Sun StorEdge Enterprise Backup Software™ 7.3

Administration Guide

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

Part No. 819-4972-10
December 2005, Revision 01

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

Copyright © 2005 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, AnswerBook2, docs.sun.com, and Solaris are trademarks or registered trademarks of Sun Microsystems, Inc. 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. Legato NetWorker® is a trademark of Legato Systems, Inc.

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 © 2004 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, AnswerBook2, docs.sun.com, et Solaris sont des marques de fabrique ou des marques déposées de Sun Microsystems, Inc. aux Etats-Unis et dans d'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. Legato NetWorker® est un marque de fabrique de Legato Systems, Inc.

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



Table of Contents

Table of Contents **3**

List of Figures **13**

List of Tables **15**

Preface **21**

Before You Read This Book 21
Using UNIX Commands 22
Typographic Conventions 23
Related Documentation 24
Accessing Sun Documentation Online 24
Sun Welcomes Your Comments 24

Chapter 1: Overview **25**

The Sun StorEdge EBS Environment 25
Sun StorEdge EBS Features 27
Common Sun StorEdge EBS Tasks 29
Sun StorEdge EBS User Interfaces 31
Sun StorEdge EBS Services 47

Chapter 2: Backing Up Data **53**

Scheduled Backups 53
Save Sets 63
Ad Hoc Backups 66
Verifying Backed-Up Data 72

Password Protecting Windows or NetWare Backup Data 73
Encrypting Backup Data 74
Compressing Backup Data 77
Special Data Handling for Sun StorEdge EBS Clients on Windows 78
Backing Up Sun StorEdge EBS Console Management Data 78
Backing Up Windows Mount Points 80
Backing Up the Windows Content Index Server 83
Backing Up Windows DHCP and WINS Databases 84
Windows Backup and Recovery Notes 85
Customizing the Backup Command 88
Considerations for Backing Up Raw Partitions 98
Backing Up a Mapped Drive 99

Chapter 3: Configuring Libraries and Devices 101

Devices and Libraries 101
Storage Nodes 102
Dynamic Drive Sharing 108
Autodetection of Libraries and Devices 117
Add (Configure) Libraries 121
Library Reconfiguration 125
Add (Configure) Stand-Alone Devices 127
Deleting Libraries and Devices 129
Library Slots 130
Troubleshooting 132

Chapter 4: Backup Groups and Schedules 133

Overview of Sun StorEdge EBS Scheduling 133
Backup Groups 134
Managing Backup Groups 140
Bootstrap Report Generation and Printout 142
Backing Up Open Files 142

Schedules 144
Backup Levels 154
Save Set Consolidation 163

Chapter 5: Browse and Retention Policies 169

About Browse and Retention Policies 169
Managing the Data Lifecycle 176
Browse and Retention Policies for Ad Hoc Backups 182
Modifying the Browse and Retention Policy on a Save Set 183

Chapter 6: Directives 185

Directives Overview 185
Creating a Global Directive Resource 186
Editing a Global Directive Resource 187
Deleting a Global Directive Resource 187
Copying a Global Directive Resource 187
Applying Directives 188
Sun StorEdge EBS User Program Local Directives 190
Preconfigured Global Directive Resources 192
Format of Directive Statements 194
Application Specific Modules (ASMs) 197

Chapter 7: Sorting Backup Data 203

Media Pools 203
Label Templates 217

Chapter 8: Archiving 225

Overview of Archiving 225
Permissions for Archiving 228
About Archive Pools 230
Archiving Data Procedures 231

Retrieving Archived Data ... 238
Archive Request Management ... 242

Chapter 9: Cloning Backups ... 243

Overview of Cloning ... 243
Volume Cloning ... 246
Save Set Cloning ... 247
Manual Clone Operations ... 256
Recovering Cloned Data ... 257
Cloning Archived Data ... 258
Directing Clones to a Special Storage Node ... 258
Considerations for File Type and Advanced File Type Devices ... 260

Chapter 10: Staging Backups ... 261

Save Set Staging ... 261
Working with Staging Policies ... 262
Sun StorEdge EBS DiskBackup ... 266

Chapter 11: Device Operations ... 297

Compatible Block Size for UNIX and Microsoft Windows ... 298
Library Device Maintenance ... 300
Device Parameter Settings ... 303
Common Device Interface ... 307
Nonrewinding Tape Device Usage (UNIX/Linux Only) ... 308
Block Size Mode (UNIX/Linux Only) ... 309
Silos ... 309
Miscellaneous Device Operations ... 317
Tips for Using Libraries ... 320

Chapter 12: Media Management 325

Storage Management Operations 325
Auto Media Management 328
Volume Operations 330
Media Handling Errors 354
Media Management in a Silo 355
Volume Save Sets 361

Chapter 13: Recovering Data 367

Choosing a Recovery Method 367
Recovering by File Selection 368
Recovering by Save Set Selection 379
Disaster Recovery 386
Recovering Encrypted Data 386
Directing Recoveries to Another Client 387
Recovering Online Indexes for a Sun StorEdge EBS Server 395
Recovering the Console Database 399
Moving a Sun StorEdge EBS Client to a New Sun StorEdge EBS Server 399
Recovering the Windows System Configuration 400
Recovering Windows Volume Mount Points 408
Recovering Windows DHCP and WINS Databases 411
Restoring a Windows Content Index Server 412

Chapter 14: Reporting 413

Gathering Report Data 413
Data Retention and Expiration Policies 415
Report Categories 418
Configuring Reports 421
Viewing Reports 426
Preconfigured Reports 436

Customizing and Saving Reports 462
Sharing Reports 464
Exporting Reports 465
Command-Line Reporting 466
Printing Reports 468

Chapter 15: Monitoring 469

Overview of Monitoring 469
Events 469
Monitoring Window 475

Chapter 16: Console Server Management 491

Managing Console Server Users 491
Moving the Console Server Component 498
Managing Product Licenses 500
Setting System Options 504
Setting Environment Variables 511
Using the Console Configuration Wizard 514
Displaying International Fonts in Non-US Locale Environments 514

Chapter 17: Sun StorEdge EBS Server Management 515

Enterprise 515
Configuring a Sun StorEdge EBS Server 527
Improving Backup Performance 528
Managing Server Access 530
Sun StorEdge EBS User Groups 533
Server Communication Issues within Microsoft Windows 540
Notifications 541
Indexes 552
Message Log Management 564

Chapter 18: Sun StorEdge EBS Client Management 567

Sun StorEdge EBS Client Overview 567
Client Configuration 568
Sun StorEdge EBS Authentication 571
Multiple Clients from the Same Computer 583
Adding or Changing the Sun StorEdge EBS Servers for a Client 585
Client Priority 586

Chapter 19: Cluster Support 587

Cluster Terminology 587
Cluster-Aware and Cluster-Unaware Sun StorEdge EBS Servers 588
Backing Up Data on UNIX Cluster Environments 589
Configuring a Virtual Client to Back Up to a Local Storage Node 592
Configuring an External Client to a Virtual Server 593
Recovering Data in a UNIX Cluster Environment 594
Tracking Scheduled Saves 595
Using a Physical Storage Node with a Virtual Server on UNIX Cluster Platforms 599
Using Autochangers and Stand-Alone Tape Devices 600
Sun StorEdge EBS Cluster Licensing 614
Client Configuration Wizard Support Requirements 616

Chapter 20: Sun StorEdge EBS Server Support for NDMP 617

Overview of NDMP 617
Supported NDMP Configurations 619
Requirements for and Limitations of NDMP Support 625
Configuring the Sun StorEdge EBS Server for NDMP Operations 633
Performing NDMP Backups 647
Performing NDMP Recoveries 655
NDMP Cloning 665
Ancillary Product Information 667

Chapter 21: SNMP Module 691

SNMP Traps 691
Configuring StorEdge EBS SNMP Notifications 692
Configuring SNMP Management Software 695

Chapter 22: Support for Microsoft Removable Storage 697

Microsoft Removable Storage 697
Types of Support for Removable Storage 699
Removable Storage Media Pools 702
Usage Scenarios 704
Configuring the Sun StorEdge EBS Server for Removable Storage 706
Library Operations with Removable Storage 710
Device Cleaning 720
Considerations for Uninstalling the Sun StorEdge EBS Software 720

Chapter 23: Support for Microsoft SharePoint Portal Server 721

Microsoft SPS 721
The SHAREPOINT Save Set 723
SPS Backup 724
SPS Recovery 726
Performing an SPS Disaster Recovery 729

Chapter 24: Support for Microsoft Automated System Recovery 731

Microsoft ASR 731
Sun StorEdge EBS Support for ASR 732
ASR Backup 734
Creating an ASR Disk 735
Using the ASR Disk to Recover a Sun StorEdge EBS Client 739
Limitations and Special Considerations 743

Chapter 25: Volume Shadow Copy Service 745

Overview of VSS 745
VSS and the Backup Process 746
Controlling VSS from NetWorker Software 748

Chapter 26: Troubleshooting 755

Before Contacting Technical Support 755
Sun StorEdge EBS Functionality Issues 758
Devices and Autochangers 771
Sun StorEdge EBS Locale and Code Set Support 776
Resource Database Notes 777
Network and Server Communication Errors 779
Sun StorEdge EBS Archiving and Retrieval 790
Storage Nodes 791
Console Error Messages and Corrective Actions 793
Console Log Files 798
Console Troubleshooting Notes and Tips 799

Appendix A: SYSTEM and VSS SYSTEM Save Sets 801

SYSTEM Save Sets 801
VSS SYSTEM Save Sets 807

Appendix B: Firewall Support 811

Firewall Support 811
Configuring Firewall Support 816
Troubleshooting Port Configuration Problems 823
Connections between the Console Server and Clients 824
Connections between the Console Server and Sun StorEdge EBS Hosts 826
Additional Firewall Support Considerations 828

Appendix C: Backing Up and Restoring a Microsoft DFS 831

Overview of a Microsoft Distributed File System 831
DFS Backups and Restores 832

Appendix D: Additional Features of the Microsoft Windows Server 837

Active Directory 837
Encrypting File System 844
Event Logs 844
Internet Information Server 845
Windows Registry 845
Sparse Files 846
Terminal Server Licensing 847
Windows Change Journal 848
Advanced Configuration and Power Interface 852

Appendix E: UNIX and Linux Platform-Specific Notes 855

Solaris 855
Linux 856
HP-UX 857
AIX 867
IRIX 868

Appendix F: Mac OS X Support 869

Support for Mac OS X 869
Mac OS X Backup Considerations 870
Recovering Files and Directories on Mac OS X 873

Glossary 875

Index 893

List of Figures

Figure 1. Sun StorEdge EBS Components 26
Figure 2. Sun StorEdge EBS Components 26
Figure 3. Console Window 32
Figure 4. Administration Window 34
Figure 5. Sun StorEdge EBS User Program 39
Figure 6. Example of Browse Window 41
Figure 7. **nwbackup** Program 43
Figure 8. **nwrecover** Program 44
Figure 9. **nwarchive** Program 45
Figure 10. **nwretrieve** Program 46
Figure 11. SYSTEM Save Sets in Sun StorEdge EBS User Backup Window 69
Figure 12. VSS SYSTEM Save Sets in Sun StorEdge EBS User Backup Window 70
Figure 13. Dynamic Drive Sharing 110
Figure 14. DDS with NDMP 114
Figure 15. How Sun StorEdge EBS Groups Are Used to Back Up Multiple Clients 135
Figure 16. Weekly Backup Cycle 146
Figure 17. Using the Sun StorEdge EBS Default Schedule for Multiple Clients 147
Figure 18. Staggered Weekly Schedules for Multiple Clients 148
Figure 19. Backups for October 1 through October 7 156
Figure 20. Backups for October 1 through October 14 156
Figure 21. Backups for October 1 through October 15 157
Figure 22. Consolidated Backup 167
Figure 23. One-Week Browse Policy 171
Figure 24. Two-Week Browse Policy 172
Figure 25. One-Week Backup Cycle and Three-Week Retention Policy 174

Figure 26. Data Lifecycle in the Client Index and the Media Database 177
Figure 27. Using Media Pool Configurations to Sort Data 206
Figure 28. Labeling a Volume by Using a Label Template 217
Figure 29. Overview of Archive Operation 227
Figure 30. Automatic Cloning Scheme 253
Figure 31. Configuration Showing an AFTD 272
Figure 32. Sample Sun StorEdge EBS DiskBackup Configuration 287
Figure 33. CIFS-Mounted Advanced File Type Device 294
Figure 34. How Library Sharing Works 319
Figure 35. A Directed Recovery from a Remote Client 388
Figure 36. Sun StorEdge EBS Servers Worldwide 517
Figure 37. Configuration Option 1 601
Figure 38. Configuration Option 2 603
Figure 39. Configuration Option 3 604
Figure 40. DDS Enabled with Robotics Controlled in the Cluster 607
Figure 41. DDS Enabled with the Robotics Controlled Outside the Cluster 612
Figure 42. Robotics Control between Clusters 613
Figure 43. NDMP Local Configuration 620
Figure 44. Backup Initiated from a Sun StorEdge EBS Server with an Attached Storage Device 622
Figure 45. File History and Data to Sun StorEdge EBS Storage Node 623
Figure 46. Backup Using SnapImage Module 623
Figure 47. NDMP Tape Server 624
Figure 48. Usage Scenario: One Shared Library 704
Figure 49. Usage Scenario: Two Shared Libraries 705
Figure 50. VSS Backup Process 747
Figure 51. Storage Node Adjacent to the Sun StorEdge EBS Client 813
Figure 52. Storage Node Adjacent to the Sun StorEdge EBS Server 815
Figure 53. Open Ports Required for Communications through a Firewall 825
Figure 54. Connecting to a Sun StorEdge EBS Server through a Firewall 827

List of Tables

Table 1. Typographic Conventions 23
Table 2. Related Documentation 24
Table 3. Windows Launched from the Console Window 32
Table 4. Windows Launched from the Administration Window 34
Table 5. Sun StorEdge EBS User Groups Requirements 40
Table 6. Sun StorEdge EBS User Toolbar Functions 40
Table 7. Main Sun StorEdge EBS Services 47
Table 8. Main Services for Sun StorEdge EBS Hosts 49
Table 9. Sun StorEdge EBS Startup Commands 52
Table 10. Accounting Department Backup Requirements 54
Table 11. Components in the All Save Set 65
Table 12. Shared Devices Attributes 115
Table 13. Common jbedit Options 126
Table 14. Group Attributes 136
Table 15. Backup Group Management 140
Table 16. Preconfigured Sun StorEdge EBS Schedules 145
Table 17. Key Components of a Schedule 150
Table 18. Time Zone Strings 153
Table 19. Backup Levels 154
Table 20. Advantages and Disadvantages of Backup Levels 158
Table 21. Regular File and Index Backup Levels 159
Table 22. Preconfigured Directives 192
Table 23. Save Environment Keywords 196
Table 24. Sun StorEdge EBS ASMs 197

Table 25. Media Pool Resource Attributes to Direct Client Indexes and Bootstrap to a Separate Media Pool 207
Table 26. Sun StorEdge EBS Hierarchy for Resolving Media Pool Conflicts 209
Table 27. Key Label Template Attributes 219
Table 28. Examples of Number Sequences for Volume Labels 221
Table 29. Using Label Template Components 222
Table 30. Save Set Status Settings 249
Table 31. Accounting Incremental Schedule 254
Table 32. Accounting Full-Clone Schedule 254
Table 33. AFTD Names 273
Table 34. Concurrent Operations within a Single Storage Node 274
Table 35. Concurrent Operation with Storage Nodes Sharing an AFTD in NAS 274
Table 36. Attributes Requiring Manual Propagation 281
Table 37. Advanced File Type and File Type Device Differences 281
Table 38. TapeAlert Severity 301
Table 39. Sun StorEdge EBS Startup Script Location on UNIX/Linux Systems 305
Table 40. DAS and StorageTek Environment Variables 314
Table 41. Sun StorEdge EBS Startup Commands 316
Table 42. Library Resource Sleep Attributes 320
Table 43. Volumes Detail Table 332
Table 44. Volume Modes 333
Table 45. Volume Save Sets Window 362
Table 46. Naming Conflict Resolution Options 374
Table 47. Naming Conflict Resolution Options 374
Table 48. Naming Conflict Resolution Options 385
Table 49. Access Permissions Used in Directed Recoveries 391
Table 50. Recover Options 393
Table 51. Preparing to Recover the SYSTEM STATE Save Set 406
Table 52. Preparing to Recover the SYSTEM DB Save Set 407
Table 53. Data Retention Policies 415
Table 54. Report Categories 418

Table 55. Date and Time Input Formats for Common Locales 424
Table 56. Report Chart Formats 431
Table 57. Sun StorEdge EBS Backup Statistics Parameters 436
Table 58. Sun StorEdge EBS Backup Statistics Basic Reports 438
Table 59. Sun StorEdge EBS Backup Statistics Drill-Down Reports 442
Table 60. Sun StorEdge EBS Backup Status Parameters 446
Table 61. Sun StorEdge EBS Backup Status Basic Reports 448
Table 62. Sun StorEdge EBS Backup Status Drill-Down Reports 451
Table 63. Event Parameters 453
Table 64. Event Basic Reports 454
Table 65. Event Drill-Down Reports 456
Table 66. Host Reports 457
Table 67. User Reports 458
Table 68. Device Reports 462
Table 69. Report Export Formats 465
Table 70. Events Columns 471
Table 71. Event Priorities 472
Table 72. Monitoring Window Tabs 475
Table 73. Groups Tab Icons 477
Table 74. Alerts Tab Icons 480
Table 75. Devices Tab Icons 481
Table 76. Operations Tab Icons 483
Table 77. Log Tab Icons 485
Table 78. Archive Requests Tab Icons 487
Table 79. recoverpsm Options 500
Table 80. Console System Options 505
Table 81. Usernames for Individual User Authentication 510
Table 82. Console Environment Variables 511
Table 83. Viewing the Enterprise 518
Table 84. User Group Privileges 535
Table 85. Preconfigured Notifications 541

Table 86. Events 545
Table 87. Actions 546
Table 88. Priorities 549
Table 89. Event Viewer Messages 550
Table 90. Indexes Window Information 555
Table 91. Index Save Sets Dialog Box Information 555
Table 92. Sun StorEdge EBS Startup Script Location 565
Table 93. Distinctions between NDMP Tape Server and DSA 621
Table 94. DSA Feature Save Set Flags 631
Table 95. Vendor-Specific Information 632
Table 96. Configuring Resources for NDMP 633
Table 97. NDMP Server Examples 667
Table 98. NetApp Application Information Variables 673
Table 99. EMC Celerra Variables 678
Table 100. EMC IP4700 Variables 679
Table 101. Auspex Variables 683
Table 102. Command-Line Options for nsrtrap 692
Table 103. Removable Storage Support by Operating System 699
Table 104. Library Hardware Supported for Removable Storage 700
Table 105. VSS Commands 753
Table 106. Boot-Time File Locations 786
Table 107. Error Messages or Symptoms 793
Table 108. Installation Log File Locations 798
Table 109. Environment Variables for the GSTD Log 799
Table 110. SYSTEM STATE Save Set Basic Components 802
Table 111. SYSTEM STATE Save Set Optional Components 802
Table 112. SYSTEM DB Save Set Basic Components 805
Table 113. Windows Databases Not in SYSTEM DB Save Set 806
Table 114. Sun StorEdge EBS Server Port Requirements 818
Table 115. Client Port Requirements 819
Table 116. Options for the **nsrports** Command 821

Table 117. Sample of rpcinfo Output 822
Table 118. Console Client/Server Communications Ports 824
Table 119. Sun StorEdge EBS Save Triggers 850

Preface

The *Sun StorEdge Enterprise Backup 7.3 Administrator's Guide* contains information on how to configure and manage the Sun StorEdge 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.

