backup data can be written to it.

mount point See "volume mount point."

multiplexing To simultaneously write data from more than one save set to the same storage device.

N

NDMP (Network Data Management Protocol) A TCP/IP-based protocol that specifies how heterogeneous network components communicate for the purposes of backup and recovery.

NDMP data server An NDMP service that either reads from or writes to disk.

NDMP host The host computer that executes the NDMP server application.

NDMP save set A save set of an NDMP client that is backed up to an NDMP tape device.

NDMP server An instance of one or more NDMP services controlled by a single NDMP control connection. Thus a data/tape/SCSI server is an NDMP server providing data, tape, and SCSI services.

NDMP service The virtual state machine on an NDMP host that is controlled by the DMA. Examples of NDMP services include:

- A general server with direct attached storage (storage appliance)
- A system with one or more tape drives
- A software process that reads two datastreams and multiplexes them into one stream

NDMP storage node A host having NDMP services or any open system having NDMP services installed. For instance, Netapp Filer and EMC Filer.

near-line storage See “library.”

NFS (Network File System) A client/server application that allows users to access shared files stored on networked computers of different types.

Sun StorageTek EBS client See “client.”

Sun StorageTek EBS Console server See “Console Server.”

Sun StorageTek EBS Management Console See “Console Server.”

Sun StorageTek EBS server The host running the Sun StorageTek EBS server software, which contains the online indexes and provides backup and recovery services to the clients on the same network. See also “online indexes.”

NetWorker Administrator A default NetWorker server user group that can add, change, or delete NetWorker server user groups.

Sun StorageTek EBS storage node See “storage node.”

NFS server A computer that contains exported file systems that NFS clients can access.

node See "member."

nonclone pool Pools that contain data that has not been cloned.

non-NDMP device A device that is not configured as an NDMP type of device.

notification A message sent to the Sun StorageTek EBS administrator about important Sun StorageTek EBS events.

nsrd The master Sun StorageTek EBS server process.

nsrhost The logical hostname of the Sun StorageTek EBS server.

O

online indexes Databases on the Sun StorageTek EBS server that contain information about client backups and backup volumes. See "client file index." See "media database."

operator The person who monitors the server status, loads backup volumes into storage devices, and executes day-to-day Sun StorageTek EBS tasks.

override A backup level that takes place instead of the scheduled one.

P

package A term used by HP-UX to denote a cluster server. Also known as an agent (Sun) or virtual server (Microsoft).

packet filtering A method of firewall protection that accepts or rejects packets entering or leaving the network based on user-defined rules. See "firewall."

parallelism A method of backing up or recovering data from several clients or save sets in parallel.

pathname Instructions for accessing a file. An absolute pathname indicates how to find a file starting from the root directory. A relative pathname indicates how to find the file starting from the current directory.

peer A Sun StorageTek EBS host that is involved in an authentication process with another Sun StorageTek EBS host.

permanent enabler A code that allows the software to be authorized for permanent use. See “temporary enabler.” See “authorization code.”

physical cluster client A Sun StorageTek EBS client that is bound to a physical host in the cluster and can have its own resources (private or local).

physical host A node or host that forms part of a cluster.

policy A set of constraints that specify how long data is available for recovery. Each Client resource has a browse policy and a retention policy.

pool A feature to sort backup data to selected volumes.

PowerSnap EMC technology that provides point-in-time snapshots of data. Sun StorageTek EBS 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.

probe A query to determine the directories or files to back up on each client.

purging Deleting file entries from the client file index.

R

recover To restore files from a backup volume to a client disk.

recyclable save set A save set whose browse and retention policies have expired. Recyclable save sets are removed from the media database.

recyclable volume A volume whose data has passed both its browse and retention policies and is now available to be relabeled.

Registry A Microsoft Windows database that centralizes all Windows settings and provides security and control over system, security, and user account settings.

remote device A storage device that is attached to a storage node.

RPC (Remote Procedure Call) The protocol that the Sun StorageTek EBS server uses to perform client requests over a network.

repository A Console database that contains configuration and reporting information.

requestor A VSS-aware application that creates and destroys a shadow copy. Sun StorageTek EBS software is a requestor. See "shadow copy."

resource A component that describes the Sun StorageTek EBS server or its clients. Clients, devices, schedules, groups, and policies are all Sun StorageTek EBS resources. Each resource has attributes that define its properties.

resource database A database of information about each resource.

resource owner A 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.

retention policy Determines how long backup data is available for recovery, though not necessarily immediate recovery. See "browse policy."

retrieve To locate and recover archived files and directories.

retry mechanism The action Sun StorageTek EBS software performs when client operations fail. This situation might occur when the rate of transmission is either low or non-existent.

roles Used to assign user privileges to the Console. There are three roles: Console Application Administrator, Console Security administrator, and the Console User. See also "user groups."

root

1. (UNIX only) The UNIX superuser account.
2. (Microsoft Windows and UNIX) The highest level of the system directory structure.

S

save The command that backs up client files and makes entries in the online index.

save set A group of files or a file system that is backed up on storage media.

save set consolidation A 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. See "level."

SSID (Save Set ID) An 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 An attribute that indicates whether a save set is browsable, recoverable, or recyclable. The save set status also indicates whether the save set was successfully backed up.

save stream The data and save set information being written to a storage volume during a backup.

scanner A command used to read a backup volume when the "online indexes" are no longer available.

server See "Sun StorageTek EBS server." See "Console Server."

server index See "client file index".

service port A port used to listen for backup and recover requests from clients through a firewall.

shadow copy A temporary, point-in-time copy of a volume created using VSS technology. See "VSS (Volume Shadow Copy Service)."

shared disk The storage disk that is shared among multiple nodes in a cluster.

shell prompt A cue in a shell window where you enter a command.

silo A repository for holding hundreds or thousands of volumes. Silo volumes are identified by bar codes, not by slot numbers.

skip A backup in which designated files are not backed up. See "level."

snapshot A point-in-time copy of data created during an instant backup.

snapshot policy A Sun StorageTek EBS resource used to configure backups using the Sun StorageTek EBS PowerSnap Module software.

snapup See "backup." Snapup is a term previously used by Avamar to refer to backup.

SNMP (Simple Network Management Protocol) A protocol used to send messages to the administrator about Sun StorageTek EBS events.

staging Moving 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 Sun StorageTek EBS server that starts in noncluster (stand-alone) mode.

stand-alone device A storage device that contains a single drive for backing up data.