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 StorEdge Enterprise Backup 7.3 Installation Guide*
- *Sun StorEdge Enterprise Backup 7.3 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

Typographic Conventions

Table 1. Typographic Conventions

Convention	Indicates	Example
boldface	Names of line commands, daemons, options, programs, or scripts	The **nsradmin** command starts the command line version of the Sun StorEdge EBS Administrator program.
italic in text	Pathnames, filenames, computer names, new terms defined in the Glossary or within the chapter, or emphasized words	The specifications are created, one for each swap file, in *c:\pagefile.sys*.
italic in command line	A variable that must be provided in the command line	**save -s** *server_name*
`fixed-width`	Examples and information displayed on the screen	`media waiting: recover waiting for 8mm 5GB tape volume name`
`fixed-width, boldface`	Commands and options that must be entered exactly as shown	**`nsr_shutdown -a`**
Menu_Name> Command	A path or an order to follow for making selections in the user interface	Volume>Change Mode>Appendable
Important:	Information that must be read and followed to ensure successful backup and recovery of data	**Important:** Use the **no_verify** option with extreme caution.
Backslash (\)	Paths for both Windows and UNIX unless specified otherwise	<*Sun StorEdge EBS_install_path*>\logs

Related Documentation

Table 2. Related Documentation

Appliation	Title	Part Number
Installation	*Sun StorEdge Enterprise Backup Software 7.3 Installation Guide*	819-4973-10
Release Notes	*Sun StorEdge Enterprise Backup Software 7.3 Release Notes*	819-4974-10

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 (TBD) of your document in the subject line of your email.

Chapter 1: Overview

This chapter covers the following topics:

- "The Sun StorEdge EBS Environment" on page 25
- "Sun StorEdge EBS Features" on page 27
- "Common Sun StorEdge EBS Tasks" on page 29
- "Sun StorEdge EBS User Interfaces" on page 31
- "Sun StorEdge EBS Services" on page 47

The Sun StorEdge EBS Environment

The Sun StorEdge 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 StorEdge EBS software provides the power and flexibility to meet these challenges.

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

1

Sun StorEdge EBS Components

Figure 1 on page 26 illustrates the main components in a Sun StorEdge EBS environment.

Figure 1. Sun StorEdge EBS Components

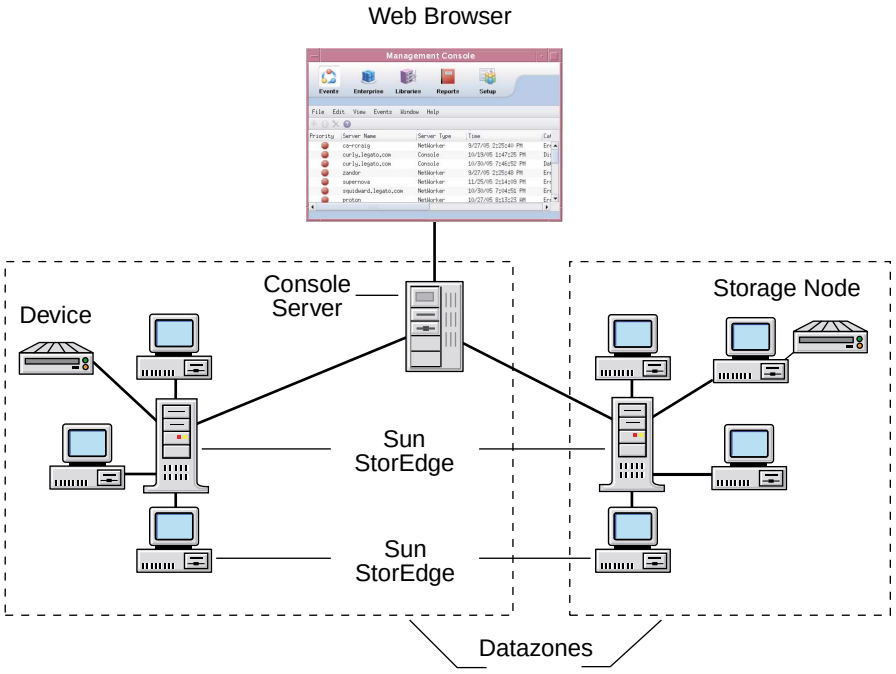


Figure 2. Sun StorEdge EBS Components

Console Server

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

Web Browser

The Console server is accessed through a graphical interface that can be run from any computer that has a web browser. Multiple users can access the Console server concurrently from different browser sessions. A computer that hosts the web browser can also be a Sun StorEdge EBS client.

Sun StorEdge EBS Server

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

Datazone

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

Sun StorEdge EBS Storage Node

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

Sun StorEdge EBS Client

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

Sun StorEdge EBS Features

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

Performance Features

Standard Sun StorEdge 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.
- Back up to file devices and optional subsequent staging to near-line or offline volumes.

1

Ease of Use

Sun StorEdge 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 StorEdge EBS tasks and functions.
- Administer and configure Sun StorEdge 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. Windows users can also access the *NetWorker Command Reference Guide*, which contains information that is equivalent to the UNIX man pages.
- Take advantage of the automatic media management feature to allow the Sun StorEdge EBS server or storage node to label and mount volumes as needed for backups.
- Use the integrated knowledge base and technical bulletins at *http://softwaresupport.emc.com* to find answers to common questions.

Scalability

Existing Sun StorEdge 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 StorEdge EBS server.
- Add the EMC Legato® License Manager (LLM) software to administer all of your network's Sun software licenses from a single server. LLM is only available from Sun.

Optional Software Additions

Optional additions to the Sun StorEdge EBS software include:

- Sun StorEdge EBS Autochanger Module
- Sun StorEdge EBS Silo Software Module
- Sun StorEdge EBS Archive Module
- NetWorker Database Modules (for backing up several types of databases)
- Sun StorEdge EBS Open File Connection Module (OFC) (Microsoft Windows NT only)
- StorEdge EBS SNMP (Simple Network Management Protocol)
- NDMP (Network Data Management Protocol) support
- Cluster support, including Sun FullTime AutoStart
- License Manager (available only from Sun)
- Advanced reporting capability

Common Sun StorEdge EBS Tasks

The following are some of the most common tasks performed with the Sun StorEdge EBS Management Console:

How to Schedule a Backup

Follow these steps to schedule backups:



1. Click the Enterprise button on the taskbar, and add Sun StorEdge EBS servers.
2. Highlight a server. Right-click and select Launch Application.



3. Click the Configuration button on the taskbar.
4. Define schedules, groups, and clients.

For more information about scheduling backups, see "Chapter 2: Backing Up Data" on page 53.

1

How to Configure Media

Follow these steps to configure media using the Sun StorEdge EBS Management Console:



1. Click the Enterprise button on the taskbar, and add Sun StorEdge EBS servers.
2. Highlight a server. Right-click and select Launch Application.



3. Click the Devices button on the taskbar.
4. In the tree view, right-click a server, and select Scan for Devices.

For more information about configuring media, see "Chapter 3: Configuring Libraries and Devices" on page 101.

How to Label Tapes

Follow these steps to label tapes from the Sun StorEdge EBS Management Console:



1. Click the Enterprise button on the taskbar, and add Sun StorEdge EBS servers.
2. Highlight a server. Right-click and select Launch Application.



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

For more information about labeling tapes, see "Chapter 12: Media Management" on page 325.

How to View Failed Backups

To see whether any backups have failed:



1. Click the Enterprise button on the taskbar, and add Sun StorEdge EBS servers.
2. Highlight a server. Right-click and select Launch Application.



3. Click Monitoring.

For information about viewing backup status information, see "Chapter 15: Monitoring" on page 469.

Sun StorEdge EBS User Interfaces

The Sun StorEdge EBS application consists of the following user interfaces:

- "Sun StorEdge EBS Management Console Interface" on page 31
- "Sun StorEdge EBS Client Interface" on page 38
- "Sun StorEdge EBS Character-Based Interface" on page 47
- "Sun StorEdge EBS Command-line Interface" on page 47

Sun StorEdge EBS Management Console Interface

The interface for Sun StorEdge EBS Management Console (also referred to as Sun StorEdge EBS Console) consists of two main windows:

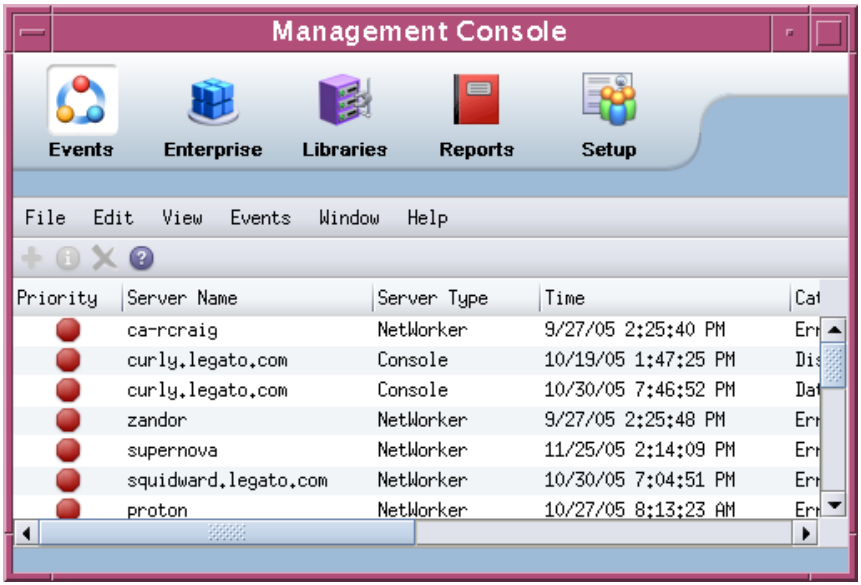
- Console window
- Administration window

1

Console Window

When Sun StorEdge EBS software is started, the Console window appears as shown in Figure 3 on page 32.

Figure 3. Console Window



The Console window is the first point of access for Sun StorEdge EBS tasks. Table 3 on page 32 lists the task-based windows that can be launched from the Console window taskbar.

Table 3. Windows Launched from the Console Window (Part 1 of 2)

Button	Window	Description
	Events	View important messages about Sun StorEdge EBS and Console servers. For more information about managed events, see "Events" on page 469. The Administrator window can also be opened from this location.
	Enterprise	Select a Sun StorEdge EBS server to manage and monitor the server and its backup clients. Click this button to open the Administration window for a Sun StorEdge EBS server.

Table 3. Windows Launched from the Console Window (Part 2 of 2)

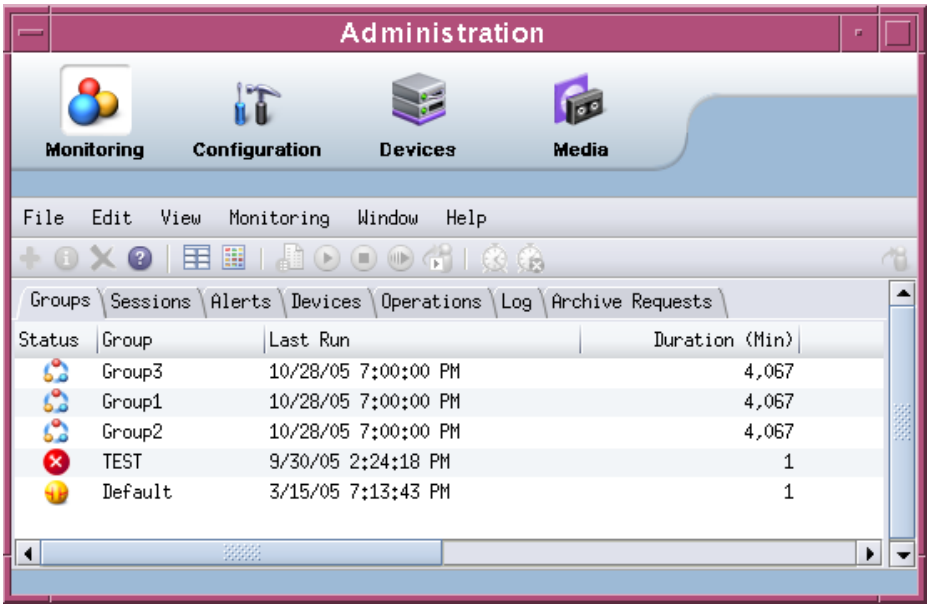
Button	Window	Description
	Libraries	View summary information and manage libraries for all Sun StorEdge EBS servers. The Administrator window can also be opened from this location.
	Reports	Configure and view Console reports. For more information about reports, see "Chapter 14: Reporting" on page 413.
	Setup	Control administrative functions: • User management — Add, edit, and delete Console user accounts, restrict user views of servers. For more information about user management, see "Managing Console Server Users" on page 491. • License management — Manage Sun StorEdge EBS licenses. For more information about license management, see "Managing Product Licenses" on page 500.

1

Administration Window

Sun StorEdge EBS servers are managed through the Administration window as shown in Figure 4 on page 34.

Figure 4. Administration Window



One can toggle between the Administration window and the Console window.

Table 4 on page 34 lists the windows that can be launched from the Administration window taskbar.

Table 4. Windows Launched from the Administration Window (Part 1 of 2)

Button	Window	Description
	Monitoring	Monitor various activities related to the Sun StorEdge EBS server. For example, one could monitor the progress of a scheduled backup and view any alerts from the Monitor window.
	Configuration	Manage the Sun StorEdge EBS server and its resources such as clients, backup schedules, and policies. For example, one could create a backup schedule, add Sun StorEdge EBS clients, and apply the backup schedule to several Sun StorEdge EBS clients.

Table 4. Windows Launched from the Administration Window (Part 2 of 2)

Button	Window	Description
	Devices	Add, configure, and operate devices, libraries, and silos for the Sun StorEdge EBS server.
	Media	Manage activities and resources related to backup volumes. For example, one could mount a backup volume or create a template for labeling backup volumes.

How to Start Console for the First Time

These steps assume that the Sun StorEdge EBS software is installed and that all of the software and hardware requirements have been met on the computer that will access Console. For more information, refer to the appropriate *Sun StorEdge EBS Installation Guide*.

To open Console for the first time:

1. Start a web browser session.
2. Enter the URL of the Console server:

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

   where:

   server_name is the name of the computer where the Console server component was 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.)

   For example:

   `http://houston:9000`
3. From the Welcome page, click Start Sun StorEdge EBS Management Console.
4. From the Security Warning screen, click Start to run Sun StorEdge EBS Management Console.
5. If JRE 1.4.2 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.

6. Complete the Java Web Start Desktop Integration dialog box by selecting one of the following 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.
7. In the Sun StorEdge EBS Management Console Login dialog box, enter the username and password.



Important: The default password for the administrator is “administrator.” For added security, change the password during the first login session. If the administrator password is lost, forgotten, or otherwise needs changing, see "Managing Console Server Users" on page 491.

8. Click OK.

   The Console Window appears.

How to Start Console after the First Time

After Console has been started the first time, start it by using any of the following methods:

- Point the browser to the same URL as in "How to Start Console for the First Time" on page 35.
- Double-click the Sun StorEdge EBS Console product name in the Java Web Start Application Manager.
- Double-click the desktop icon, if one was set up through the Java Web Start Application Manager.

How to Open the Administration Window

To select a Sun StorEdge EBS server and open the Administration window:

1. From the Console window, click Enterprise.
2. From the left pane, click a Sun StorEdge EBS server in the Enterprise list.
3. From the right pane, click the application.
4. From the Enterprise menu, select Launch Application.

   The Administration window is launched as a separate application.

How to Sort 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: 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 the Events button.
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. For further information about reports, see "Chapter 14: Reporting" on page 413.

How to Perform MultiColumn Sorts of a Table

To sort on two or more columns within a table:

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

For example, given a large table of Events, one could 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.

How to Display or Hide 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 unselect the column.
3. Click OK.

How to Set User Interaction Preferences

One can set various user preferences such as the user interface font, font size, window sizes, and table settings.

To set user preferences:

1. From the View menu, select User Preferences.
2. Set various user preference options as appropriate.
3. Click OK.

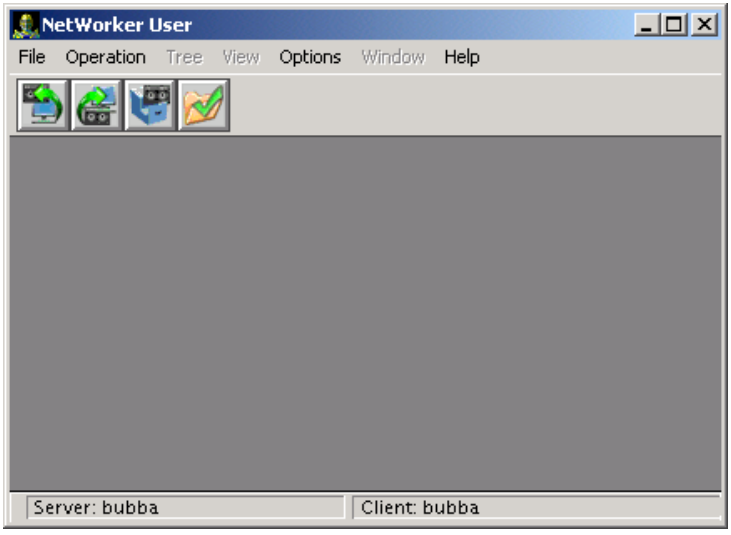
Sun StorEdge EBS Client Interface

This section describes the interface from which users can recover data and perform ad hoc backup and archive operations. Ad hoc operations are not scheduled. Instead, they are performed when a user wants to back up or archive one or more files on the Sun StorEdge EBS client immediately. Scheduled backup and archive operations are set up through the Console interface. For information about the Console interface, see "Sun StorEdge EBS Management Console Interface" on page 31.

Windows Client Interface

The Sun StorEdge EBS User program shown in Figure 5 on page 39 is the Windows client interface.

Figure 5. Sun StorEdge EBS User Program



How to Start the User Program

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

- Click the Windows Start button and select Programs>Sun StorEdge EBS Group>Sun StorEdge EBS User.
- From the Administration window, click Monitoring and then from the Monitoring menu, select Launch Sun StorEdge EBS User Application.

  Note: The Sun StorEdge EBS client package must be installed on the host from which you are attempting to start the Sun StorEdge EBS User program.

To start the Sun StorEdge EBS User program, you must belong to the appropriate Windows groups. Table 5 on page 40 lists the groups that you must belong to in order to run the Sun StorEdge EBS User program.

1

Table 5. Sun StorEdge 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

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

Toolbar Buttons

The Sun StorEdge EBS User program has a toolbar with buttons for common User program tasks. Table 6 on page 40 describes the function of each button.

Table 6. Sun StorEdge EBS User Toolbar Functions

Button	Name	Function
	Backup	Starts an ad hoc (unscheduled) backup of the client's data to a Sun StorEdge 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.

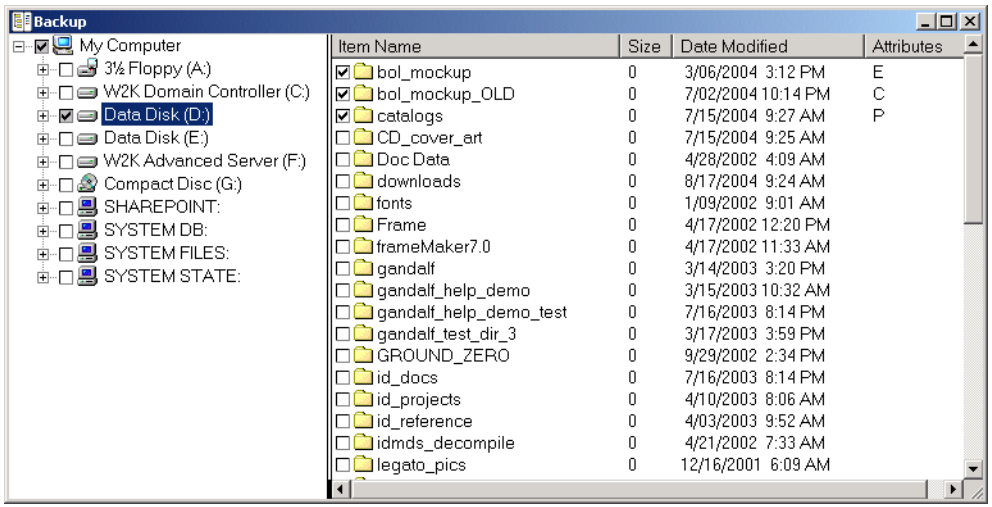
Browse Window

The Sun StorEdge EBS User program browse window appears when one of the following is selected:

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

The browse window, shown in Figure 6 on page 41, displays the directory tree of the file system that is being browsed.

Figure 6. Example of 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 check mark beside an item name indicates that the item is marked for backup, recovery, archiving, or verification.

The Attributes column indicates any special handling option that has been applied:

- P — The item is marked for password-protection. For more information, see "Password Protecting Windows or NetWare Backup Data" on page 73.
- E — The item is marked for password-protection and encryption. For more information, see "Encrypting Backup Data" on page 74.
- C — The item is marked for compression. For more information, see "Compressing Backup Data" on page 77.

1

How to Connect to a Sun StorEdge EBS Server

A typical user that runs the Sun StorEdge EBS User program only needs to connect to the Sun StorEdge EBS server that performs scheduled backups. However, to perform a directed recovery or to back up files to another server, it may be necessary to connect to a different Sun StorEdge EBS server.

Note: Before the Sun StorEdge EBS User program can connect to a Sun StorEdge EBS server, the client computer must be set up as a Client resource on that server. For information about creating a Client resource, see "Task 7: Create a Backup Client Resource" on page 61.

To connect to a Sun StorEdge EBS server:

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

UNIX Client Interface

The UNIX version of the Sun StorEdge EBS client interface uses separate graphical interface programs for performing each of the following tasks:

- Ad hoc backups (**nwbackup**)
- Recoveries (**nwrecover**)
- Archiving (**nwarchive**)
- Retrieval of archived data (**nwretrieve**)

Note: The Sun StorEdge EBS client package must be installed on the host from which you are attempting to start any of the client graphical interface programs.

The Ad Hoc Backup Program

Use the **nwbackup** program to perform an ad hoc backup initiated from the client computer, rather than from the Sun StorEdge EBS server's scheduled group backup.

To start the **nwbackup** program, perform one of the following tasks:

- From the Administration window, click Monitoring and then from the Monitoring menu, select Launch Sun StorEdge EBS User Application.
- Enter the following command:

  nwbackup

  To specify the Sun StorEdge EBS server to connect to, enter the **-s** *server* option with **nwbackup**:

  nwbackup -s *server*

Figure 7 on page 43 displays the **nwbackup** program.

Figure 7. nwbackup Program



1

The Recover Program

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

To start the **nwrecover** program, enter the following command:

```
nwrecover
```

- To specify the Sun StorEdge EBS server to connect to, enter the **-s** *server* option.
- To specify the Sun StorEdge EBS client, use the **-c** *client* option:

  `nwrecover -s` *server* **`-c`** *client*

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

Figure 8 on page 44 displays the **nwrecover** program.

Figure 8. nwrecover Program



The Archive Program

Use the **nwarchive** program to manually archive files to a Sun StorEdge EBS server.

Note: Archive functionality is enabled by license. See the Sun StorEdge EBS Roadmap for licensing information.

To start the **nwarchive** program, enter the following command:

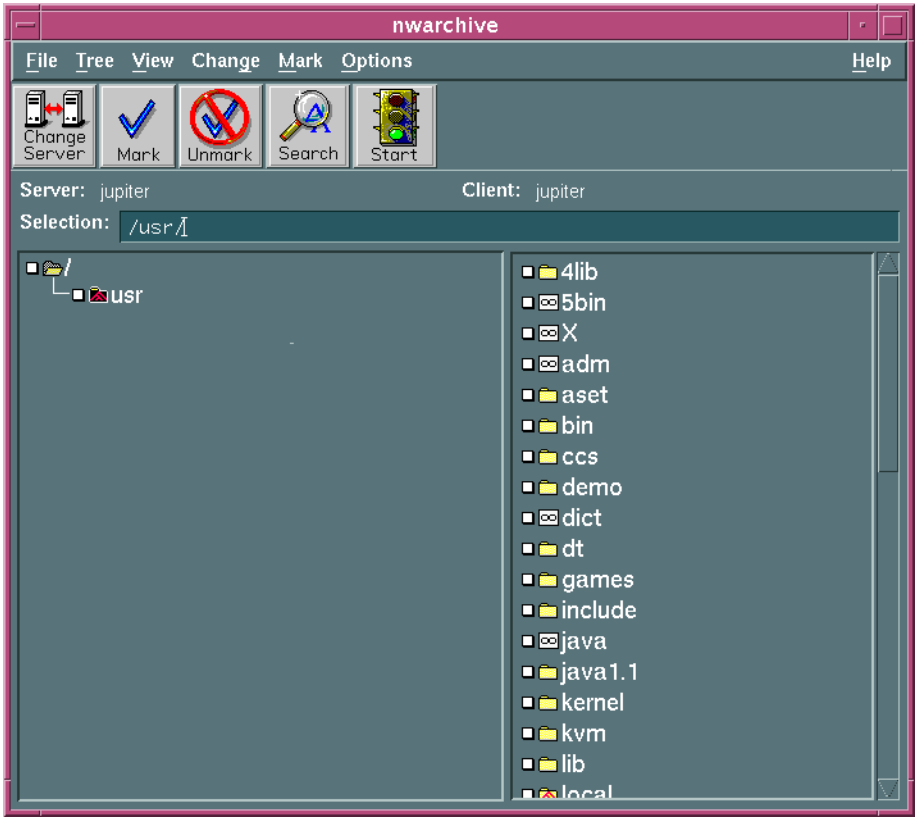
```
nwarchive
```

To specify the Sun StorEdge EBS server to connect to, enter the **-s *server*** option with **nwarchive**:

```
nwarchive -s server
```

Figure 9 on page 45 displays the **nwarchive** program.

Figure 9. nwarchive Program



1

The Archive Retrieval Program

Use the **nwretrieve** program to retrieve archived files.

To start the **nwretrieve** program, enter the following command:

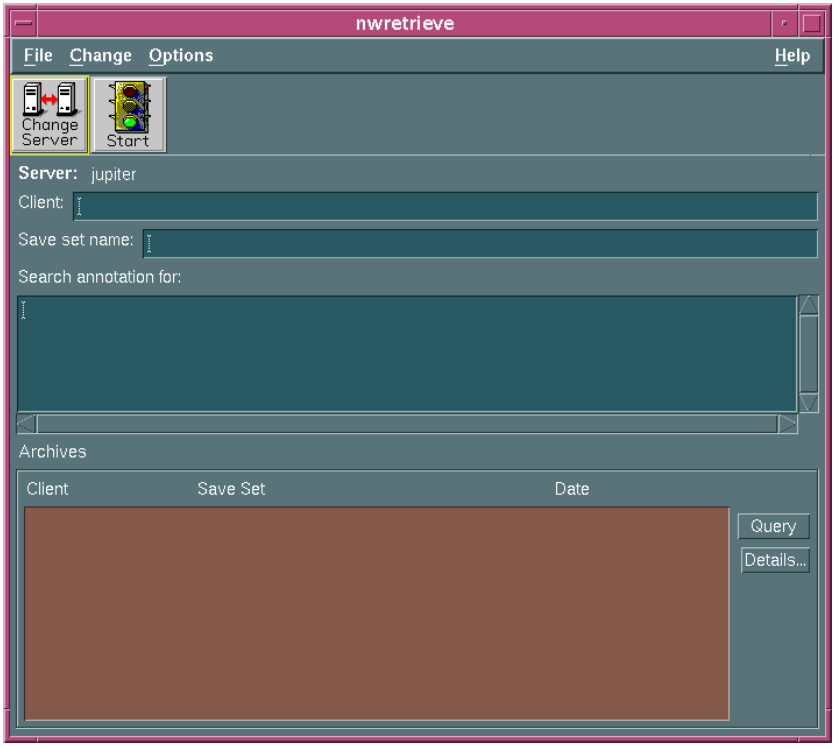
```
nwretrieve
```

To specify the Sun StorEdge EBS server to connect to, enter the **-s** *server* option with **nwretrieve**:

```
nwretrieve -s server
```

Figure 10 on page 46 displays the **nwretrieve** program.

Figure 10. nwretrieve Program



Sun StorEdge EBS Character-Based Interface

The Sun StorEdge EBS character-based interface (**nsradmin**) can be used to perform configuration and management tasks for a Sun StorEdge EBS server.

To start the **nsradmin** interface, enter the following command:

```
nsradmin
```

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

Sun StorEdge EBS Command-line Interface

Sun StorEdge EBS client and server tasks can also be performed by entering commands at the prompt. For information about these commands, refer to the *NetWorker Command Reference Guide* or the UNIX man pages.

Sun StorEdge EBS Services

Table 7 on page 47 describes the main services that provide Sun StorEdge EBS software functionality. For more information about these services, refer to the *NetWorker Command Reference Guide* or the UNIX man pages. For information about service port requirements when configuring a firewall, see "Appendix B: Firewall Support" on page 811.

Table 7. Main Sun StorEdge EBS Services (Part 1 of 2)

Service	Function
ansrd	Monitors an active **recover** session. This agent process is spawned by **nsrd** in response to a **recover** session.
gstd	Controls other services provided by the Console server. This master service is called the Generic Services Toolkit (GST) service.
hagentd	Manages the Sun StorEdge EBS host agent client modules and maintains the Sun StorEdge EBS host-agent resource database.
nsrlcpd	Provides a uniform library interface to the media management service **nsrmmgd**.
nsrd	Provides save and recover operations. This master service controls other services on the Sun StorEdge EBS server, clients, and storage nodes.

1

Table 7. Main Sun StorEdge EBS Services (Part 2 of 2)

Service	Function
nsrexecd	Authenticates and processes the Sun StorEdge EBS server's remote execution requests and executes the save programs on the client.
nsrindexd	Inserts entries in the client file index based on information passed by the **save** program.
nsrjobd	Allows for monitoring and long term recording of Sun StorEdge EBS activities. All functionality requiring monitoring or remote execution is handled with help of the **nsrjobd** service.
nsrmmd	Provides device support, generates mount requests, and *multiplexes* save set data during a multiclient backup. The **nsrd** service can start one **nsrmmd** service for each enabled device. An additional **nsrmmd** service is started while mount requests are pending.
nsrmmdbd	Provides media and save set database management services to the local **nsrd** and **nsrmmd** service, and records entries in the media database.
nsrmmgd	Provides an RPC-based service that manages all jukebox operations on behalf of the **nsrd** service.
nsrpm (Windows only)	Works with the energy-saving features of Windows to keep Sun StorEdge EBS components running properly during power-saving modes. It can be disabled if not needed.

Manually Stopping and Starting Sun StorEdge EBS Services

Sun StorEdge EBS hosts might need to be manually stopped and restarted in conjunction with configuration and troubleshooting activities.

Table 8 on page 49 lists the main Sun StorEdge EBS services that are required for each type of Sun StorEdge EBS host.

Table 8. Main Services for Sun StorEdge EBS Hosts

Sun StorEdge EBS Host	Sun StorEdge EBS Service
Console server	**gstd nsrexecd**
Sun StorEdge EBS server	**nsrd**, **nsrexecd**, **nsrindexd**, **nsrjobd**, **nsrlcpd, nsrmmd**, **nsrmmdbd, nsrmmgd**
Sun StorEdge EBS client	**nsrexecd**
Sun StorEdge EBS storage node	**nsrexecd**, **nsrlcpd, nsrmmd**
License Manager (available only from EMC)	**lgtolmd**

Note: The Sun StorEdge EBS service **nsrmmd** is only present if one or more devices are enabled.

Stopping and Starting the Console Server

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

How to 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 LEGATO GST Service (**gstd**), then select Stop.

1

How to 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 StorEdge EBS client is running.

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

   a. Right-click Sun StorEdge EBS Remote Exec Service.
   b. Select Start.
4. Right-click EMC LEGATO GST Service (**gstd**), then select Start.

How to Stop the Console Server on UNIX

To stop the Console server:

1. Log in as root.
2. Enter the following command:

   `/etc/init.d/gst stop`

How to Start the Console Server on UNIX

To start the Console server:

1. Log in as root.
2. Verify that the Sun StorEdge EBS client is running.

   For example, enter the following command:

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

   If the Sun StorEdge EBS client is running, a message similar to the following appears:

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

   If the Sun StorEdge EBS client is *not* running, start it. For information about starting the Sun StorEdge EBS client, see "Start a Sun StorEdge EBS Host on UNIX" on page 52.
3. Start the Console server by entering the following command:

   `/etc/init.d/gst start`

Stopping and Starting a Sun StorEdge EBS Server, Client, or Storage Node

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

How to Stop a Sun StorEdge EBS Host on Windows

To stop a Sun StorEdge EBS 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. Select Sun StorEdge EBS Remote Exec Service (**nsrexecd**), right-click, and select Stop.

How to Start a Sun StorEdge EBS Host on Windows

To start a Sun StorEdge EBS 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 StorEdge EBS client or storage node:

   Right-click the Sun StorEdge EBS Remote Exec Service (**nsrexecd**) and select Start.

How to Stop a Sun StorEdge EBS Host on UNIX

To stop the Sun StorEdge EBS services:

- Log in as root and enter the following:

  `nsr_shutdown`

1

How to Start a Sun StorEdge EBS Host on UNIX

To start Sun StorEdge EBS services:

- Log in as root and enter the appropriate startup command listed in Table 9 on page 52.

Table 9. Sun StorEdge 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*

How to Shut Down a Sun StorEdge EBS Host on Mac OS X

To shut down the Sun StorEdge EBS host:

1. Log in as a Mac Administrator and Open the Mac OS X Terminal application utility.
2. Shut down the Sun StorEdge EBS service by entering the following command:

   ```
   # SystemStarter stop NetWorker
   ```

How to Start a Sun StorEdge EBS Host on Mac OS X

To start the Sun StorEdge EBS host:

1. Log in as a Mac Administrator and open the Mac OS X Terminal application utility.
2. Start the Sun StorEdge EBS client by entering the following command:

   ```
   # SystemStarter start NetWorker
   ```

Chapter 2: Backing Up Data

This chapter covers the following topics:

- "Scheduled Backups" on page 53
- "Save Sets" on page 63
- "Ad Hoc Backups" on page 66
- "Verifying Backed-Up Data" on page 72
- "Password Protecting Windows or NetWare Backup Data" on page 73
- "Encrypting Backup Data" on page 74
- "Compressing Backup Data" on page 77
- "Special Data Handling for Sun StorEdge EBS Clients on Windows" on page 78
- "Backing Up Sun StorEdge EBS Console Management Data" on page 78
- "Backing Up Windows Mount Points" on page 80
- "Backing Up the Windows Content Index Server" on page 83
- "Backing Up Windows DHCP and WINS Databases" on page 84
- "Windows Backup and Recovery Notes" on page 85
- "Customizing the Backup Command" on page 88
- "Considerations for Backing Up Raw Partitions" on page 98

Scheduled Backups

This section explains how to plan and create scheduled backups.