STL Silo Tape Library.

storage device The hardware that reads and writes data during backup, recovery, or other Sun StorageTek EBS operations.

storage node A storage device physically attached to another computer whose backup operations are controlled by the Sun StorageTek EBS server.

T

tape service An NDMP DSP service that controls access to tape storage. A system can simultaneously host multiple tape services corresponding to multiple backup streams.

target sessions The number of backup sessions accepted by a backup device.

temporary enabler A code that allows operation of the software for an additional 45 days beyond the evaluation period. See “enabler code” and “permanent enabler.”

TCP/IP (Transmission Control Protocol / Internet Protocol) A standard set of communication protocols that connects hosts on the Internet.

trap Unsolicited notification sent from the SNMP agent to the network manager’s SNMP event manager.

U

update enabler A code that updates software from a previous release. Like other temporary enabler codes, it expires after 45 days.

user

1. A Sun StorageTek EBS 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 authentication

A feature that validates user sign-on attempts. Sun StorageTek EBS can validate sign-on attempts against either a central database, such as an LDAP database, or a local a Console database. See also "host authentication."

user alias

The username seen by the Sun StorageTek EBS server when a Console user connects to the Sun StorageTek EBS server.

user groups

Used to assign privileges to Sun StorageTek EBS users. See also "roles."

V

versions

The date-stamped collection of available backups for any single file.

virtual cluster client

A Sun StorageTek EBS client that is not permanently bound to a physical machine, but is instead managed by a cluster manager. It can be bound to more than one physical machine in the cluster and can own its own data disks. It is also referred to as a logical cluster client or a virtual client.

virtual server

A Microsoft term to denote a cluster server. Also known as an agent (Sun)or a package (HP-UX).

volume

1. The physical storage medium, such as magnetic tape, optical disk, or file system to which backup data is written.
2. An identifiable unit of data storage that may reside on one or more computer disks.

volume ID

An internal identification that Sun StorageTek EBS software assigns to a backup volume.

volume mount point

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, and a single disk or partition to be linked to more than one directory tree.

volume name The name assigned to a backup volume when it is labeled. See "label."

volume pool See "pool."

VSS (Volume Shadow Copy Service) Microsoft technology that creates a point-in-time snapshot of a disk volume. Sun StorageTek EBS software backs up data from the snapshot. This allows applications to continue to write data during the backup operation, and ensures that open files are not omitted

VSS component A subordinate unit of a writer. See "writer."

W

writer A 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 "metadata" and "VSS component."

Index

A

aborting a recover 919
ACPI
    OnNow 1014
    recover considerations 1015
    scheduled backup considerations 1015
    Sun StorEdge EBS support 1014
actions, notifications 623
Active Directory
    authoritative restore 1001, 1002, 1003, 1004
    backup 998
    domain controller 998
    explained 998
    recover 999
        prerequisites 465
    SYSTEM STATE save set 965, 998
adding
    annotation 544
    enabler code 885
    folder 593
    host 591
    license 885
    multiple hosts 595
addresses
    server, changing 948
Administration window 51
administration window
    opening 53
administrator
    adding 611
    privileges 599, 610
    Windows groups and Sun StorageTek EBS privileges 612
administrators
    group 61
Advanced Configuration and Power Interface. See ACPI
Advanced File Type Device
    concurrent access 328
    concurrent operations 329
    notification 325
    retention policy 325
    simultaneous recovery of save sets 328
aes encrypted data
    implementing 96
    recovering 441
AIX
    EMC AutoStart, using with 682
alert priority 543, 551, 556
alerts
    lists priority, category, time, and message 546
Alias attribute 938
aliases, host sharing restriction 916
AlphaStor
    firewall issues 985
annotation
    adding 544
    attribute 542
    icon 544
    viewing 544
Annotation Details 519
ansrd service 66, 67
application data
    recovering on Sun STorageTek EBS client UNIX 1036
application information field

Auspex 799
EMC 794, 796
NetApp 791
application specific module (ASM) 235
archive
changing time of 280
request, starting automatically later 559
stopping request in progress 558
Archive attribute 195
Archive button 61
archiving
archive pools
errors 951
archive requests
creating 278
defined 278
deleting 281
disabling 285
editing 280
archive services
enabling 276
archive volume pool 271
cloning 290
compared to backups 270
copying an archive request 279
disable scheduled archive 559
enabling 273, 274
grooming files 270
indexed 272
licensing 270
managing 285
manual 276
naming archive requests, errors 951
NDMP 745
NetWorker User program
toolbar function 61
nonindexed 273
nsrarchive program considerations 951
permissions 273
permissions
Archive feature 274
pools
configuration 250
creating 259
types of pools 274
remote requests, failure 950
requesting an archive 557
requesting an archive's status 547
requirements 270, 887
retrieving from client machine 281
save sets 270
multiple 951
retrieving 282
scheduled 278
scheduling 285
starting 285
stopping 285
time change 280
tracking entries 271
troubleshooting 950
viewing details 286
volumes, cloning 303
ASM (application specific module) 235
ASR save set 815, 969
ASR. See Automated System Recovery
attribute
Administrator 612
annotation 542
auth code 885
category 542
enabler code 885
Grooming
limitation 278
message 542
note 542
priority of managed event 542
See also specific attribute name
server name 542
time 542
auth code attribute 885
authentication
considerations for 659
login 562
nsrauth 658
oldauth 658
authorization code
error message 954
Auto Media Management
cloning, considerations 289
recyclable volumes 210
auto media verification
media position errors 913
Autochanger parallelism 607

autochangers
AIX considerations 1029
attributes
new 926
Auspex support 798
autodetection 926
control ports access 930
destination component 927
EMC Celerra support 793
EMC IP4700 support 795
installation, HP-UX considerations 1023
maintenance commands 926
NDMP support 752, 753
NetApp support 789
sharing
NDMP clients 754
sharing with NDMP support 754
STK-9840 1029
Automated System Recovery
ASR save set 815, 969
cluster database, special handling 825
COM+ database, special handling 825
disk quota database, special handling 825
documentation 814
FAT16 partitions 816
OEM recovery CD 817
overview 814
recovering a client 823
scheduled backup of ASR save set 818, 819, 821
Sun StorageTek EBS software support 814
usage overview 815
verifying client recovery 827
WMI database, special handling 825
Autorestart attribute 171, 681
Autostart attribute 170