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: Planning Scheduled Backups for the Accounting Computers

Company XYZ wants to ensure that all of the computers in the Accounting department are backed up according to the requirements listed in Table 10 on page 54. This table also maps each requirement to specific Sun StorEdge EBS features.

Table 10. Accounting Department Backup Requirements (Part 1 of 2)

Requirement	Sun StorEdge EBS Feature	For More Information
Backups occur at the same time.	• Backup Schedule • Backup Group	"Task 1: Set Up a Schedule for Backups" on page 56 "Task 2: Set Up a Group for Backup Clients" on page 56
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 57
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 57
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 58 "Task 5: Set Up a Pool to Sort Backup Data" on page 59
To avoid unnecessary backups, do not back up files with a .tmp extension.	Directives	"Task 6: Set Up Directives for Special Processing" on page 60

Table 10. Accounting Department Backup Requirements (Part 2 of 2)

Requirement	Sun StorEdge EBS Feature	For More Information
The same files and folders are backed up on each accounting computer.	Client resource	"Task 7: Create a Backup Client Resource" on page 61
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 57 "Task 7: Create a Backup Client Resource" on page 61

Setting Up a Scheduled Backup

You can create a scheduled backup quickly by using the Client Configuration Wizard. For more information, see "Using the Sun StorEdge EBS Client Configuration Wizard" on page 62.

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

- "Task 1: Set Up a Schedule for Backups" on page 56
- "Task 2: Set Up a Group for Backup Clients" on page 56
- "Task 3: Set Up Policies for Quick Access and Long Term Storage" on page 57
- "Task 4: Set Up a Label Template to Identify Volumes" on page 58
- "Task 5: Set Up a Pool to Sort Backup Data" on page 59
- "Task 6: Set Up Directives for Special Processing" on page 60
- "Task 7: Create a Backup Client Resource" on page 61

Note: If during the completion of any task in this section, the following error message appears:

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

it means that you do *not* have Sun StorEdge EBS permissions to make configuration selections. For information about getting permissions, see "Managing Server Access" on page 530.

Note: For specific information about backing up Sun StorEdge EBS clients on Mac OS X, see "Appendix F: Mac OS X Support" on page 869.

Task 1: Set Up a Schedule for Backups

A schedule can be applied to each client backup. For more information about schedules, see "Chapter 4: Backup Groups and Schedules" on page 133.

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, enter 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 StorEdge EBS server.
- Sort data to specific backup volumes. To sort data, groups are used in conjunction with backup pools.

For more information about groups, see "Chapter 4: Backup Groups and Schedules" on page 133.

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, enter a name for the group.
5. In the Comment attribute, enter a description of the group.
6. For the Start Time attribute, enter 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, enter the name of the printer on which bootstrap save set information will be printed.
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 StorEdge 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.

For more information about browse and retention policies, see "Chapter 5: Browse and Retention Policies" on page 169.

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, enter 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, enter 15 months.
5. In the Comment attribute, enter a comment about the policy.
6. In the Number of Periods attribute, enter the number of periods applied to the policy. For example, if you choose months for the Period attribute and three for the Number of Periods attribute, then the policy lasts for three 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 (three). 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

Every backup volume requires a unique label for identification. The Sun StorEdge EBS server creates a unique label for each volume by applying a label template. For more information about label templates, see "Chapter 7: Sorting Backup Data" on page 203.

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, enter a name for the label template.
5. In the Comment attribute, enter a description for the label template.

6. In the Fields attribute, enter 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 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. For more information about pools, see "Chapter 7: Sorting Backup Data" on page 203.

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, enter 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, enter 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. In the Groups attribute, select the groups whose backup data will be directed to the volumes belonging to this pool.

   Note: Pools also allow backups to be sorted by backup clients, save sets, and backup levels. For more information about sorting criteria, see "How Sun StorEdge EBS Software Uses Media Pool Criteria to Sort Data" on page 205.

10. Click OK.

Task 6: 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. For more information about directives, see "Chapter 6: Directives" on page 185.



Important: 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. For more information, see "Preconfigured Global Directive Resources" on page 192.

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, enter a name for the directive.
5. In the Comment attribute, enter a description for the directive.
6. In the Directive attribute, enter the directive instructions. For example, to skip all files on */tmp*, enter the following:

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

7. Click OK.

Task 7: Create a Backup Client Resource

A client is both a physical computer with Sun StorEdge EBS client software installed on it and a Sun StorEdge 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 StorEdge EBS client computer can have several Client resources. 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.

For more information about multiple Client resources, see "Multiple Clients from the Same Computer" on page 583.

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, enter the hostname of the Sun StorEdge EBS client computer.
5. In the Comment attribute, enter a description of the client. If multiple Client resources are being set up for the same host, enter 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, enter 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 the following:

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

    Enter All to back up *all* client data, including ACLs (access control lists). For 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. For more information about the All save set and the SYSTEM or VSS SYSTEM save sets, see "Save Sets" on page 63.

    For information on how to back up a mapped drive, see "Backing Up a Mapped Drive" on page 99.

    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. For more information, see "Preconfigured Global Directive Resources" on page 192.

11. From the Group attribute, select a group from the list.
12. From the Schedule attribute, select a schedule from the list.
13. 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.

Using the Sun StorEdge EBS Client Configuration Wizard

The Sun StorEdge EBS Client Configuration Wizard provides the ability to:

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

The wizard supports Sun StorEdge EBS servers and clients in a stand-alone or cluster environment.

Client Configuration Wizard Requirements

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

- The Sun StorEdge EBS client (**nsrexecd**) must be installed and running on the host from which the wizard is launched.
- The wizard user must:
  - Have root (UNIX) or Administrator (Windows) privileges.
  - Have Configure NW privileges on the Sun StorEdge EBS server where the scheduled backup is to be configured.
- The wizard host must be listed in the *servers* file on the client machine that is being configured for a scheduled backup.
- Multiple wizard hosts cannot access the same client machine at the same time.

Accessing the Client Configuration Wizard

To access the Client Configuration Wizard:

1. From the Administration window, click Configuration.
2. In the Configuration window, click Clients.
3. From the Configuration menu, select Launch Client Wizard.

For complete information about scheduling a client backup, see "Setting Up a Scheduled Backup" on page 55.

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 StorEdge 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.

A save set is defined when a Client resource is created. For information about creating a Client resource, see "Task 7: Create a Backup Client Resource" on page 61.

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 enter the names of predefined save sets when configuring the Sun StorEdge 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, the following 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, the following predefined save sets are available:

- All
- VSS SYSTEM BOOT
- VSS SYSTEM FILESET
- VSS SYSTEM SERVICES
- VSS USER DATA
- VSS OTHER
- VSS ASR DISK

For Windows NT 4.0, the following predefined save sets are available:

- All
- SYSTEM STATE
- REPAIRDISK



Important: To properly protect Sun StorEdge EBS client computers, all of the SYSTEM or VSS SYSTEM save sets must be backed up and recovered at the same time. Failure to do so will yield unpredictable results.

For more information about SYSTEM, VSS SYSTEM, SHAREPOINT, and ASR save sets, see "Appendix A: SYSTEM and VSS SYSTEM Save Sets" on page 801.

For more information about Sun StorEdge EBS support for Sharepoint, see "Chapter 23: Support for Microsoft SharePoint Portal Server" on page 721.

For more information about Sun StorEdge EBS support for ASR, see "Chapter 24: Support for Microsoft Automated System Recovery" on page 731.

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 11 on page 65 provides a list of the save sets included in the All save set.

Table 11. Components in the All Save Set (Part 1 of 2)

Operating System	Files
UNIX	Backs up all file systems listed in the following locations: • Solaris and IBM DYNIX / ptx: */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 11. Components in the All Save Set (Part 2 of 2)

Operating System	Files
Windows Server 2003 with VSS licensed and enabled	VSS SYSTEM BOOT VSS SYSTEM FILESET VSS SYSTEM SERVICES VSS USER DATA VSS OTHER VSS ASR DISK All local, physical drives
Windows 2000	SYSTEM STATE SYSTEM DB SYSTEM FILES SHAREPOINT ASR All local, physical drives
Windows NT 4.0	SYSTEM STATE (includes the registry) REPAIRDISK All fixed disks
Mac OS X	All local and mounted volumes

Ad Hoc Backups

The Sun StorEdge EBS server backs up client data regularly by using scheduled backups. It is the preferred backup method because the backup occurs automatically, and data can be recovered more easily. You can also start a scheduled backup at any time. For information about starting a scheduled backup group, see "How to Start a Group Immediately" on page 478.

By comparison, ad hoc backups enable users to make quick backups of a few files. Ad hoc backups are more difficult to recover because they do not generate bootstrap files or backup indexes.

How to Perform an Ad Hoc Backup on UNIX

To start an ad hoc backup on UNIX:

1. Start the **nwbackup** program by entering the following command:

   nwbackup

The **nwbackup** program provides a graphical interface to perform ad hoc backups. For general information about the **nwbackup** program, see "The Ad Hoc Backup Program" on page 42.

2. Select the directories and / or files and select Mark from the Mark menu.

   To clear an item, select Unmark from the Mark menu.

3. Click Start to begin the ad hoc backup.

   When the Sun StorEdge EBS server has successfully finished the backup, the **nwbackup** program displays a message similar to the following:

   ```
   Backup completion time: 2-15-06 3:27p
   ```



Important: Certain types of corrupt files or errors on computer disk volumes are not detected. Sun StorEdge EBS may back up this corrupt data. To avoid this situation, run diagnostic programs regularly.

How to Perform an Ad Hoc Backup on Windows

There are considerations to be aware of when performing an ad hoc backup of SYSTEM or VSS SYSTEM save sets. For more information, see "Considerations for Ad Hoc Backups of the SYSTEM and VSS SYSTEM Save Sets" on page 69.

To start an ad hoc backup on Windows:

1. In the Sun StorEdge EBS User program, click Backup.

   For general information about the Sun StorEdge EBS User program, see "Windows Client Interface" on page 39.

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.

   Note: If performing a Sun StorEdge EBS User backup on a Sun StorEdge EBS server, see "Excluding File Type Devices from an Ad Hoc Backup on Windows" on page 68.

4. Click Start to begin the ad hoc backup.

   The Backup Status dialog box displays the progress of the backup.

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

```
Backup completion time: 2-15-06 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 StorEdge EBS log file.



Important: Certain kinds of corrupt files or errors on computer disk volumes are not detected. Sun StorEdge EBS may back up this corrupt data. To avoid this situation, run diagnostic programs regularly to correct disk volume errors.

Excluding File Type Devices from an Ad Hoc Backup on Windows



Important: When performing a Sun StorEdge EBS User backup on a Sun StorEdge EBS server that is backing up to a file type device, do *not* include the file type device in the backup. The backup file would grow until there is no more disk space.

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

1. Start the Sun StorEdge EBS 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.

For more information about file type devices, see "Chapter 11: Device Operations" on page 297.

Considerations for Ad Hoc Backups of the SYSTEM and VSS SYSTEM Save Sets

This section discusses considerations to be aware of when performing an ad hoc backup of the SYSTEM or VSS SYSTEM save sets. These save sets are used to back up Windows system files. For more information about the SYSTEM and VSS SYSTEM save sets, see "Appendix A: SYSTEM and VSS SYSTEM Save Sets" on page 801.

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

Ad Hoc Backups of the SYSTEM and VSS SYSTEM Save Sets from the Sun StorEdge EBS User Program

In the Sun StorEdge EBS User Backup window, each of the SYSTEM or VSS SYSTEM save sets appears as a distinct node in the left pane. Expanding any of these nodes reveals its components in the right pane, as shown in Figure 11 on page 69 and Figure 12 on page 70.

Figure 11. SYSTEM Save Sets in Sun StorEdge EBS User Backup Window

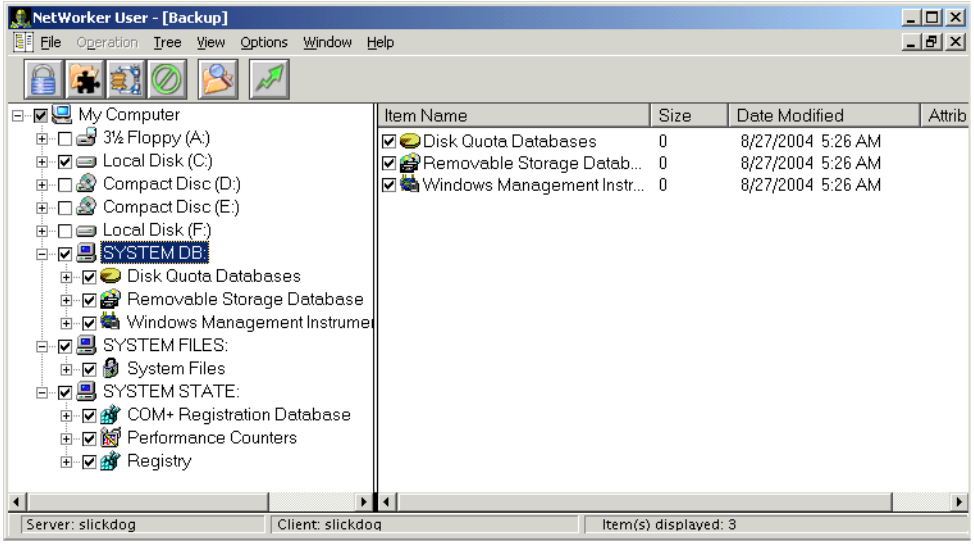
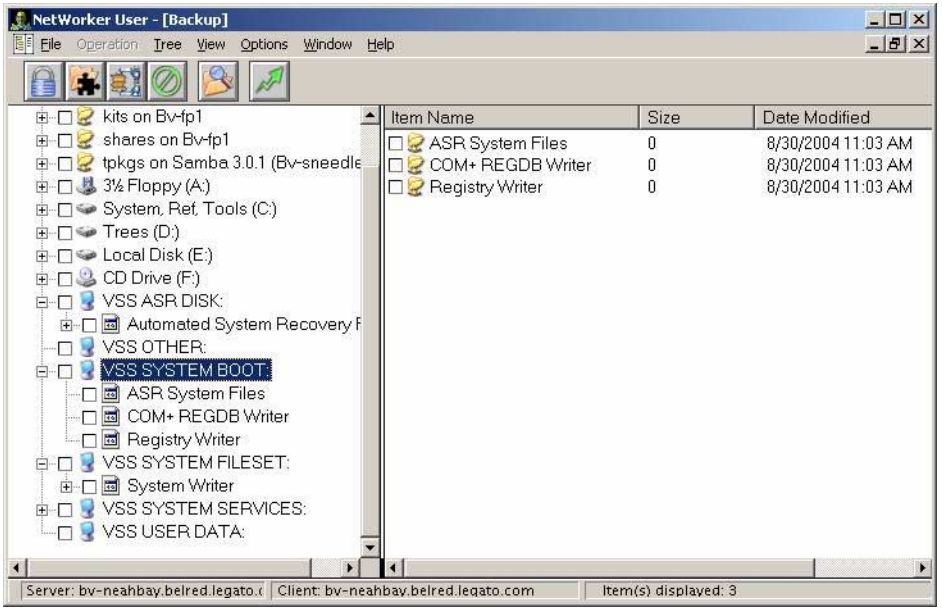


Figure 12. VSS SYSTEM Save Sets in Sun StorEdge EBS User Backup Window



The nodes that appear in the right pane reflect save set components that were eligible for backup at the time the Sun StorEdge EBS User program was started. After the Sun StorEdge 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 Sun StorEdge EBS User program starts.

Note: The Sun StorEdge EBS 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.

How to Perform an Ad Hoc Backup from the Command-prompt

An ad hoc backup can also be performed from the command-prompt by using the **save** command. For example, to back up */myfile* to the server *jupiter*, enter the following:

```
save -s jupiter /myfile
```

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

Ad Hoc Backups of the SYSTEM and VSS SYSTEM Save Sets from the Command-prompt

An ad hoc 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 StorEdge EBS client 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 StorEdge EBS_server_name*] **"SYSTEM STATE:"**

save [**-s** *Sun StorEdge EBS_server_name*] **"SYSTEM FILES:"**

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

save [**-s** *Sun StorEdge EBS_server_name*] **"SYSTEM DB:"**

save [**-s** *Sun StorEdge EBS_server_name*] **"SHAREPOINT:"**

save [**-s** *Sun StorEdge EBS_server_name*] **"ASR:"**

With VSS Enabled

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

To back up all components of the Windows system state:

save [**-s** *Sun StorEdge EBS_server_name*] **"VSS SYSTEM BOOT:"**

save [**-s** *Sun StorEdge 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 StorEdge EBS_server_name*] **"VSS SYSTEM SERVICES:"**

save [**-s** *Sun StorEdge EBS_server_name*] **"VSS USER DATA:"**

save [**-s** *Sun StorEdge EBS_server_name*] **"VSS OTHER:"**

save [**-s** *Sun StorEdge EBS_server_name*] **"VSS ASR DISK:"**

Requirements

Do not select individual components of any of the SYSTEM or VSS SYSTEM save sets for backup.

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 Sun StorEdge EBS log file.

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

Verifying Backed-Up Data

Note: This feature is not available on UNIX clients.

Use the Sun StorEdge EBS Verify feature to ensure that backup data on the Sun StorEdge 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 StorEdge 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 StorEdge EBS server from completing a successful backup. The Verify command 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 Sun StorEdge EBS 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.

Password Protecting Windows or NetWare Backup Data

The password-protection described in this section is based on the **pw1** Application Specific Module (ASM). PW1 password-protection can be used with backup data that is encrypted with the **pw2** ASM. PW1 password-protection and PW2 encryption is available for Windows and NetWare clients only.



Important: For information about password-protection and encryption that can be used with UNIX clients, see "Implementing Encryption with the AES ASM" on page 75.

Setting Up PW1 Password-Protection

Setting up password-protection is a two-part process:

- "Task 1: Enable PW1 Password-Protection" on page 73
- "Task 2: Apply PW1 Password-Protection" on page 74

Task 1: Enable PW1 Password-Protection

To enable PW1 password-protection on Windows:

1. In the Sun StorEdge EBS User program, select Password from the Options menu to open the Password dialog box.
2. In the Password attribute, enter a password (cannot exceed 20 characters).
3. In the Retype Password attribute, reenter the same password and click OK.

To remove an existing password:

1. Delete the contents of the Password and Retype Password text boxes.
2. Click OK.



Important: To recover a password-protected file, you *must* provide the password that was in effect when that version of the file was backed up. Keep password changes to a minimum.

Task 2: Apply PW1 Password-Protection

Password-protection can be applied to a scheduled backup or an ad hoc backup.

Applying PW1 Password-Protection to a Scheduled Backup

To password-protect data for a scheduled backup:

1. Create a local or global directive that includes the **pw1** ASM.
2. Apply the directive to the client save set.

   For information about creating and applying local or global directives, see "Chapter 6: Directives" on page 185.

   Note: When you password-protect a data item for a scheduled backup, password-protection is, by default, applied to the data item for an ad hoc backup. To remove password-protection for an ad hoc backup, see "Special Data Handling for Sun StorEdge EBS Clients on Windows" on page 78.

Applying PW1 Password-Protection to Ad Hoc Backups

To password-protect data for an ad hoc backup on Windows:

1. Mark the data to be password protected.
2. From the File menu, select the Special Handling option in the Sun StorEdge EBS User program.

   For more information, see "Special Data Handling for Sun StorEdge EBS Clients on Windows" on page 78.

Encrypting Backup Data

Encryption is implemented with Application Specific Modules (ASMs). Use either one of the following ASMs to encrypt backup data:

- **aes** ASM (UNIX and Windows only)
- **pw2** ASM (Windows and NetWare hosts only)

Important: To encrypt backup data, use either AES or PW2 encryption. Do not use both types of encryption for the same backup data. The type AES encryption works for UNIX and Windows host file systems. The PW2 type encryption works for Windows and NetWare hosts only.

Implementing Encryption with the AES ASM

Backup and archive data on UNIX and Windows hosts can be encrypted with the **aes** 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: AES data encryption is supported in the Sun StorEdge EBS 7.3 release. AES encrypted data cannot be recovered by using Sun StorEdge EBS 7.2.*x* or earlier.

How to Set the AES Pass Phrase for a Sun StorEdge EBS Server

To set the AES 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 enter a pass phrase in the Datazone pass phrase attribute.
5. Click OK.

How to Apply AES Data Encryption to Clients in the Datazone

To implement AES encryption:

1. Create a global directive that includes the **aes** ASM. For example, to encrypt all backup data in a folder, the **aes** ASM would be written as:

   ```
   << / >>
   +aes: *
   ```

   For more information about ASM syntax and creating a global directive, see "Creating a Global Directive Resource" on page 186.

2. Apply the global directive to individual clients by using the Directive attribute of the Client resource. For information about editing a Client resource, see "How to Edit a Client" on page 568.

Implementing Encryption with the PW2 ASM

Encryption with the **pw2** ASM is available for Windows and NetWare clients only. The PW2 encryption feature encrypts selected data items so that the data can be recovered with the password.

To apply encryption with the **pw2** ASM:

1. Ensure that the Sun StorEdge EBS User program's password feature is enabled. For information about enabling password-protection, see "Task 1: Enable PW1 Password-Protection" on page 73.

2. Apply encryption.

   To apply PW2 encryption to scheduled backups:

   a. Create a local or global directive that includes the **pw2** ASM.

   b. Apply the directive to the client save set.

      For information about creating and applying local or global directives, see "Chapter 6: Directives" on page 185.

      Note: When you encrypt a data item for a scheduled backup, encryption is also applied to the data item if it is selected for an ad hoc backup. To remove encryption for an ad hoc backup, see "Special Data Handling for Sun StorEdge EBS Clients on Windows" on page 78.

   To apply PW2 encryption to an ad hoc backup:

   a. From the Sun StorEdge EBS User program, mark the data to be encrypted.

   b. From the File menu, select the Special Handling option.

For more information, see "Special Data Handling for Sun StorEdge EBS Clients on Windows" on page 78.

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. *I*f 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. For information about creating and applying directives, see "Chapter 6: Directives" on page 185.

Applying Compression to an Ad Hoc Backup

To compress data for an ad hoc backup on Windows:

1. Mark the data to be compressed.
2. Select the Special Handling option from the File menu in the Sun StorEdge EBS User program.

   For more information, see "Special Data Handling for Sun StorEdge EBS Clients on Windows" on page 78.

To compress data for an ad hoc backup on UNIX:

1. Start the **nwbackup** program by entering the following command:

   nwbackup

2. Mark the data to be compressed.
3. Click Start.
4. On the Backup Options dialog, select the attribute titled, Do you want to compress this backup?
5. Click OK.

Special Data Handling for Sun StorEdge 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 Sun StorEdge EBS 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.

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

Backing Up Sun StorEdge EBS Console Management Data

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

Scheduling Backups for the Sun StorEdge EBS Console Database

If a Sun StorEdge EBS server was specified during the installation of the Console server, a Client resource was created to back up the Sun StorEdge 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 StorEdge 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 StorEdge EBS server attribute, enter the name of the Sun StorEdge 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, enter the name of the Console server.
6. Click OK.

   A Client resource will be 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 name of the computer 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 StorEdge EBS server to back up the Console server database through the Console Configuration Wizard. For more information, see "Using the Console Configuration Wizard" on page 514.

For more information about creating and tailoring a Client resource, see "Setting Up a Scheduled Backup" on page 55.

Note: Only full, incremental and skip backup levels are supported. All other backup levels (1, 2, through to 9) are mapped to an incremental backup.

Performing an Ad Hoc Backup of the Console Database

Note: Before performing an ad hoc backup on UNIX, add the following path to the LD_LIBRARY_PATH environment variable on the Console server:

"*Console_install_dir***/sybasa/lib"**

To perform an ad hoc backup of the Sun StorEdge EBS Console database, enter the following at the command-prompt:

savepsm -I "*Console_install_dir***"** *save options*

where *Console_install_dir* is the installation directory for the Console server. For example:

- On Solaris, the default installation directory is */opt/LGTOnmc*
- On Linux, the default installation directory is */opt/lgtonmc*
- On Windows, the default installation directory is *C:\program files\legato\management\GST*

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

savepsm -I "`C:\program files\legato\management\GST`**"** *save options*

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

Managing the Size of the Console Database Log File

The Console database transaction log files are automatically truncated whenever a schedule or an ad hoc backup of the Console database is performed. To truncate the transaction log file without performing a backup, run 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 an ad hoc 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.

   For more information about editing a backup schedule, see "How to Edit a Schedule" on page 151.

When performing an ad hoc backup, a full backup is performed by default. For information about ad hoc backups, see "Performing an Ad Hoc Backup of the Console Database" on page 79.

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 StorEdge EBS backup and recovery of mount points require special handling, as explained in the following procedures.

Including Mount Points in Scheduled Backups

If you specify *All* in the client's Save Set attribute, all data and any mount points on that client are included in scheduled backups. To include mount points and their data in scheduled backups without specifying save set *All*, 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 the following 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 the following in the client's Save Set attribute:

D:*mountpoint_name1*
D:*mountpoint_name1**mountpoint_name2*
D:*mountpoint_name1**mountpoint_name2**mountpoint_name3*

For more information about configuring scheduled backups, see "Scheduled Backups" on page 53.

Performing an Ad Hoc Backup of a Mount Point and Its Data

To back up a mount point and its data:

1. Start the Sun StorEdge EBS 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.

For more information about performing ad hoc backups, see "Ad Hoc Backups" on page 66.

Performing an Ad Hoc Back Up of Nested Mount Points and Their Data

To perform an ad hoc 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 Sun StorEdge EBS 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 select *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:*, *mountpoint_name1*, and *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 above

The back up 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.



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

Before a CIS backup, the Sun StorEdge EBS software does 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

In case of a problem with a CIS backup, do the following:

- 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.

Backing Up CIS on Windows NT 4.0

With CIS 2.0, certain files in the *catalog.wci* folder are not backed up unless the CIS service is stopped before the backup begins.

To back up CIS and its catalog files on a Windows NT 4.0 computer:

1. From the System Control Panel, stop the Content Index Server.
2. Start the Sun StorEdge EBS backup.
3. When the backup is complete, restart CIS.

Directing Sun StorEdge EBS Software to Skip CIS Catalog Backups

The Sun StorEdge EBS software can skip backing up the *catalog.wci* folder entirely. To do so, create a Directive resource in the Sun StorEdge EBS Console by entering the following text:

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

For more information about directives, see "Chapter 6: Directives" on page 185.

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 StorEdge EBS SYSTEM DB save set. However, you can configure Sun StorEdge EBS software to protect these databases by using the following procedures.

Note: If VSS is licensed and enabled, the DHCP and WINS databases are automatically included when performing a backup of save set All, so the following procedures are optional.

For information about scheduling and recovering backups, see "Scheduled Backups" on page 53 or "Chapter 13: Recovering Data" on page 367.

How to Back Up a DHCP Database

To back up a DHCP database, ensure that the following directory is included the save sets for the Sun StorEdge EBS client that is the DHCP server:

%SystemRoot%\System32\dhcp

How to Back 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 StorEdge EBS client that is the WINS server.

Note: For information about the Microsoft WINS administrative tools, refer to the Microsoft documentation.

Windows Backup and Recovery Notes

This section contains notes about the backup and recovery of data on Windows clients.

Short Filename Support

On Windows Server 2003 and Windows XP Professional, Sun StorEdge 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 StorEdge EBS client:

1. From the Administration interface window, click Configuration.
2. In the left pane, click Clients.
3. Right-click on 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 StorEdge EBS clients cannot save or restore these system-assigned short filenames when running on Windows 2000 or Windows NT 4.0.

Hard Link Support

The Sun StorEdge EBS server backs up and recovers files with hard links. However, the hard links of files created by using a 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 StorEdge EBS client:

1. From the Administration interface window, click Configuration.
2. In the left pane, click Clients.
3. Right-click on 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 StorEdge 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 filename *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*.

Failed Backup and Recovery Attempts

The Sun StorEdge EBS log file, located in <*Sun StorEdge EBS_install_path*> \ *logs* \ *networkr.log*, contains a record of every file that was part of an attempted ad hoc backup or recovery. This file is overwritten with the next ad hoc backup or recovery, so rename the *networkr.log* file if it contains information that you want to save.

Permissions Requirements for Disk Quota Database Backup

Sun StorEdge 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 StorEdge EBS client that uses the Windows Disk Quota feature, during SYSTEM DB or VSS SYSTEM SERVICES backup, Sun StorEdge 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 the following 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 StorEdge EBS client host computer.
2. Using Windows Explorer, perform the following 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 StorEdge EBS_install_path*> \ *logs* directory.

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

1. In Windows Explorer, navigate to the <*Sun StorEdge EBS_install_path*> \ *logs* directory.
2. Right-click the <*Sun StorEdge 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 StorEdge EBS operations performed by members of the Windows Backup Operators group will now be logged.

Customizing the Backup Command

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

1. Shuts down a mail server or database *before* the Sun StorEdge 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 either 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. See "Using the save Command with a Customized Backup Program" on page 88.
- By entering **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. See "Using the savepnpc Command with a Customized Backup Program" on page 94.

Consider the following questions to determine the best level of customization:

- Is there client data that does *not* need to be backed up every time (for example, company email)?
- Should the Sun StorEdge EBS server send special messages (in addition to the Savegroup Completion Reports) about the backups it executes?

Using the save Command with a Customized Backup Program

Additional processing instructions can be specified by entering the name of a 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.



Important: 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 program 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 program or batch file must adhere to the following criteria:

- The program name must begin with either the prefix *save* or *nsr* and must not exceed 64 characters.
- The program must reside in the same directory as the Sun StorEdge EBS **save** command.
- Always specify the full path of the **save** command in the script.
- The Sun StorEdge 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 StorEdge EBS server cannot complete the remaining instructions.
- When you invoke the Sun StorEdge EBS **save** command, invoke the command with the following arguments: **save** "**$@**". Doing so enables the **save** command in the batch file to accept the arguments usually passed to it by the Sun StorEdge EBS **savefs** program during a routine backup operation.

How to 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 StorEdge 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 the following order:

a. Run a preprocessing command before each save set backup (optional).

b. Back up the data by using the Sun StorEdge 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, enter the name of the backup program.

3. Back up the client to ensure that the newly created backup command works.



Example: The save Backup Command on Windows

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

The backup program consists of the following parts:

- Pre-Backup Command. Uses 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 StorEdge EBS commands to run the required backup process.
- Post-Backup Command. Uses the **Set** DOS command to create 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 StorEdge EBS_install_path*> \ *logs* directory. For information about viewing execution file details from the Monitor Groups tab, see "Monitoring Window" on page 475.

The following is an example of a backup script:

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

ECHO =====Sun StorEdge 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 StorEdge 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 StorEdge EBS backup commands.

:save

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

ECHO =====Sun StorEdge 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
```

The following information is displayed from the Monitor Groups tab and logged to the *sevegrp.log* file after the backup process is done. This information verifies the execution of the three processes:

```
--- Successful Save Sets ---
:* jupiter:c:\inetpub =======START BATCH
FILE============
* jupiter:c:\inetpub ===Sun StorEdge 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 StorEdge 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 StorEdge 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: The save Backup Command on UNIX

The following script backs up a ClearCase version object base (VOB). The script file must reside in the same directory as the Sun StorEdge EBS **save** command (for example, on a Solaris system, the **save** program is installed in the */usr/sbin/nsr* directory). Enter 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.



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

The following 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 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. For more information about the **savepnpc** command, refer to the *NetWorker Command Reference Guide* or the UNIX man pages.

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, enter **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 StorEdge EBS_install_path>\res* (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 the following default backup type, **preprocessing**, **postprocessing**, **timeout**, and **abort precmd** attributes:

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

Note: The **abort precmd with group** attribute is available in UNIX only. This command 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 the **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 any preprocessing commands listed for the **precmd** attribute in the *<group-name>.res* file. Then it performs the save by using the options specified for the **savepnpc** command itself. After the last save has completed, the **savepnpc** program performs any postprocessing commands listed for the **pstcmd** attribute.

When editing a *<group-name>.res* file, the following 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.
- 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*.

  The following 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.
- Text written to standard output:
  - Text written during preprocessing appears in the Sun StorEdge 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 *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.

How to 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 StorEdge EBS_install_path> \ tmp \ <group-name>.res to *<Sun StorEdge 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 StorEdge EBS_install_path> \ tmp \ <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 the following location on the Sun StorEdge EBS client:

- UNIX: */nsr/logs*
- Microsoft Windows: *<Sun StorEdge EBS_install_path> \ logs*

The format of the *savepnpc.log* file is similar to the following:

```
09/03/06 13:56:43 preclntsave: All command(s) ran
successfully.