B

back ups, manual 100
backing up Console 99
backing up renamed directories 83, 121
Backup button 61
backup configuration wizard 75
backup groups. See groups
backup levels
1-9 190
described 190
options 190
overriding 188
planning 191, 192
types
consolidated 183, 190, 194, 198, 199
full 190, 193, 194, 917
incremental 190
level 194
skip 190
usage 193, 194
backup schedules. See schedules
Backup Type Summary 499
backups
Backup and Recover Server service 619
balancing resources 167, 184
bootstrap 178
client-initiated 250
commands 110
example 112
savepnpc program 116
completed 177
consolidated 183
customization scripts 110
cycle, using levels 182
directives 224
failed 108
filesystems 185
force incremental 175
groups. See groups
hard links 107
incremental
pool for 249
large client filesystems 671
large filesystems 185
levels 193, 202
log file 108
managing 176
manual 74, 88
policies 219
pool for 250
NDMP 744
consideration with Auspex 797
NetWorker User program
browse windows 61
nonscheduled 74, 88
nsrndmp_save command, using 765, 766

online indexes 194
operations
stopping 912
operators group 61, 618
performance when cloning 289
permissions, backup operators group 618
pools 242
previewing 176
recoveries 909
RPC errors 920
save sets. See save sets
server, forced 921
time intervals, setting 174
troubleshooting 909
types of backups. See backup levels
virtual machines 710
VMware Consolidated Backups (VCB) 710
volumes
clones 295
Windows NT registry 91
bar chart 490
basic reports 479
Managed Event 519
best practices for cloud backups 156
BlueArc
requirements 800
BMR support
configuring a client 657, 877
establishing the HomeBase Server 604, 887
recovery process 441, 878
specifying profile options 657, 878, 887
bootstrap
emailing 178
NDMP limitation 744
pools 245
printing 178
failure 921
boot-time file 946
browse policies
about 205, 206
client file index growth 630
clones, storage nodes 306
defined 206
save set recoveries 454, 455
usage 206
browse windows
described 61
toolbar 61
browser
unresponsive 954

C

capture of managed event 541
case sensitivity 532, 581
category attribute of managed event 542
Category Summary, Managed Event basic report 519
centralized
license management 50
Certificate Server
recovery prerequisites 465
SYSTEM STATE save set 965
certified protocols 946
changing
name of folder 595
password 570
changing password, user 570
changing servers 62
characters
illegal 594, 595
characters, not permitted 565, 594, 595
chart formats 486
CHKDSK, running 89
client 650
of Sun StorageTek EBS server 99
Sun StorageTek EBS
definition of 650
client alias, changing 916
client backup configuration wizard 75
client file index
backup level, pool for 249
browse policies 206, 637
checking 918
clones and storage nodes 306
cross-checking 633
defined 629
entries
adding 630
removing 630, 636, 638, 639
growth 630
location, designating 634
managing size 636
operations

backups 194
checking 631
moving 636
recovery 424, 452
policies 636
pools 245
restoration 455
save sets
cycles, removing 638
entries 210
removing 636
size 629
management 636
notification 919
Client parallelism 605
Client Retries attribute 171
Client Status over Time, Sun StorageTek EBS
Backup Status Report 508, 513
Client Summary
Sun StorageTek EBS Backup Statistics Report
498, 511, 516
client-initiated backups
pool for 250
clients
aliases, problems 915
backup commands 110
backups
ad-hoc 67
scheduled 67
client ID
creating new client 457
client/server communication errors 941
cloning
examples 295, 296, 297
configuration 650
defined 650
DHCP 618
DNS name resolution 618
editing 650
groups 166
large filesystems 185
manual backups 74
multiple 671
NetWorker User program 60
operations
archive retrieve 281
backing up renamed clients 918
backups 671
creating 650
editing 650, 651, 652, 660, 662, 666
indexes, moving 636
installation 650
manual backups 74
recovery, failure 918
policies, multiple 215
priority 674
programs 67
save sets 671, 672
services 67
Solaris binary location 1019
unsynchronized client clock 911
clone pools, configuring 250
Clone Volumes dialog box 291
cloning
archives 303
automatic 295
defined 288
destination volume, defined 288
examples 295, 296, 297, 299
manual 292
online indexes, storage nodes 306
recovery
save sets 301
volumes 301
save sets 292
automated 295
manually 295
performing 292
source volume, defined 288
storage nodes
online indexes 306
volume copy, comparison 289
volumes 291, 295
automatic 295
creating 291
details, viewing 291
process 291
cloning to cloud 161
cloud
best practices 156
cloning 161
compared to other device types 155
data consumption information 161
prerequisites 156

reports 526
setting up a cloud device 157
staging 161
support for 155
cluster
performance issues 706
Cluster Server
recovery prerequisites 465
SYSTEM STATE save set 965
clusters
backups 680, 682
clients 680
data corruption 694
devices 692
jbconfig program considerations 695
licensing 887, 888
MC/ServiceGuard 682
recovery types 687
storage node 692
tape libraries 692
virtual client backing up to a local storage node 684
COM+ Database
recovery prerequisites 465
SYSTEM STATE save set 964
command
export 582
gst 69
savepsm 99
command line
reporting 480, 535, 537
reporting program 536
Common Device Interface
SCSI command 360
Completion Data Retention 476
completion data retention 476
completion message retention 476
compression, data 98
computer damage, recovery from 572
configuration
IP issues 947
configuring
Host reports 524
reports 482
Sun StorageTek EBS Backup Statistics reports 496, 509, 515
connection
problem 955
refused 955
connections, server
problem resolution 946
Console
environment variables 576
Console client
starting, after the first time 53
Console Configuration Wizard 582
console security administrator
resetting administrator password 571
Console software
as Sun StorageTek EBS client 99
improving performance 576
License Manager 886
load on CPU 576
logging on 52
managing window 49
URL 52
user roles in 562
Console window
opening 52
console window 49
consolidated backups 194, 198, 199
consolidating save sets
nsrssc command 201
contacting server, problem with 955
Content Index Server
backing up 104
backup 105
defined 104
recover 104, 470
SYSTEM DB save set 104, 470, 967
control zone 588
copying host 592
corrupted database 955
CPU load 576
creating
folder 593
host 591
label templates 266
pools 251
staging policies 315
criteria for organizing hosts 588
critical priority 543, 551, 556
cross-checking online indexes 633
cross-platform