09/03/06 13:57:43 preclntsave: All savesets on the
worklist are done.
```

Considerations for Backing Up Raw Partitions

The Sun StorEdge 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. For more information, see "Precautions When Using rawasm to Back Up UNIX Raw Partitions" on page 201.

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 StorEdge EBS_server_name* **\\.***PhysicalDrive0*

save -s *Sun StorEdge EBS_server_name* **\\.***C:*

Backing Up a Mapped Drive

To back up a mapped drive, follow these considerations:

- To specify a drive to back up in either a scheduled or ad hoc 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*, enter:

  \ \ *jupiter* \ *accounts*

- For scheduled backups, do the following:
  - 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.

Chapter 3: Configuring Libraries and Devices

This chapter covers the following topics:

- "Devices and Libraries" on page 101
- "Storage Nodes" on page 102
- "Dynamic Drive Sharing" on page 108
- "Autodetection of Libraries and Devices" on page 117
- "Add (Configure) Libraries" on page 121
- "Library Reconfiguration" on page 125
- "Add (Configure) Stand-Alone Devices" on page 127
- "Deleting Libraries and Devices" on page 129
- "Library Slots" on page 130
- "Troubleshooting" on page 132

Devices and Libraries

Sun StorEdge EBS software automatically recognizes devices that are being used for backups and recoveries. The maximum number of configured devices for any Sun StorEdge 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. For more information about file type devices, advanced file type devices, and the Sun StorEdge EBS DiskBackup option, see "Sun StorEdge EBS DiskBackup" on page 266.

Extensive handling of devices through a graphical user interface is new in the Sun StorEdge EBS 7.3 release, as is the ability to autoconfigure the following:

- SCSI devices
- NDMP devices
- Silos

The details that describe a particular device can be viewed, and often changed, in the device's Properties. Select View > Diagnostic Mode to see the full range of settings and options available in the Properties window. In the Properties window, click Field Help for a description of the various fields and attributes on a given Properties tab.

The following 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 StorEdge 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 StorEdge EBS servers for which they have access permission.

Storage Nodes

Storage nodes (including the Sun StorEdge 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 StorEdge EBS server maintains the client file index and media database. With the Sun StorEdge EBS software, client data can be routed directly to a storage node's storage devices without the data first going to the Sun StorEdge EBS server. A storage node may be a client of the Sun StorEdge EBS server, although this is not a requirement.

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

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

Only users who have the Configure Sun StorEdge EBS privilege can add to or change the configuration of the Sun StorEdge EBS server, media devices, and libraries. For more information, see "Sun StorEdge EBS User Groups" on page 533.

Prerequisites

To operate the Sun StorEdge EBS software with storage nodes, the following requirements must be met:

- Sun StorEdge EBS servers running the 7.3 release are not backward-compatible with storage nodes running older releases of Sun StorEdge EBS software. Storage nodes must be at least at the Sun StorEdge EBS 7.3 release level.
- On UNIX systems, the following software must be installed on the storage nodes. For Sun Solaris, the packages must be installed in the order listed:
  - Sun StorEdge EBS client software
  - Sun StorEdge EBS storage node software
  - (Optional) EMC AlphaStor® software. This enables multiple Sun StorEdge EBS servers to share the storage node. The AlphaStor software must be used to manage the libraries, drives, and volumes on that node. 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 Legato Software Compatibility Guide* for more information.
- On Windows systems, the Storage Node Option must be installed. This installs both the Sun StorEdge EBS client and storage node software.

  If multiple Sun StorEdge EBS servers share the same devices, then EMC AlphaStor software must be installed. The AlphaStor software must be used to manage the shared libraries, drives, and volumes on that node. 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 Legato Software Compatibility Guide* for more information.

Licensing

Obtain and enter the following enabler codes on the Sun StorEdge EBS server to use the Sun StorEdge EBS software with storage nodes:

- An enabler code for each individual storage node (included on the storage node Enabler Certificate).
- An enabler code for each library or silo connected to each storage node (included in the library and silo Enabler Certificates).



Important: If the Sun StorEdge EBS storage node software is configured and used on a remote computer in evaluation mode, and the evaluation period expires before the storage node enabler code is installed, it will not be possible to recover backed-up data to the storage node, and library features are disabled. To recover that data, the remote volume must be moved to a local drive connected to the server, and the recovery performed from there.

Storage Node Configuration

After the required software and enabler codes are installed, each storage node device must be configured (added to the Sun StorEdge EBS server). This procedure is similar to the tasks required for other devices. See "How to Configure a Storage Node" on page 104.



Important: Do not attempt to convert an existing Sun StorEdge EBS server to a Sun StorEdge EBS storage node. The potential for conflict exists when the resource database, media database, and client file indexes from the retired server are merged into the databases and client file indexes of the new Sun StorEdge EBS server.

How to Configure a Storage Node

To configure a Sun StorEdge 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. Enter a storage node name in the Name field.
   b. Check a box (NDMP, SCSI, or silo) in the Type of Storage Node field. The default is scsi.

      If NDMP is selected, the Remote Username and Password must be entered.
4. In the Device Management area, configure the following:
   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.

      Enter only one SCSI target per line.
5. In the Remote Host area, enter 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.

   Only one remote user with password is allowed per storage node. This applies to NDMP and advanced file type devices (AFTDs). For more information about AFTDs, see "Sun StorEdge EBS DiskBackup" on page 266.
6. Click OK when done.

   The new storage node appears in the navigation tree.

Timeout Attribute for Storage Node Operations

An attribute named nsrmmd Control Timeout, which is set during Sun StorEdge EBS server configuration, configures the amount of time a Sun StorEdge EBS server waits for a storage node request to be completed.

The Nsrmmd Control Timeout attribute controls how long the server's **nsrd** daemon or service waits for a storage node request to be completed. If the timeout value is reached without completion of the request, the operation discontinues and an error message is logged.

This attribute is found in the Sun StorEdge EBS server's Properties, on the Media tab, and is viewable only in Diagnostic Mode. The default value assigned to Nsrmmd Control Timeout is 5 minutes. Any value (in whole minutes) within a range of 2 to 30 minutes can be specified.

Timeouts for Storage Node Remote Devices

Timeouts for 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, on the Advanced tab, when in Diagnostic mode.

The Storage Node Devices area of the tab includes the following attributes:

- Server NetWork Interface
- 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.

Storage Node Affinity

The choice of which Sun StorEdge 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 StorEdge 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 StorEdge 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.

Bootstrap Backup on a Storage Node

When the server's index and the bootstrap save set are backed up, the data writes to a device that is local to the Sun StorEdge 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 StorEdge EBS release 7.0 and later, all devices created on storage nodes (with the exception of 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. All storage nodes prior to Sun StorEdge EBS release 7.0 behave as though the Dedicated Storage Node attribute is set to the default No.



Important: 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.

Tips for Troubleshooting Storage Nodes

If a backup fails, the following message might be displayed:

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

The problem could be related to storage node affinity.

Possible reasons for receiving this message include:

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

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

To fix this problem, do one of the following:

- 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.

Dynamic Drive Sharing

Dynamic Drive Sharing (DDS) is a feature that provides Sun StorEdge EBS software with the ability to recognize shared drives. DDS enables Sun StorEdge EBS software to skip the shared drives that are in use and route the backups or recoveries to other available shared drives.

Introduction to DDS

DDS controls application requests for media and allows the Sun StorEdge 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 StorEdge EBS system provides the following benefits:

- Reduces storage costs — A single tape drive can be shared among several storage nodes. In fact, since Sun StorEdge 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.

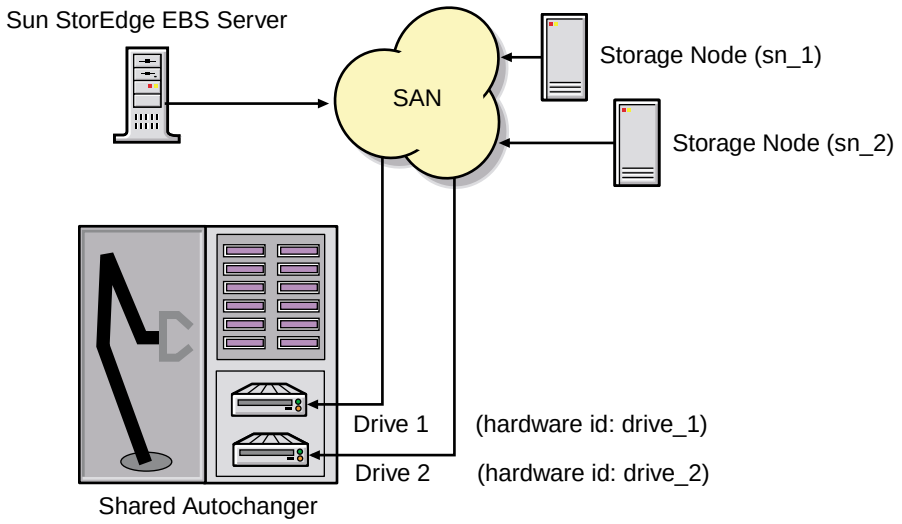
DDS Configuration Overview

Figure 13 on page 110 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, what is more important, can recognize when a shared drive is in use.

Under such a configuration, two DDS licenses are required, one for each drive.

Figure 13. Dynamic Drive Sharing



Block Size Compatibility

With DDS enabled, drives can be shared between storage nodes on different platforms, such as UNIX and Microsoft Windows. For Sun StorEdge 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 which is the smallest of all block sizes supported by all the storage nodes sharing the drive.

For more information about how to set block sizes for individual drives or tapes on different platforms, see "Compatible Block Size for UNIX and Microsoft Windows" on page 298.

Block Size Incompatibility between UNIX and Microsoft Windows

Incompatible block size settings between UNIX and Microsoft Windows storage nodes could result in any of the following 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 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 StorEdge EBS software will switch from using file and block positioning simply to reading the tape to reach the correct position. The data will probably 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, meaning 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 running 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 StorEdge 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 SCI 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 may be registered with a given drive at any given time, but only one may hold the active reservation. Sun StorEdge 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. For more information regarding CDI, see "Common Device Interface" on page 307.

Once a Reserve / Release setting (other than None) has been selected, for newer hardware, 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 the following 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 the following restrictions for this feature:

- It is available on CDI platforms only. Consequently, since CDI is not supported within an NDMP environment, Reserve / Release does not support NDMP or IRIX.
- It is supported only on a specific platform set and configuration. For details, refer to the Sun StorEdge EBS Release 7.3 *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 StorEdge EBS Reserve / Release feature to work.
- Even if all of the enterprise's Sun StorEdge 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

This section explains the requirements for 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.

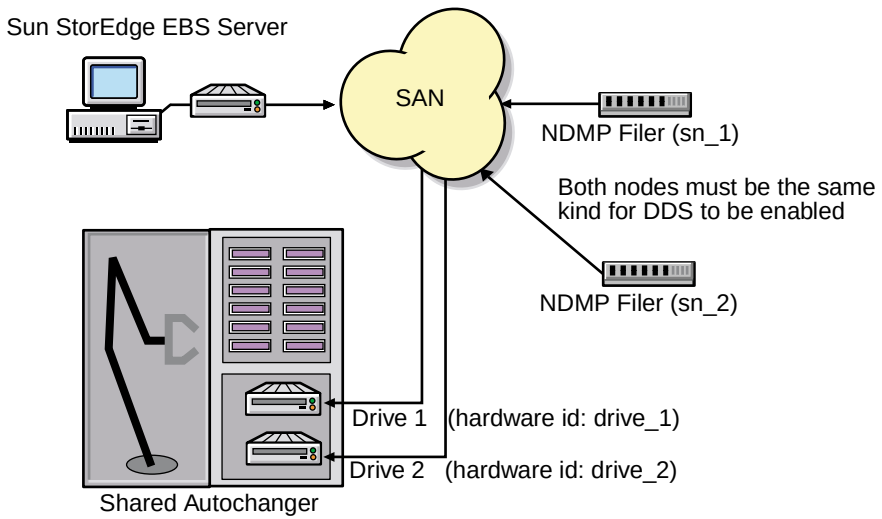
For example, DDS can be enabled in the following configurations:

- 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 preventing this.

Figure 14 on page 114 illustrates a basic DDS configuration with NDMP.

Figure 14. DDS with NDMP



DDS Attributes in the Device Properties

Attributes used in DDS are found in a device's Properties window:

- Hardware ID
- Shared Devices

Hardware ID Attribute

The Hardware ID tracks drives 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 can be added when manually configuring a device. It is not editable by users.

The Hardware ID can be viewed in a device's Properties window, 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 sharing the same hardware ID with a single action. Table 12 on page 115 lists this attribute's allowed values and their descriptions.

Table 12. Shared Devices Attributes (Part 1 of 2)

Value	Description
Enable All	When selected, enables all devices using the same hardware ID.

Table 12. Shared Devices Attributes (Part 2 of 2)

Value	Description
Disable All	When selected, disables all the devices 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 for 10 to 15 minutes 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: Selecting zero means that the device never times out and the tape must be ejected manually.

High Availability and DDS

The Sun StorEdge 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 StorEdge 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: Host Crash Requires User Intervention

The scenario is as follows: 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).

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 StorEdge EBS software be able to access the same path through which the volume was initially loaded.

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.

Autodetection of Libraries and Devices

Devices already known to the Sun StorEdge 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 StorEdge EBS server. Note the following:

- 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 StorEdge EBS server.
3. In the Name column of the Host detail table, double-click Sun StorEdge EBS.

The Sun StorEdge EBS Administration window for the selected server opens. Note that while multiple Sun StorEdge EBS Administration windows can be open simultaneously, each one displays information about only one host or server.

4. In the Administration window, click the Devices button.
5. In the navigation tree:
   a. Right-click the server name, and select Scan for Devices.
   b. Click the storage node to be scanned.

      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.
   c. Fill in any required information, such as whether to scan for SCSI or NDMP devices and whether to search all LUNs.
   d. Click Start Scan.

      To monitor the scan activity, click the Monitoring button, 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 StorEdge 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 StorEdge 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.

Configuration of 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 Legato *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.

The following 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. For information about how to add or delete a device, see "How to Add a Library Device" on page 122.



Important: 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

The following 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:devctl

Found StorEdge HBA device:
/devices/pci@1f,0/pci@1/pci@1/pci@f/SUNW,qlc@5/fp@0,0:devctl

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, the following 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, complete the following 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, 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. 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 have been created before devices can be configured to be used by them. See "Storage Nodes" on page 102 for 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.

Devices must be updated to the most recent firmware and drivers.

The following device types can be configured automatically:

- NDMP
- SCSI
- Silo

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

- AlphaStor
- RS (Removable Storage) Service
- Controlling IBM tape libraries using 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.

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 the Devices button.
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 that if Configure All Libraries is started from the server folder instead of from the Storage Node folder, then all storage nodes on the Sun StorEdge 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 the following 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).)

     Storage nodes to which libraries can be configured (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.

   – 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, the Back button becomes active and allows the user to return and adjust input.

How to Add (Configure) a Silo Device

To configure a new silo resource to a storage node automatically:

1. In the server's Administration interface, click the Devices button.
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 that if Configure All Libraries is started from the server folder instead of from the Storage Node folder, then all storage nodes on the Sun StorEdge 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 the following 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 to which to configure libraries (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.

– Update storage node properties if required.

a. 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.

b. Enter the Hostname of the silo controller.

c. Enter the Type of silo controller. The default is ACSLS Silo (StorageTek).

d. If a DAS silo controller was chosen, enter the DAS client name, which should match the storage node name.

e. The Test Silo Controller Connectivity button (optional) is useful to see whether the connection to a silo controller works. If used, 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, allowing the user to return to the input screen to adjust input.

Device Resource Queuing for AlphaStor

Because the Sun StorEdge 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 software with the **jbconfig** command. For information about configuring resource queuing, refer to the *AlphaStor Administrator'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 *Sun StorEdge EBS Command Reference Guide*.

Library Reconfiguration

Reconfigure a library in order to add or remove access paths to the drives in the library. This includes access paths that allow libraries to be shared.

The reconfiguring of stand-alone or file type devices is not supported. Instead, delete the stand-alone or file type device, then create a new one.

How to Reconfigure 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 the Devices button.
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.

   (It is also possible to open the Storage Nodes folder, open the library folder, and 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.

The jbedit Command

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

The **jbedit** program supports all direct-attached SCSI/SJI, SAN, and NDMP libraries. It does not currently support AlphaStor or silo 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 125. 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 StorEdge EBS privileges.

Using the jbedit Program

Table 13 on page 126 lists the most commonly used **jbedit** program options.

Table 13. 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.

For a detailed description of the **jbedit** command, its options, and associated diagnostic messages, refer to the *Sun StorEdge EBS Command Reference Guide* or the UNIX man page.

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. See "Storage Nodes" on page 102 for 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.

Devices must be updated to the most recent firmware and drivers.

How to Autodetect and Configure a Stand-Alone Tape Drive

It is also possible to configure a new stand-alone tape drive, automatically, by using Scan for Devices:

1. In the server's Sun StorEdge 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.

   Check the scan status by clicking the Monitoring button, then selecting the Log tab. Return to the Devices navigation tree.
4. 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.
5. Right-click the stand-alone drive to be configured, and select Configure Drive.

   A Configuration dialog box appears.
6. Click Yes to confirm that the drive should be configured.

   The new drive is automatically configured.

How to Add a Stand-Alone Device Manually

To configure (add) a new stand-alone device on a storage node:

1. In the server's Sun StorEdge EBS Administration interface, click the Devices button.
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*.
4. Also in the Identity area, configure the following:
   a. Add an optional, descriptive comment in the Comment field.
   b. In the Media Type field, select a media type.
5. In the Status area, configure the following checkboxes, if applicable:
   a. Read Only
   b. Auto Media Management
6. In the Cleaning area, configure the following, if applicable:
   a. Cleaning Required
   b. 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 the following:
   – 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.

For specific instructions and restrictions on backing up to a file type device, see "Sun StorEdge EBS DiskBackup" on page 266.

Configuring Sun StorEdge EBS DiskBackup

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

Sun StorEdge EBS software offers two variations of this method:

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

For information about the Sun StorEdge EBS DiskBackup solution and these device types, see "Sun StorEdge EBS DiskBackup" on page 266. For how to configure file type and advanced file type devices, see "Sun StorEdge EBS DiskBackup Configuration" on page 287.

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 the Devices button.
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.

   An error message appears, asking the following:

   ```
   "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.

A library's devices remain, and can still respond to Sun StorEdge EBS operations (such as monitoring, labeling, deletion, and so on) after the library definition is deleted. Deletion of a library deletes the library, not its devices. This behavior is the same as in previous releases.

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 the Devices button.
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 that 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 StorEdge 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.

The entries can be a range of slot numbers or a single slot number. The range of available slots entered must be the same for all of the storage devices in the library. It is not possible to enter a different range for each device located in an library.

Important: Ensure that volumes have been placed in all the available slots of the library so that the Sun StorEdge EBS server can proceed uninterrupted with an automatic backup.

Note: 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 StorEdge EBS Administration interface, select View>Diagnostic Mode from the menu bar.
2. Click the Devices button.
3. Open the Libraries folder in the navigation tree.

   The Libraries detail table appears.
4. In either the navigation tree or in the Library’s details 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 the plus sign (+) 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, do the following 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 henceforth be skipped when tapes are loaded.

Troubleshooting

This section contains troubleshooting information.

Autoconfiguration Failure

Common causes for autoconfiguration failure include the following if, after a library has been autoconfigured:

- The library is not listed.
- The library is listed, but is listed as being unconfigured.

Common causes of this could be:

- The device driver is 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:

- Check Monitoring > Log for relevant messages.
- From the command-prompt, run 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.

Chapter 4: Backup Groups and Schedules

This chapter covers the following topics:

- "Overview of Sun StorEdge EBS Scheduling" on page 133
- "Backup Groups" on page 134
- "Managing Backup Groups" on page 140
- "Bootstrap Report Generation and Printout" on page 142
- "Backing Up Open Files" on page 142
- "Schedules" on page 144
- "Backup Levels" on page 154
- "Save Set Consolidation" on page 163

Overview of Sun StorEdge EBS Scheduling

Together, the following two items enable the scheduled backup of client data:

- Group
- Schedule

Groups (backup groups) specify the time of day when a backup occurs. Typically, backup times occur after regular work hours. All clients assigned to a backup group will be backed up at the time specified by the backup group.

Schedules provide the ability to specify the day of the week or date during the month when the backup occurs. The type of backup, such as full, incremental, or levels 1-9 are also specified in the schedule.

Backup Groups

Backup groups specify the starting time for a client's scheduled backup. 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.
- Sort data to specific backup volumes. To sort data, groups are used in conjunction with backup pools. For more information, see "Chapter 7: Sorting Backup Data" on page 203.

How the Sun StorEdge EBS Server Uses Groups

When a Client resource is created, it is assigned to a backup group. The clients in each backup group begin their automatic scheduled backups according to the start time of the group. Balance the backup loads by taking the client's backup schedule into account when determining which clients to include in a specific group.



Example: Using Groups to Balance Client Backups

Figure 15 on page 135 illustrates how the Sun StorEdge EBS server uses two backup groups to back up multiple client save sets. In Figure 15 on page 135, 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.

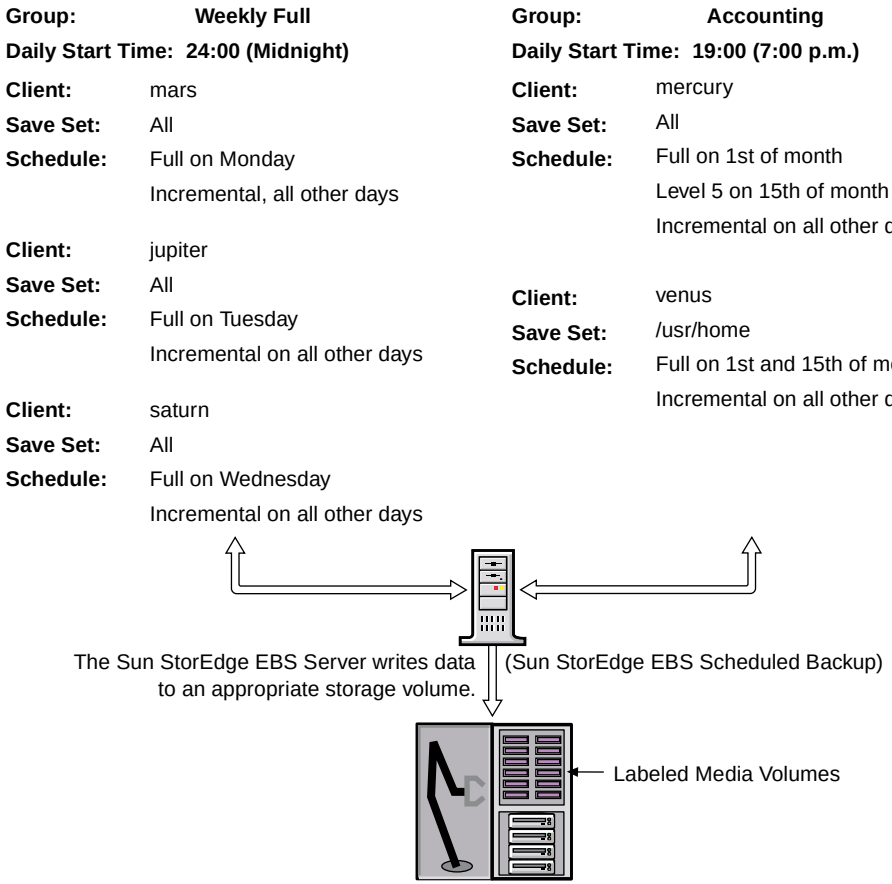
Because each client runs its full backup on a different day of the week, the server is not overloaded.

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 back up 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 what 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 StorEdge EBS server uses pools to organize, track, and store save sets. The server uses groups to determine what time clients start their scheduled backups.

Figure 15. How Sun StorEdge EBS Groups Are Used to Back Up Multiple Clients



Preconfigured Groups

The Sun StorEdge EBS product ships with a single preconfigured group named *Default*. To ensure that all data is backed up, the Sun StorEdge EBS server automatically adds all clients to the Default group. However, you must enable the Default group for the Sun StorEdge 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 14 on page 136 displays some of the key attributes for the Group resource.

Table 14. Group Attributes (Part 1 of 2)

Attribute	Action
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 the specified start time. You must enable the Autostart attribute for the Default group, and any other group you create, before a scheduled backup can be run.
Schedule	This optional attribute can be used to select a Schedule resource for the group. If this attribute is set, it overrides the schedules selected in the Client resource's Schedule attribute for all clients in this group.
Interval	This attribute dictates how often a group starts a scheduled backup. The default value is 24:00 (once a day), but you can change this value to start backups more often. For instructions on modifying this attribute, see "How to Set the Backup Group Time Interval" on page 139.
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 be restarted when the Sun StorEdge EBS server is restarted provided that the period of time specified in the Restart Window attribute has not elapsed.

Table 14. Group Attributes (Part 2 of 2)

Attribute	Action
Restart Window	If Autorestart is enabled, this attribute specifies the period of time in which an incomplete backup can be restarted. If the period of time specified in this attribute 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 StorEdge 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 longer than the Inactivity Timeout value, the server will consider the client as stopped responding. 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 or for save sets with large sparse files, increase the timeout value if the backup consistently aborts due to an inactive job.

Aborted Backup Groups

If the backup of a save set fails, then the Sun StorEdge 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 is 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 StorEdge EBS backup hangs until it times out. The Sun StorEdge EBS server marks the save set "aborted," and when the NFS server restarts, the backup continues and is completed.

How to Create a Group

For information about creating a group, see "Task 2: Set Up a Group for Backup Clients" on page 56.

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, enter a name for the new group.
5. Edit the attributes and click OK.

How to Set the Backup Group Time Interval

The Sun StorEdge 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 138.
2. Select the Advanced tab.
3. In the Interval attribute, enter a value in the hh:mm format.

   For best results, use time interval values that make it easy to determine the backup group time, such as 24, 12, or 6 hours.
4. Click OK.

Note: An increase in the backup group time interval (for example, changing the interval from once every 24 hours to once every 12 hours) can add strain on a network, the Sun StorEdge EBS server, and associated devices.

How to Limit 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. 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 StorEdge 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.

To force incremental backups on groups:

1. Select the group to edit. For information about editing a group, see "How to Edit a Group" on page 138.
2. Select the Force Incremental attribute.
3. Click OK.

For information about how to configure backup groups to occur more than once every 24 hours, see "How to Set the Backup Group Time Interval" on page 139.

Managing Backup Groups

Table 15 on page 140 lists backup group management tasks and where to find more information about these tasks.

Table 15. Backup Group Management

Group Management Task	For More Information
View backup group details	"How to View Group Backup Details" on page 479
Start a group immediately	"How to Start a Group Immediately" on page 478
Stop a group immediately	"How to Stop a Group" on page 478
Restart a group	"How to Restart Groups" on page 478
Preview a backup group	"How to Preview a Backup Group" on page 140

How to Preview a Backup Group

You can simulate a backup for a specific group. This feature outputs the following information:

- File System to be backed up
- Backup level
- Backup pool

Preview a backup to identify potential problems before the Sun StorEdge EBS server runs a backup group.

To preview a backup group:

1. Select the group to edit. For information about editing a group, see "How to Edit a Group" on page 138.
2. Select the Advanced tab.
3. In the Option attribute, select Preview.
4. Click OK.

To see the simulated results of the backup, open the *daemon.log* file located in the <*Sun StorEdge EBS_install_dir*> \ *logs* directory.

How to Estimate 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. For information about editing a group, see "How to Edit a Group" on page 138.
2. Select the Advanced tab.
3. In the Option attribute, select Estimate and No Save.
4. Click OK.

To see the estimated save set sizes, open the *daemon.log* file located in the <*Sun StorEdge EBS_install_dir*> \ *logs* directory.

Notes:

- 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.

Backup Status Reports

When the backup is completed, several types of backup status reports are generated. For more information about these reports, see "Chapter 14: Reporting" on page 413. Information about the status of backed-up groups is also written to the *savegrp.log* file, which is located in the *<Sun StorEdge EBS_install_dir> \ logs* directory.

Bootstrap Report Generation and Printout

When the backup group includes the Sun StorEdge EBS 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 printout is generated whether the scheduled backup is initiated automatically or manually. The bootstrap information is essential for recovery from a disaster. For information on how the bootstrap is used during a disaster recovery operation, refer to the *NetWorker Disaster Recovery Guide*.

Note: If the Sun StorEdge EBS server is not a member of an active group, the bootstrap is always created when any group backup is run, even if the group has a level of "skip."

The bootstrap is printed to the Sun StorEdge EBS server's default printer. To change the default printer, edit the Printer attribute in the Group resource. If the bootstrap report fails to print to the specified printer, its contents can be viewed in the *daemon.log* file, which is located in the *<Sun StorEdge EBS_install_dir> \ logs* directory.

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 loss of data. They might be skipped, improperly backed up, or locked, preventing user access.

There are two varieties of open files:

- Open files owned by the operating system
- Open files owned by a specific application

The following sections describe how the Sun StorEdge EBS server handles open files during a backup.

Open 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 StorEdge EBS software, from writing to or reading from the open file.

The Sun StorEdge EBS software normally skips locked files and returns the message:

`save:` *file_name* `cannot open`

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. For more information, see "Customizing the Backup Command" on page 88.

Open Files Owned by a Specific Application

The Sun StorEdge EBS software by itself cannot normally back up an open file belonging 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 StorEdge EBS software alerts you by displaying the following message in the Groups tab of the Monitoring option:

`warning:` *file_name* `changed during save`

To make sure the changed file is backed up, do *one* of the following:

- Restart the backup group.
- Perform an ad hoc backup. For information, see "Ad Hoc Backups" on page 66.

Note: NetWorker Modules can back up these types of files correctly.

Backing Up Open Files with VSS

In Sun StorEdge EBS releases 7.1 or later, if a VSS client license is present, Sun StorEdge EBS 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. For more information about VSS, "Chapter 25: Volume Shadow Copy Service" on page 745.

Schedules

Each Client resource is backed up according to a schedule. A Client resource's backup schedule tells the Sun StorEdge EBS server what level of backup (for example, full or incremental) to perform on a given day. For instance, on Fridays you might perform a full backup on a Client resource and the rest of the week perform incremental backups. The time of day the backup operation 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 backup *level* you specify. 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. For more information about backup levels, see "Backup Levels" on page 154.

Preconfigured Sun StorEdge EBS Schedules

The Sun StorEdge 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 16 on page 145 describes the preconfigured schedules.

Table 16. Preconfigured Sun StorEdge EBS Schedules

Schedule Name	Sun StorEdge 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.
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

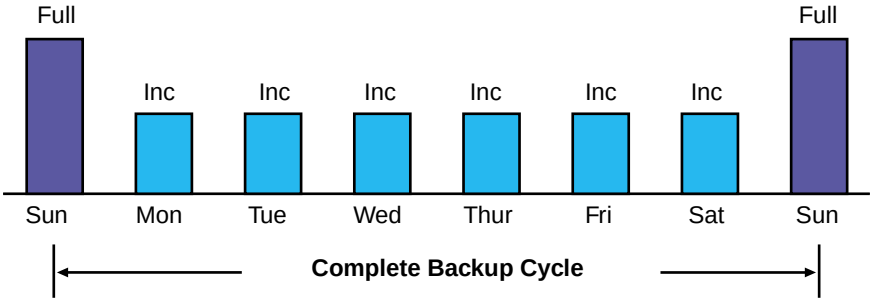
The period of time from one full backup to the next full backup is called a *backup cycle*. The following examples demonstrate how to use schedules for different backup cycles and client backup needs:

- "Example: Weekly Backup Cycle" on page 146
- "Example: Default Schedule with Multiple Clients" on page 147
- "Example: Staggered Weekly Schedules for Multiple Clients" on page 148

Example: Weekly Backup Cycle

Figure 16 on page 146 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 16. Weekly Backup Cycle



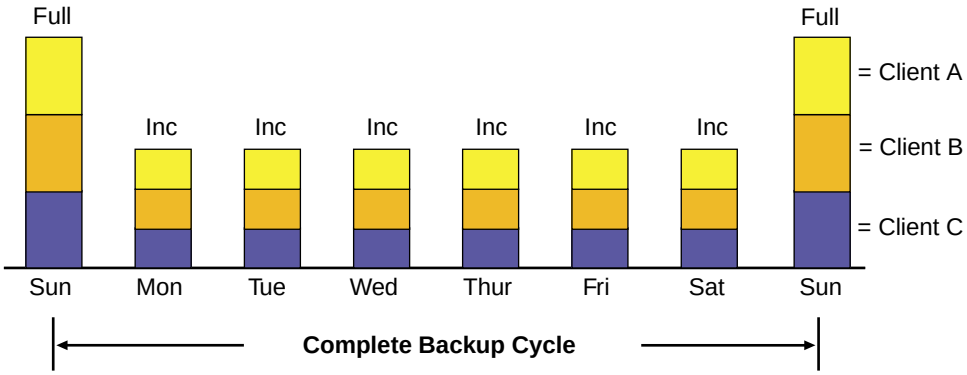
Use backup schedules to balance and stagger the load on a Sun StorEdge EBS server. Depending on the size of a network, you could apply the same schedule to all clients. For example, if no one works over the weekend and you want to run full backups during this time, you could apply the Default schedule to all of the clients. The Default schedule tells the Sun StorEdge EBS server to perform full backups on Sunday, and incremental backups the rest of the week.



Example: Default Schedule with Multiple Clients

Figure 17 on page 147 illustrates how the Default schedule works for three clients.

Figure 17. Using the Sun StorEdge EBS Default Schedule for Multiple Clients



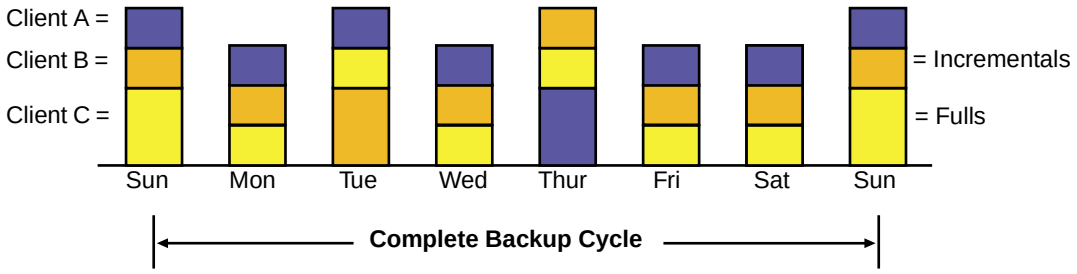
Note: If you have a short backup window period and need to create a full backup, consider using the consolidation backup. A consolidation backup uses the most volume space because it holds both a level 1 and a consolidated full backup. However, data recovery is faster because of the full backup. For more details, see "Save Set Consolidation" on page 163.

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: Staggered Weekly Schedules for Multiple Clients

Figure 18 on page 148 illustrates how to use a staggered backup schedule for multiple clients.

Figure 18. 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 StorEdge EBS server.

4

Schedule Planning Considerations

Deciding which schedules are most appropriate for an environment requires planning.

When you create backup schedules, consider the following:

- 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 only expect to be able to recover data from 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 StorEdge EBS server is determined by the browse and retention policies associated with each client. For more information about browse and retention policies, see "Chapter 5: Browse and Retention Policies" on page 169.

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.



Important: 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, enter the value All in the Save Set attribute.

Key Components of a Schedule

Table 17 on page 150 describes the key components of a Schedule resource.

Table 17. 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 the following levels in a backup schedule: • Full • Incremental • Level [1 – 9] • Consolidated • Skip The Overrides option in the pop-up menu 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. For more information on how backup levels work, see "Backup Levels" on page 154 and "Save Set Consolidation" on page 163.

Note: The Force Incremental attribute in the backup group, determines the level used by the Sun StorEdge 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. For more information, see "Limit Full Backups When the Time Interval Is Less Than 24 Hours" on page 139.

How to Create a Schedule

For information about creating a schedule, see "Task 1: Set Up a Schedule for Backups" on page 56.

To assign a backup schedule to a client, select an existing schedule in the Schedule attribute of the Client resource. For information about setting up a Client resource, see "Task 7: Create a Backup Client Resource" on page 61.

How to Edit 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.

How to Delete 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.

How to Copy 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, containing the same information as the schedule that was copied, except for the Name attribute.
4. In the Name attribute, enter 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 / 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. For more information, see "How to Edit a Client" on page 568.
2. Clear the Scheduled backup attribute.
3. Click OK.

Client Time Zone Adjustment for Windows Clients

The Sun StorEdge EBS software allows you to schedule backups over long-distance networks. In some cases, scheduling can occur over two or more time zones. To incorporate time zones into your backup schedules, adjust each client time zone setting by editing the client's *AUTOEXEC.BAT* file.

How to Adjust 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. Enter the following in the *AUTOEXEC.BAT* file:

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

Table 18 on page 153 describes the time zone strings that can be adjusted in the *AUTOEXEC.BAT* file.

Table 18. 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 StorEdge EBS software presumes a default setting of TZ = EST5EDT.

Backup Levels

Because it may not be practical or efficient for you to run full backups every day, you can specify the level of the backup to be performed during scheduled backups. Limiting how often full backups are performed helps 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 19 on page 154 describes the five kinds of backup levels:

Table 19. 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 of level 2, level 1, or full backup. For example, if the most recent backup was at level full, then a level 3 backup backs up all files that changed since the full backup. However, if the most recent backup was at level 2, then a level 3 backup backs up only those files changed since the level 2 backup. • A level 9 backs up all the files that have changed since the most recent backup of any level except level 9. **Note:** The Sun StorEdge 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. For more information, see "Save Set Consolidation" on page 163.
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.

How Sun StorEdge 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. For more information on directives, see "Chapter 6: Directives" on page 185.

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



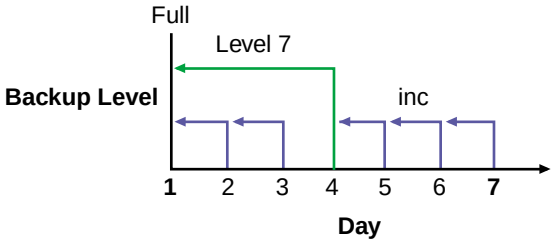
Example: Backup Levels (Part 1)

As shown in Figure 19 on page 156, 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 19. Backups for October 1 through October 7

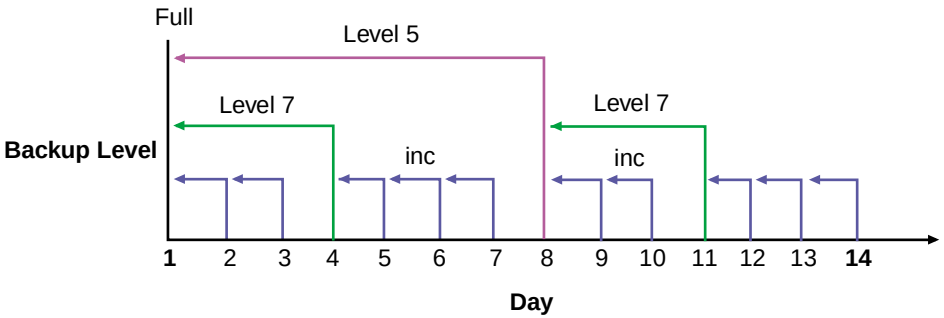




Example: Backup Levels (Part 2)

Figure 20 on page 156 continues the example illustrated in Figure 19 on page 156 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 20. Backups for October 1 through October 14

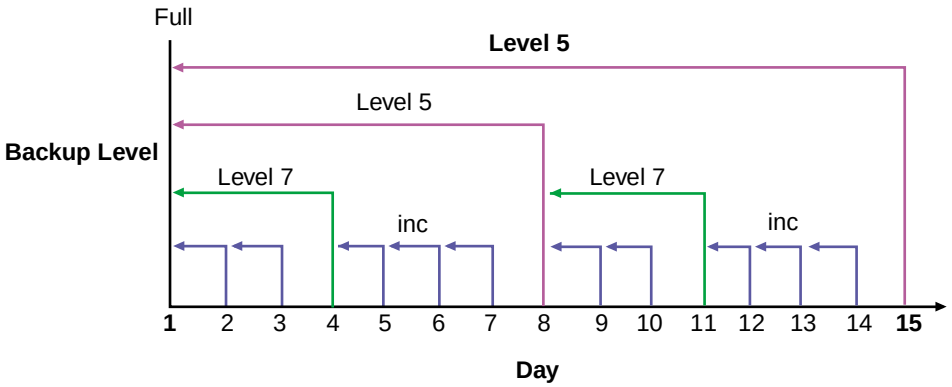


Also, a level 7 backup on October 11 backs up all of the data that changed since the last lower-numbered backup (in this case, the level 5 backup on October 8). To recover from a disk crash on October 11, you need three volumes: the full volume from October 1, the level 5 volume from October 8, and the new level 7 volume.