name resolution 618
CSV 534
cutting and pasting host 593

D

daemon log file 641
daemon.log 643
Daily Backup Status 508
Daily Save Set Summary, Sun StorageTek EBS Backup Status Report 506
Daily Summary 499
Daily Summary, Sun StorageTek EBS Backup Status Report 506
damaged computer 572
Dantz servers
monitoring 540
data
compression 98
encryption with aes ASM 96
life cycle
cloned data 290
managing 213
recovery, name conflicts 432
relocating, errors 920
sorting
into pools 243, 246, 248
to storage devices 251
verify 94
data compression 928
Data Retention dialog box 477
database
backing up 99
corrupted 955
corruption 955
delete failed 955
failed to store 955
fetch operation 955
database, fetch operation failure 955
db-output 957
dbsrvr9 960
dbstop output 957
DDS (dynamic drive sharing)
clustering 698
robotics control 701
debug level 576, 581
debug messages
logging stopped 956
default backup schedules 181
default pool 242, 248, 250
definition of
Sun StorageTek EBS client 650
delete operation failed 955
deleting
folder 594
host 592
label templates 267
license 885
managed event note 544
multiple hosts 595
note 544
problem 955
staging policies 318
user 570
device 337
lists of 546
related messages 546
devices
device drivers
maintenance commands 926
device ordering 361
correcting device order problems 363
detecting device order problems 362
disk label errors 921
filesystem
staging 314
labeling errors 921
nonrewinding 931
pools 251
volumes, clones 295
DFS (Distributed File System)
junctions 992
not in SYSTEM save sets 967
recover 995
DHCP (Dynamic Host Configuration Protocol)
database 105, 470
static IP address for Sun StorageTek EBS server 949
DHCP (dynamic host configuration protocol) 943
clients 618
dialog box
Data Retention 477
DinoStor TapeServer
configuration 802

directed recoveries 451
directed recovery
    access 445
    advantages 443
    defined 442
    NDMP 777
    performing 436, 437, 452
    recover program 449, 450
    UNIX clients 452
    use of 443
directives
    ASM (application specific modules) 235
    copying 226
    creating 225
    defined 224
    deleting 226
    editing 225
    naming restrictions 915
    preconfigured 230
disabling
    managed event capture 541
disabling user datagram protocol 578
disallowed characters 565, 594, 595
disappearing managed event 545, 957
disaster recovery 441
disk quota database
    recovery prerequisites 466
    SYSTEM DB save set 967
disk space
    insufficient 954
disk space, gstd log file 576, 581, 959
disk space, gstd.log 581
display problem 955
displaying
    annotation 544
    reports 494
DNS
    hostname alias, troubleshooting 915
DNS (Domain Name System) 943
    host name determination 618
document view 494
domain controller
    Active Directory, configured by 998
    defined 998
    encryption keys
        not supported 1004
    recovering system state on 999
drag-and-drop, column 54
drill-down reports, Managed Event 520
DSA 737
    DSA and NDMP Tape Server distinctions 740
duplicating host 592
dynamic addressing 619
dynamic host configuration protocol (DHCP) 943

E

ECB (Event Control Block) 914
editing
    clients 650
    folder 595
    label templates 266
    pools 258
    staging policies 317
    user 570
emailing the bootstrap report 178
EMC AutoStart 682
EMC, application information values 794
emergency priority 543, 551, 556
enabler code 885
enabler code, entering 885
enabler code, problem 955
enabler codes
    removing 883
enabling
    debug information 576, 581
    JavaScript 961
    software 885
encrypting data
    aes 96
Encrypting File System
    backup 1004
    directed recover 1004
    encryption keys 1004
    explained 1004
    not in SYSTEM save sets 968
    recover 1004
entering
    enabler code 885
    license 885
enterprise
    button 50
enterprise hierarchy

adding folder 593
adding host 591
copying folder 595
copying host 592
deleting folder 594
deleting host 592
folder 588
host 588
managing host 590
managing multiple hosts 595
moving folder 595
moving host 593
renaming folder 595
viewing 590
Enterprise Summary report 521
environment variable 581
GST_DEBUG 959
GST_MAXLOGSIZE 581, 959
GST_MAXLOGVERS 581, 959
setting 581
environment variables
format 642
NSR_DEV_BLOCK_SIZE_MEDIA_TYPE 350, 357
NSR_DEV_LOAD_POLL_INTERVAL_MEDIA_TYPE 358
NSR_DEV_LOAD_TIME_MEDIA_TYPE 358
NSR_DEV_LOAD_TRY_TIMEOUT_MEDIA_TYPE 359
NSR_DEV_TAPE_FILE SIZE_MEDIA_TYPE 358
NSR_DEV-DEFAULT_CAPACITY_MEDIA_TYPE 359
error
database store 955
relocating data 920
RPC (remote procedure call) 920
error messages 954
copy violation 922
destination component 927
disk label 921
environmental variables 952
hardware address not found 1026
illegal record size 914
media verification 913
no space left in tape 1030
nwrecover program 916
print server 625
RPC errors 945
save sets 917
server, unavailable 945
xview 922
Event Control Block (ECB) 914
Event Details 519
event logs
backup 1006
explained 1006
not in SYSTEM save sets 968
recover 1006
Event Time 518
Event Viewer 627
Events by Category 521
Events by Priority 521
Events by Server 521
example
organizing hosts 589
sorting managed event 54
export command 582
export formats 534
exporting
non-ASCII characters 535
reports 534
external authentication to LDAP server 562

F

failed
operation 955
failed operation 955
failover, save group requirements 681
FAT16 partitions, ASR limitation 816
fetch operation 955
Figure 133
file conversion, sparse to fully-allocated 923
file handling
indicators 62
file index missing, message 918
File Manager 916
File Replication Service
backup 1006
SYSTEM STATE save set 965
filemarks 928
filename, support for short names 106
files

compressing 98
encrypting 96
HOSTS 618
log 957
naming conflicts on recovery 432
open files, backing up 179
relocating during recovery 433
servers 673, 938
verify 94
filesystem devices 314
filesystems
backups, large 671
firmware, verification 933
folder 588
adding 593
deleting 594
editing 595
Force Incremental attribute 186
force incremental attribute 175
formats
export 534
full backups 193

G

grooming 278
limitation 278
grooming files 270
group
backup, stopping 549
restarting backup 549
start immediately 549
viewing control details 549, 554
viewing status 546
Group Details, Sun StorageTek EBS Backup Status Report 506
Group resource
Group Status by Server, Sun StorageTek EBS Backup Status Report 508
Group Status over Time, Sun StorageTek EBS Backup Status Report 508
Group Summary
Sun StorageTek EBS Backup Statistics Report 498
Group Summary by Client and Server, Sun StorageTek EBS Backup Status Report 506
Group Summary by Server
Sun StorageTek EBS Backup Statistics Report 499
Group Summary by Server, Sun StorageTek EBS Backup Status Report 506
Group Summary, Sun StorageTek EBS Backup Status Report 505
groups 167
attributes 170
backup operators 618
backups
management 176
previewing 176
bootstrap 178
client policies, multiple 215
completed 177
containing bootstrap 914
defined 166
filesystems, large 185
force incremental, setting 175
listing of per server 546
naming restrictions 915
NetWorker User program privileges 60
operations
copying 173
deleting 173
editing 172
previewing 176
time intervals, setting 174
types 170
gst command 69
GST server process 960
GST_DEBUG 959
GST_MAXLOGSIZE 581, 959
GST_MAXLOGVERS 581, 959
gstd 960
environment variable 581
log 576, 581, 959
process 69
size of log 576, 581, 959
gstd log file 959
gstd service 66
gstmodconf command 595, 597

H

hagentd service 66

hard links, backup and recovery 107
hardware
    upgrade 572
HomeBase Server 877
host
    adding 591
    copying 592
    deleting 592
    managing 590
    moving 593
    organization of 588
    transfer affidavit 574
Host List 521
host name determination
    DHCP clients 618
    TCP/IP 618
Host Reports 478
    types and their configuration 524
hostname alias 915
hostname file 596
HOSTS file 618
hosts table, setting up 942
How 482
HP MC/ServiceGuard
    clusters 682
HP-UX
    creating device files 1024
    EMC AutoStart, using with 682
    installing autochangers 1023
    pass-through drivers 1023, 1024
HTML 534
HTTP service port 52
hung browser 954

I

icon
    annotation 544
    priority 543
IIS (Internet Information Server)
    defined 1006
    recovery prerequisites 466
    SYSTEM STATE save set 964
Importing save set information to a different server 457
Inactivity Timeout attribute 171
incremental backup
    NDMP limitation 745
incremental backups
    pool for 249
Index Save Sets dialog box 632
indexed archives 272
individual user authentication 578
info priority 543, 551, 556
information, sorting table 54
inquire program 752, 1021
install log 957
Interactive view 485
interface
    overview 48
Internet Information Server See IIS
Interval attribute 171
Invalid Object id, error message 955
ioscan program 1023
IP
    errors
        hosts table 942
        troubleshooting 941
    name search order, setting 943
IP configuration, issues 947

J

Java Runtime Environment 537
JavaScript 961
jbconfig program
    clusters 695
    hanging 926
jbexercise program
    NDMP, not supported 745
jobsd service 66
JRE 53, 536, 537
Jukebox parallelism 607

L

label templates
    attributes 264
    components 265
    creating 266
    deleting 267
    editing 266
    naming restrictions 915
    naming strategies 264
    number sequences 264

preconfigured 261
labeling
tips 265
labels
silos, volumes 260
launch button 53
lcmap, improving performance of 706
LDAP login authentication 562
level of debug 576, 581
Level Summary 498
Library parallelism 607
license
adding 885
deleting 885
License allocation failed, error 956
License Manager 886
licensing
archiving 270
clusters 887, 888
copy violation 922
License Manager 886
modules 888
Linux, EMC AutoStart, using with 682
list command 793
List Report, User 523
LMHOSTS 618
local backups with NDMP
introduction 737
localized environments 483
location of
gst command 69
lockbox for pass phrases 657
log file 581, 957, 959
log files
backup/recovery attempts 108
cloning information 288
management 641
viewing 904
logging on
to Console software 52
logging, events 627
login
authentication 562
name 565
password 565
lsdev program 1024, 1025

M

Mac OS X 1032
managed application 588
managed event
define 540
deleting note 544
disabling capture of 541
disappearing 545, 957
polling 576
priority 542
response time
modifying 576
sorting example 54
Managed Event basic reports 519
Managed Event configuration
parameters 518
Managed Event drill-down reports 520
Managed Event Reports 478
Managed Events button 50, 52
managed node 588
adding 591
copying 592
deleting 592
moving 593
managing
host 590
license 886
manual backups 100
manual backups. See backups, manual 74
Manual Restart option 681
Many 976
Max active devices 607
Max parallelism, media libraries 607
Max parallelism, pools 609
Max sessions 609
MC/ServiceGuard
clusters 682
media database
backups 194
cloned data 290
clones and storage nodes 306
compression 637, 640
cross-checking 633
entries, removing 638, 639
managing size 629, 637
restoration 455

retention policies 206
save sets, entries 210
Media Library parallelism 607
media pools. See pools
media position errors 913
memory
NDMP requirements 746
message attribute of managed event 542
message logs
failed backup/recovery attempts 108
management 641
reducing size 642, 643
message retention 476
messages file
Event Viewer 627
Microsoft Automated System Recovery. See Automated System Recovery
Microsoft Windows
backup operators group 618
Event Viewer 627
Windows 2000
databases
not in SYSTEM save sets 967
SYSTEM DB save set 966
Windows Management Instrumentation
SYSTEM DB save set 967
Mirapoint
requirements 800
mminfo program
reports 221
modifying
log size 581
response time 576
user 570
modules
NDMP 762
NetWorker database modules 888
monitoring
Dantz servers 540
Retrospect servers 540
Monthly Summary 499
Monthly Summary, Sun StorageTek EBS Backup Status Report 505
moving host 593
multiple hosts
adding or deleting 595
Multiplexing 608
multiplexing
performance issues 43

N

name
conflicts on recover 432
folder, editing of 595
resolution cross-platform 618
servers, issues 948
name resolution 943
name servers, disabling 943
naming restrictions 915
native login authentication 562
NDMP (Network Data Management Protocol) 736
application information variables 761
archives 745
Auspex
autochanger handle 798
autochangers
configuration 753
list command 793
sysconfig program 789
backups
command line options 765, 766
parallel 744
bootstrap limitation 744
changing location of temporary space directory 747
clients
definitions 744
configuration
options 737
Sun StorageTek EBS server 750
device pathname 752
directed recoveries 777
EMC Celerra
autochanger handle 793
EMC IP4700
autochanger handle 795
Force Incremental setting 762
incremental backup limitation 745
inquire program 752
memory requirements 746
NetApp

autochanger handle 789
nsr resource attributes 745
operations
browsing 744
recoveries 770
recoveries, directory level 771, 774
prerequisites 764
protocol limitations 746
save set operations
consolidation 745
recovery 745
savegroup Parallelism setting 763
scanner program 745
sharing autochangers 754
silo support 757
storage node support for NDMP client backup 768
storage node support for NDMP clients 747, 748
Sun StorageTek EBS
changes with NDMP 744
index database 746
index processing and memory requirements 746
temp storage space requirements 746
Nested Mountpoints 467
network
DNS connection issues 946
IP configuration issues 947
TCP/IP certified 946
NetWorker
User program. See NetWorker User program
NetWorker Backup Statistics drill-down reports 517
NetWorker User program
backing up registry, Windows NT 91
browse windows 61
changing servers 62
compression 98
connecting to a server 63
encryption 98
manual backups 88
overview 60
password protection 98
privileges 60
server connection 62
starting 60
toolbar 61
node
adding 591
copying 592
deleting 592
managed 588
moving 593
nonindexed
archives 273
note
attribute 542
deleting from managed event 544
notifications 627
defined 619
deleting 629
operations
customizing 622
preconfigured 620
printing 624
priorities 626
programs 623
SNMP
configuring 808
creating 810
modifying 809
nsrtrap 809
NSR_AVOID_ARCHIVE 195
nsr_getdate program 219, 220
nsr_render_log 904
nsr_shutdown program 643, 914
nsradmin program
editing the nsrla.res database 653
starting 65
nsrarchive program 271
nsrauth authentication 658
nsrck program 452
nsrck service 66
nsrd daemon 643
nsrd service 66, 67, 643
nsrexecd 70
nsrexecd daemon 643
nsrexecd service 67, 643
nsrim service 66
nsrindexd service 66
nsrjb program
troubleshooting, HP-UX 1027
nsrla.res database 653

nsrlcpd service 67
nsrlogd service 66, 67
nsrlpr program 624
nsrmm program 220, 455, 637, 914
nsrmmd daemon 358, 359
    NDMP
        unsupported options 745
nsrmmd service 66, 67
    NDMP
        unsupported options 745
nsrmmdbd service 66
nsrmmgd service 67
nsrndmp_recover program 770, 771, 774
nsrndmp_save program 768
    backing up data 765, 766
nsrpm service 67
nsrssc program 201
nsrtrap 808, 809
    command line options 749, 809
    verbose mode 809
numeric order 54
nwrecover program 63, 284
    browse policy 206

O

object, reference 955
OEM recovery CD, ASR limitation 817
oldauth authentication 658
online indexes
    cross-checking 633
    entries
        checking 631, 918
        removing 630, 638
    information
        refreshing 634
        viewing 631
    management
        manual 629
        size 636
    moving 636
    recovery
        location 915
    restoration 454
    save sets, viewing 632
    size considerations 630
    volumes
        removing 638, 639
OnNow
    Advanced Configuration and Power Interface (ACPI) 1014
    recover considerations 1015
    scheduled backup considerations 1015
    Sun StorEdge EBS support 1014
open files, backing up with VSS 180
operation failed 955
optimizing Console 576
organizational criteria 588
organizational structure, labeling for 265
organizing
    Sun StorageTek EBS Servers
        example 589
out of memory 954

P

page cannot be displayed 955
Parallelism 604
parallelism
    performance 43
password
    changing 570
pathname restrictions 916
Path-to-Tape support 782
PDF report format 534
performance
    cloning 289
    features 43
Performance Counters
    SYSTEM STATE save set 964
performance of Console software 576
permissions
    Archive feature 273
    backup operators group 618
Persistent binding 361
Persistent naming 362
pie chart 491
pie report format 490
plot chart 491
policies 216
    backups, manual 219
    browse
        about 205, 206
        data life cycle 213

defined 206
modifying 220
usage 206
clients 215
data life cycle 213
multiple 215
naming restrictions 915
overriding 219
planning 184
retention 209
about 205, 206
data life cycle 213
defined 206
modifying 220
usage 206
volume relabeling 206
setting expiration 477
policy
deleting 216
polling 577
pool
consolidated backup 246
copying resource 258
Pool parallelism 609
Pool Type attribute, with archive pools 275
pools
archive 250
errors 951
archives, creating 259
archiving 259
bootstrap 245
client file index 245
clones 250
configuration
archive 250
clone 250
creating 251
criteria 243, 246
data sorting 243, 248
default 242, 248, 250
clone pool 250
defined 242
devices 251
editing 258
expression matching 245
incremental backups 249
manual backups 250
precedence 246
restrictions 243
save set consolidation 202
sorting data 251
volume labels 260
port
HTTP service 52
port range
troubleshooting 949
portmappers
verifying 944
POSIX hard links, problems recovering 107
PostScript report format 534
Power Monitor service 1014
precedence for pools 246
preconfigured notifications 627
priorities, notifications 626
priority
of managed event 542
symbol 543, 551, 556
Priority attribute 542
Priority Summary 519
privileges
administrator 599
problem
contacting server 955
process, stopping and restarting 69
processor load 576
Procom NetFORCE
requirements 799
Program not registered 956
program not registered 956
programming thread 577
properties
user 570
Properties dialog box
NetWorker User program 62
protocols, certified 946
provider 831
ps command 960

R

rearranging enterprise hierarchy 593
rearranging information in table 54
Recover button 61
recover program 68

clusters, using with 688
media database 210
options 449, 450
retention policy 209
recoveries 425
archives 281
Backup and Recover Server services 619
backup operators group 618
client file index 452
clients, renamed 918
clone volumes 301
cluster 687
conflict resolution 432
conflicts 431
directed 451
access 445
advantages 443
defined 442
recover program 449, 450
usage 443
disaster-related 441
failed 108
failure 916
files, finding for recovery 434
hard links created by POSIX 107
index-based
advantages 424
interrupted backups, from 917
log file 108
methods, comparing 424
NDMP 770, 771, 774
NetWorker User program
browse windows 61
planning 184
relocation 433
save set recovery 436
save sets
client file index 454, 455
media database 455
single file 68
type, determination 424
VMware Consolidated Backups (VCB) 728
volumes
required 435
recovery
aborting 919
ASR 823
recycling. See volumes
reference object 955
refused connection 955
registering program 956
registry
backups 1007
backward compatibility 1007
explained 1007
REGISTRY save set 1007
SYSTEM STATE save set 964, 1007
relocating data 920
remote access
recoveries 939
Remote Access list 939
remote archives, failure 950
remote backups with NDMP TapeServer software
performing 805
Remote NDMP backups with DinoStor
TapeServer
performing 803
remote NDMP backups with SnapImage software
performing 805
Removable Storage Manager
recovery prerequisites 466
SYSTEM DB save set 967
removing
folder 594
host 592
license 885
user 570
renamed directories, backing up 83, 121
renaming folder 595
report
background processing 495
chart types 489
document mode 487
export formats 534
interactive mode 485
restricted views 494
User List 523
reports
basic and drill-down 478
clones 295
command line reporting program 536
customized 480
daemon log file 641
daemon.log 643

date and time formats 482
Managed Event Basic 519
Managed Event Drill-Down 520
message logs
management 641
reducing size 642
save set policies 221
saved 531
viewing 494
Reports button 50, 52
requirements
NetWorker User groups 60
resetting administrator password 571
resolved events 545
resource
archive
changing archive time 280
request
copying 279
status 547
label template
copying 267
notification
copying 628
deleting 629
policy
deleting 216
pool
copying 258
directing data from consolidated backup 246
staging
copying a policy 317
user group
copying 617
creating 616
customizing privileges 614
deleting 617
editing 617
preconfigured 613
resources
configuring for NDMP 750
response time
managed event 576
Restart Window attribute 171
restrictions on pathnames 916
retention policies 209
about 205, 206, 209
clones, storage nodes 306
defined 206
usage 206
volumes, relabeling 206
retention, completion data 476
retention, completion message 476
retention, save set 476
retrieval
troubleshooting 950
retrieving archives 281
retrieving save sets 282
Retrospect servers
monitoring 540
RPC (remote procedure call)
errors 920
rpcinfo program 944

S

save program 68, 110, 219, 220, 688
save set
All save set 86
Save Set Data Retention
and Sun StorageTek EBS Backup Statistics Reports 497
Save Set Details 503
Sun StorageTek EBS Backup Statistics Report 500, 516
Save Set Details by Client 500
Sun StorageTek EBS Backup Statistics Report 500
Save Set Details by Client, Sun StorageTek EBS Backup Status Report 507
Save Set Details, Sun StorageTek EBS Backup Status Report 507
Save Set Name 504
Save Set Output, Sun StorageTek EBS Backup Status Report 507
save set retention 476
Save Set Summary
Sun StorageTek EBS Backup Statistics Report 498, 516
save sets
archives, retrieving 282
backup commands 110

backups 671
client combination 672
client file index, entries 210
client priority 674
cloning
automatically 296
described 292
manually 295
status 292
consolidation
files, large 199
limitations 200
NDMP 745
nsrssc command 201
pools 202
usage 198, 199
defining 672
indexes, viewing 632
information, viewing 632
load balancing 671
media database, entries 210
multiplexing 608
policies
modification 220
reports 221
predefined 85
recoveries 436, 437, 438
client file index 455
conflicts 440
media database 455
NDMP 745
online indexes 454
relocating 439
requirements 437
usage 436
volumes 438, 439
recovery
compared to index-based recovery 424
REGISTRY, backward compatibility 1007
staging. See staging
status
clone 292
retention policy 209
suspected 301
SYSTEM
manual backups 90
point-in-time recovery
command prompt 463
recover
command prompt 461
SYSTEM DB
backup levels 197
basic components 966
databases not included 967
optional components 967
recovery prerequisites 466
SYSTEM FILES
backup levels 196
basic components 965
SYSTEM STATE
backup levels 196
basic components 964
optional components 964
recovery prerequisites 465
VSS ASR DISK
backup levels 198
components 973
VSS OTHER
backup levels 198
components 972
VSS SYSTEM
manual backups 90
point-in-time recovery 463
VSS SYSTEM BOOT
backup levels 197
components 970
VSS SYSTEM FILESET
backup levels 197
components 970
VSS SYSTEM SERVICES
backup levels 198
components 971
VSS USER DATA
backup levels 198
components 972
savefs program 68
savegroup completion report 621
Savegroup parallelism 606
savegrp program 67
backup limitation 921
save set consolidation from command line 201
savegrp service 67

savepnpc program
message logging by 119
using with customized backup program 116
savepsm 99
savestream 242
scalability 45
SCANDISK, running 89
scanner program
NDMP 745
record size 914
recovering clone volumes 301
recovering save sets from volumes 213
recreating online indexes 438
recreating volume entries 438
retention policy 209, 210
volumes, read-only 914
Schedule attribute 170
Schedule resource 166
schedules 180
attributes 186
backup cycle 182
balancing 184
copying 187
default 181, 183
deleting 187
editing 187, 201
large filesystems 185
load balancing 671
naming restrictions 915
overriding 188
planning 184
preconfigured 181
staggered 183
usage 180
Schedules window 186
scheduling backups 99
SCSI address selection for HP-UX 1023
SCSI ID 136
security
application authentication 658
backup server access 609
backup server user groups 613
console server users 562
lockbox for pass phrases 657
login authentication 562
server
problem contacting 955
setting up 599
web address 52
server name
attribute of managed event 542
Server parallelism 606
Server Summary 519
Sun StorageTek EBS Backup Statistics Report 498, 499, 511
servers
address, changing 948
administrators
adding 611
backup
operators Group 618
client/server communication errors 941
clients
tasking rights 673
DCHP 618
DNS name resolution 618
dynamic addressing 619
errors, binding to 945
file 673, 938
index
backup, failure 921
management 629
notifications
priorities 626
ping command, testing 944
testing 944
service
ansrd 66
gstd 66
hagentd 66
jobsd 66
nsrck 66
nsrd 66
nsrim 66
nsrindexd 66
nsrlcpd 67
nsrlogd 66, 67
nsrmmd 66, 67
nsrmmdbd 66
nsrmmgd 67
nsrpm 67
port, HTTP 52
savegrp 67
services

Backup and Recover Server 619
described 66
Power Monitor service 1014
stopping and starting 68
session management 44
sessions
lists of backup, recover, or browse sessions 546
Setting 477
setting
data retention policies 477
environment variable 581
expiration policies 477
severity of managed event 542
shadow copy 830
short filename support 106
silos
NDMP support 757
Simple Network Management Protocol. See SNMP
single file recovery 68
size of
gstd log 959
log file 581
sleep state
defined 1014
SnapImage
autochanger handle 803
snapshot 830
snapshot policy
creating 218
SNMP (Simple Network Management Protocol) 808
configuring 808, 811
defined 808
notifications
configuring 808
creating 810
modifying 809
nsrtrap 808, 809
traps 808
Software Administration button 50
Solaris
EMC AutoStart, using with 682
troubleshooting 1019
sorting managed event, example 54
sorting table 54
Source Client dialog box 436, 438, 440
sparse files 1008
not in SYSTEM save sets 968
sparse files, converting 923
Special Handling dialog box 98
stacked bar report format 490
stacking bar chart 491
staging
defined 314
filesystem devices 314
policies
creating 315
deleting 318
editing 317
to cloud 161
Start Time attribute 170
starting
server process 69
starting services 68
status
viewing
group status 546
status values, scanned-in 454
stopping Console server 69
stopping services 68
storage nodes
affinity
problems 952
clusters
virtual storage nodes, support 692
timeouts 952
troubleshooting 952
virtual server 691
volumes, clones 295
Store Index Entries attribute, with archive pools 275
storing database failed 955
Sun 619
Sun StorageTek EBS
interfaces 63
NDMP
effects on 744
server
organization of 588
startup commands 910
Sun StorageTek EBS Backup Statistics basic reports 498, 511, 513, 516

Sun StorageTek EBS Backup Statistics drill-down reports 500, 513, 517
Sun StorageTek EBS Backup Statistics Reports 478
Sun StorageTek EBS Backup Statistics reports
types and their configuration 496, 509, 515
Sun StorageTek EBS Backup Status drill-down report 507
Sun StorageTek EBS Backup Status Reports 478
Sun StorageTek EBS Backup Status reports 502
Sun StorageTek EBS client
definition of 650
Sun StorageTek EBS DiskBackup
advanced file type concurrent operations 329
file type and advanced file type differences 337
notification 325
retention policy 325
Sun StorageTek EBS server
organizing 588
support
for firewall 976
suspected save sets 301
symptom of problem 954
sysconfig command 789
System File Protection
backup 965
defined 965
explained 965
recover 965
SYSTEM FILES save set 965
system standby 1014
SYSVOL
SYSTEM STATE save set 965

T

tables
display or hide columns in 55
multicolumn sorting 55
rearranging columns 54
sorting 54
Target sessions 608
TCP/IP
certification 946
changing Sun StorageTek EBS server address 948
DHCP clients 618
host name determination 618
troubleshooting hostname alias problems 915
technical support, troubleshooting information 902
temporary enabler code
expired 954, 956
temporary notes 542
Terminal Services Licensing
backup 1008
defined 1008
recover 1009
prerequisites 466
SYSTEM DB save set 967
This 925, 1023
thread, programming 577
three-party NDMP backups
introduction 742
performing 766
time attribute of managed event 542
time range 481
Together 166
toolbars
User program 61
tracking
cloned data 290
online index information 630
traps
categories 811
SNMP 808
troubleshooting 576, 581
aborted recover 919
AIX
STK-9840 1029
archive pools 951
archive requests
naming 951
archives 950
multiple save sets 951
nsrarchive program 951
remote request failure 950
auto media verification 913
autochangers
AIX considerations 1029
attributes 926
autodetected scsi errors 926

control port access 930
destination component 927
HP-UX considerations 1026, 1027
maintenance 926
backups 909
backups levels 917
backups, stopping 912
bootstrap printing, failure 921
client file index
size growth 919
clients
alias 915
Solaris, location 1019
unsynchronized errors 911
connectivity issues 948
daemons 909
devices
maintenance 926
nonrewinding 931
disk label errors 921
DNS connection problems 946
DNS hostname alias 915
ECB counter 914
file conversion 923
firmware 933
hosts table 942, 943
HP-UX
error messages 1026
unloading drives 1027
unsupported media 1027
illegal characters 915
IP configuration issues 947
IP errors 941, 942
IRIX 1030
licensing, copy violation 922
name resolution 943
name server issues 948
name servers, disabling 943
nsrexec processes 909
online indexes 915
packet receive buffer 914
pathname restrictions 916
port range configuration 949
portmappers, verifying 944
recoveries 909
interrupted backups 917
new installation 916
online indexes 915
remote access 939
recovering POSIX hard links 107
renamed client backups 918
retrievals 950
routers 933
rpcinfo command 944
scanner program 914
server errors, binding to 945
server index 921
Solaris 1018
storage nodes 952
technical support information 902
xview errors 922

U

uasm program 915
unresponsive browser 954
unsynchronized clock errors 911
URL
Console software 52
user
authentication 578
deleted 532
deleting 570
editing 570
properties 570
user authentication 578
user datagram protocol
disabling 578
user groups 613
creating 616
privileges 614
user interface
overview 48
setting preferences 59
User List Report 523
user privileges 60
User Reports 478
user roles in Console server 562

V

vanishing managed event 957
variable
case sensitivity 581
for gstd log 959

GST_DEBUG 959
GST_MAXLOGSIZE 581
GST_MAXLOGVERS 581, 959
setting 581
VCB 710
backups 710
recovery 728
verification
NetWorker User program
browse windows 61
of files 94
Verify button 61
verifying ASR recovery 827
versions of the gstd log file 581
view
document view 494
viewing
annotation 544
enterprise hierarchy 590
reports 494
virtual clients 710
Automatic discovery 722
virtual machines
configuring a NetWorker client 710
Virtualization 724
VMware 710, 876
Automatic discovery of VMware environments 722
Configuring a NetWorker client 718
Licensing 710, 892
Notification of changes to VMware environment 723
Visual representation of environment 724
VMware Consolidated Backups (VCB) 710
backups 710
recovery 728
volume pool
See also pools
archive 271
defined 242
Volume Shadow Copy Service
commands 839
controlling from Administration window 834
controlling from command-prompt 838
controlling from Sun StorageTek EBS client 836
limitations with Microsoft Exchange Server 427
limitations with Microsoft SQL Server 427
overview 830
writers 833
volumes
archive 303
basic copy, problems with 289
client file index
removing 638, 639
cloning 291
archive data 303
creating 291
performance 289
recovery 301
labeling 260
maximum size 265
tips 265
modes
types 392
nonrewinding 931
recoveries, required for 435
recycling 637
relabeling 387
removing 637
save sets
recoveries 438, 439
verify 94
VSS. See Volume Shadow Copy Service
VTL
path-to-tape support 782

W

waiting priority 543, 551, 556
warning priority 543, 551, 556
web
browser
unresponsive 954
Console
server name 52
WINS (Windows Internet Naming Service) 618
WINS (Windows Internet Naming Service)
database 105, 470
Wizard
Console Configuration 582
wizard

client backup configuration 75
writer 833

X

xview, errors 922

Z

zone, control 588