Example: Backup Levels (Part 3)

Figure 21 on page 157 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 21. Backups for October 1 through October 15



How the Sun StorEdge EBS Server Uses 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 the following to help you decide what best suits your environment:

- Full backups take more time to complete than incremental backups.
- If you have only one 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 20 on page 158 lists advantages and disadvantages of each backup level.

Table 20. 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 the following 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 StorEdge 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 StorEdge EBS server's backup is a level 5, the backups of the server index and media database is also a level 5.

  However, when the server's backup level is incremental, the backups of the server index and media database are 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 StorEdge EBS Online Indexes

The Sun StorEdge 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 21 on page 159 compares the level at which the Sun StorEdge EBS server backs up regular files and the online indexes.

Table 21. 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 StorEdge EBS software to help determine if the file should be backed up. The Sun StorEdge EBS software backs up a file if the Archive attribute is set. The function can disable this behavior by setting the environment variable NSR_AVOID_ARCHIVE to No. To set this as a system environment variable, use the Control Panel's System applet.

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

After backing up the file, the Sun StorEdge EBS software turns off the Archive attribute. After restoring the file, the Sun StorEdge 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. For more information about these save sets, see "Appendix A: SYSTEM and VSS SYSTEM Save Sets" on page 801.

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

SYSTEM Save Sets

Following 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 the following manner:

- As a single, indivisible unit. Therefore, components of the NetWorker SYSTEM STATE save set (such as COM+ and the registry) cannot be backed up or restored individually as separate entities.
- At level full. Therefore, individual components of the SYSTEM STATE save set cannot be backed up at any level other than full. Requesting an incremental backup of the SYSTEM STATE save set, for example, always results in a full backup of the save set.

SYSTEM FILES

As with the other SYSTEM save sets, when a backup of the SYSTEM FILES save set occurs, it is always at level full. However, a requested backup of the SYSTEM FILES save set is performed only if one or more of the system-protected files have changed since the specified date and time. If no system-protected files have changed, none will be backed up and no corresponding save set entry is made in the server's media index. However, on

an incremental or level 1-9 backup of the SYSTEM FILES save set, if *any* system-protected files have changed since the specified time, *all* system-protected files are backed up.

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

Following 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 the following manner:

- As a single, indivisible unit.

  Components of the NetWorker VSS SYSTEM BOOT save set (such as COM+ and the registry) cannot be backed up or restored individually as separate entities.
- At level full.

Individual components of the VSS SYSTEM BOOT save set cannot be backed up at any level other than full. Requesting an incremental backup of the VSS SYSTEM BOOT save set, for example, always results in a full backup of the save set.

VSS SYSTEM FILESET

As with the other VSS SYSTEM save sets, when a backup of the VSS SYSTEM FILESET save set occurs, it is always at level full. However, a requested backup of the VSS SYSTEM FILESET 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

This section describes the save set consolidation feature.

How Save Set Consolidation Works

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 StorEdge 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 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 the following 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 used by backups.

Limitations of Save Set Consolidation

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

- 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.



Important: 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 an ad hoc consolidation from the command-prompt.

How to Schedule Save Set Consolidation

To set up save set consolidation on a schedule:

1. Open the schedule to which save set consolidation will be applied. For more information, see "How to Edit a Schedule" on page 151.
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.

How to Perform an Ad Hoc 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 *NetWorker Command Reference Guide* or the UNIX man pages.

How to Direct 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. See "Task 2: Set Up a Group for Backup Clients" on page 56.
2. Create a Pool for consolidation backups. See "How to Create a Media Pool" on page 213.
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.

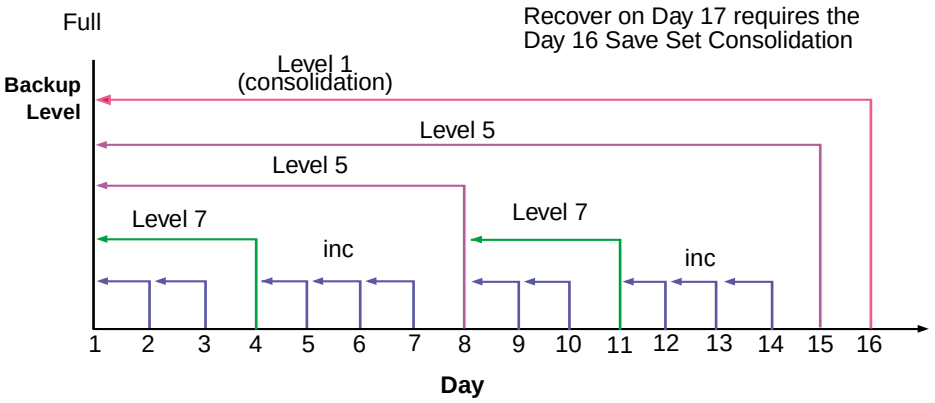


Example: Consolidation of Level 1 and Full Backups

Figure 22 on page 167 illustrates how the Sun StorEdge EBS server merges a level 1 backup with an existing full level backup to create another full level backup.

On October 16, the consolidated level 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 full level backup. To recover from a disk crash on October 17, use the volume that was created on October 16.

Figure 22. Consolidated Backup



Chapter 5: Browse and Retention Policies

This chapter covers the following topics:

- "About Browse and Retention Policies" on page 169
- "Managing the Data Lifecycle" on page 176
- "Browse and Retention Policies for Ad Hoc Backups" on page 182
- "Modifying the Browse and Retention Policy on a Save Set" on page 183

About Browse and Retention Policies

The *browse policy* determines how long files are maintained in the client's file index on the Sun StorEdge EBS server. During the period of the browse policy, users can browse backed-up data from the Sun StorEdge 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 StorEdge 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 the following occurs:

- Storage volume is relabeled.
- Entries are manually deleted.

The Sun StorEdge 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 StorEdge EBS client computer. Users can browse and mark files, and initiate data recovery. The Sun StorEdge 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.

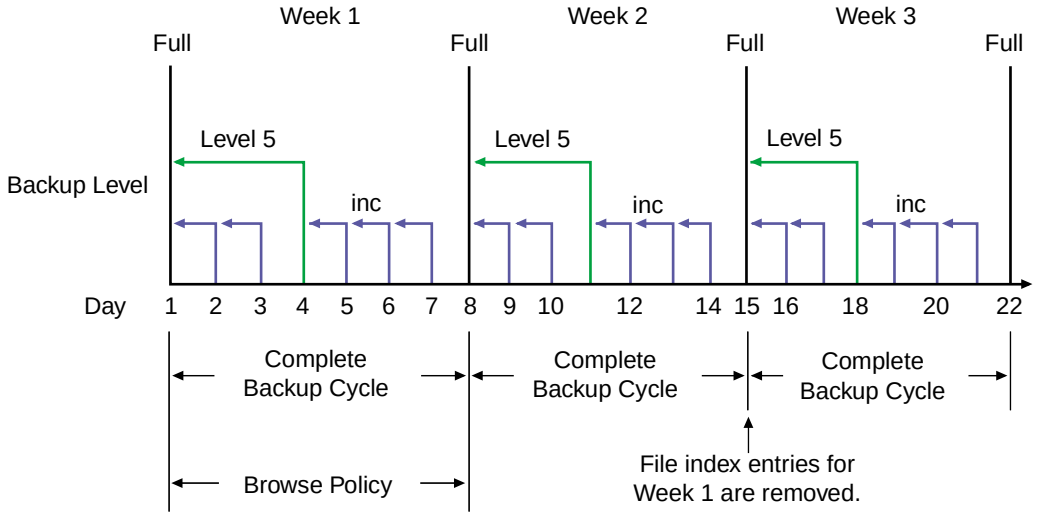
The following examples demonstrate how a browse policy affects data availability in the client file index. For more information about schedules, see "Schedules" on page 144, and for more information about backup levels, see "Backup Levels" on page 154.



Example: One-Week Browse Policy

In Figure 23 on page 171, 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 incrementals and level 5 that depend on the full backup expire.

Figure 23. One-Week Browse Policy



To further illustrate, suppose that on October 12, you recover information backed up on October 5. 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 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 browse policy (one week)
- One complete backup cycle (one additional week)

That is, until the level 5 backup on October 4 and all incremental backups dependent on the full backup pass their browse policy. In Figure 23 on page 171, entries from the Week 1 backup cycle are removed from the client file index on October 15.



Example: Two-Week Browse Policy

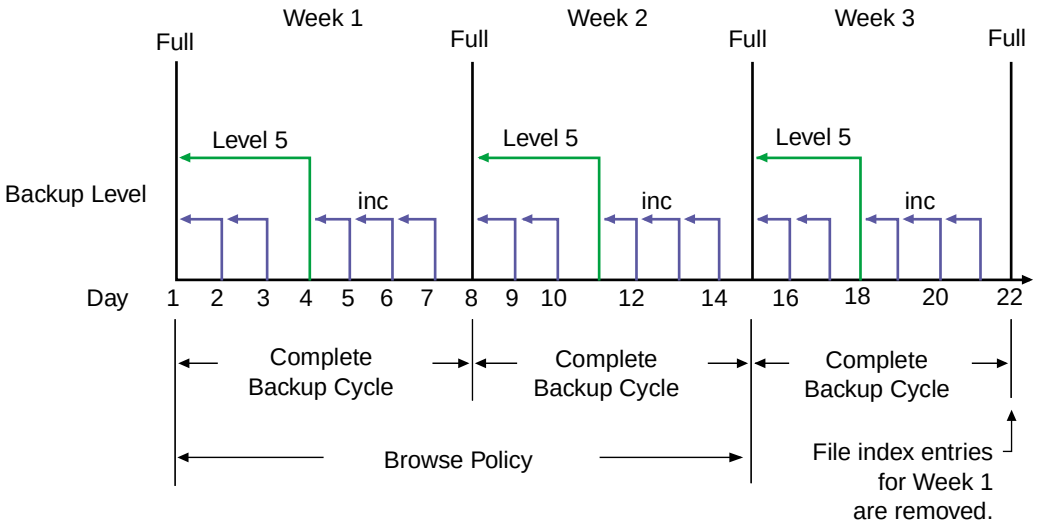
In Figure 24 on page 172, the browse policy is two weeks, which is twice as long as the backup cycle (one week). In this example, on October 18 a user can still find browsable entries in the client file index from backups created on October 4. The backup performed on October 5 is an incremental backup 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 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 24. Two-Week Browse Policy



Retention Policies

The Sun StorEdge 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 StorEdge 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.

The change in status to *recyclable* means that the volume can be overwritten if conditions are right. The volume can be relabeled and used if 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. For more information about auto media management, see "Auto Media Management" on page 328.

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.

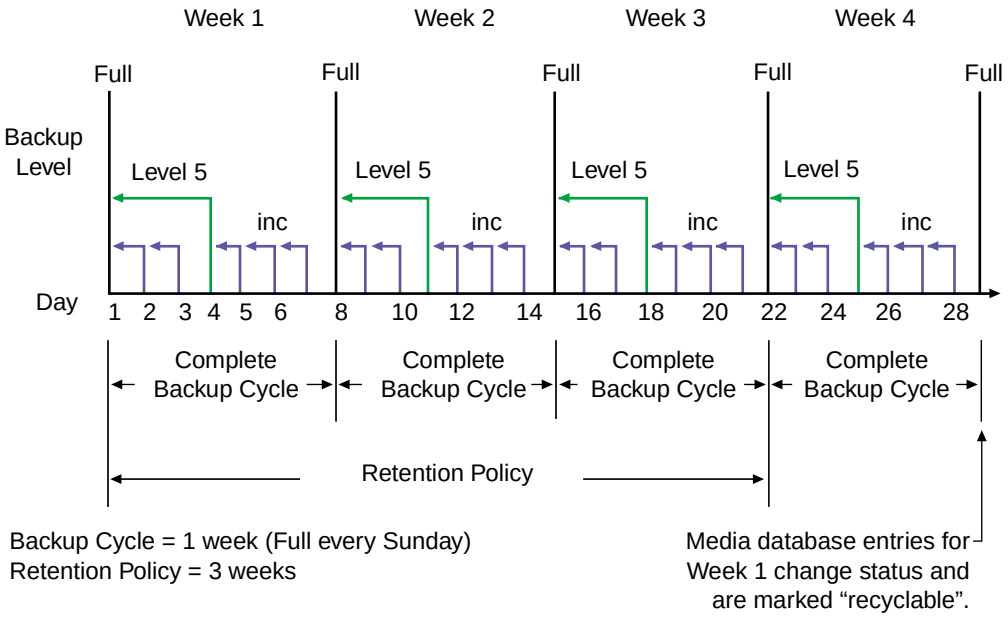
- If you re-create the entries in the client file index, a user with the proper permissions can recover data by using the Sun StorEdge EBS client computer.
- If you re-create the save set’s entries in the media database, a UNIX root user or a member of the Windows Administrators group can recover data by using save set recovery.

For more information about re-creating entries in the client file index or the media database, see "How to Restore a Save Set Entry in the Online Indexes" on page 397.

Example: Three-Week Retention Policy

Figure 25 on page 174 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 25. 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 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.



Important: Once a volume is relabeled, data on the volume cannot be recovered.

For more information on these topics, see the following sections:

- Storage volume modes, see "How to View Volume Status Information" on page 331
- Schedules, see "Schedules" on page 144
- Backup levels, see "Backup Levels" on page 154

Retention Policies for Client File Index Save Sets

The client file indexes that reside on the Sun StorEdge 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 StorEdge EBS client to which the client file index corresponds. For example, if Sun StorEdge 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 StorEdge EBS server. This ensures that if a Sun StorEdge 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. The retention policy for a save set is the longer of the retention policy for the save set client or the media pool to which the save set is backed-up.

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. For more information, see "Specifying Browse and Retention Policies for Clone Data" on page 245.

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 the command-line program. For more information, see "Browse and Retention Policies for Ad Hoc Backups" on page 182.

Managing the Data Lifecycle

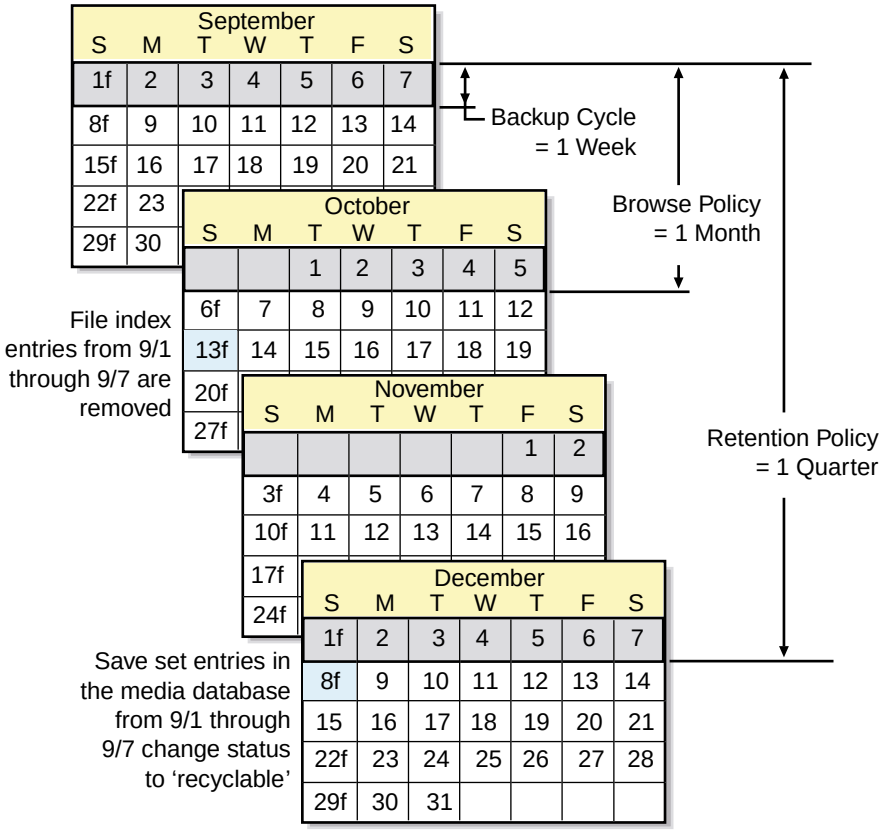
Browse and retention policies control the growth of the client file index and the media database, and how long data remains available for recovery.

Figure 26 on page 177 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 StorEdge EBS client computer. As long as the save set's file entries remain in the client file index, the status of the source save sets is browsable. After the save set status changes from browsable to recoverable, you cannot perform file recovery directly.

The status for each save set backed up during the September 1 through September 7 cycle remains recoverable until their retention policies expire, *plus* however long it takes for all the dependent save sets to pass 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.

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 StorEdge EBS client computer. For more information, see "Recovering by Save Set Selection" on page 379 and "How to Restore a Save Set Entry in the Online Indexes" on page 397.

Figure 26. 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. If auto media management is enabled, the volume may be relabeled automatically by the Sun StorEdge EBS server to satisfy a volume mount request. After the volume is relabeled, all existing data on the volume is unavailable for recovery. For more information on this feature, see "Auto Media Management" on page 328.



Important: 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 StorEdge 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 StorEdge 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: 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 StorEdge EBS software contains these preconfigured browse policies:

- 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. For more information, see "How to Create a Policy" on page 179.

How to Create a Policy

For information about creating a policy, see "Task 3: Set Up Policies for Quick Access and Long Term Storage" on page 57.

How to Edit 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.

How to Delete a Policy

Preconfigured policies cannot be deleted. For more information, see "Preconfigured Policies" on page 179.

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 PowerSnap Module documentation.

Note: VSS backups do not use snapshot policies. For more information about VSS, see "Chapter 25: Volume Shadow Copy Service" on page 745.

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 StorEdge 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.

How to Create a Snapshot Policy

To create a snapshot policy:

1. In the Administration window, click Configuration.
2. In the expanded left pane, select Shapshot Policies.
3. From the File menu, select New.
4. In the Create Snapshot Policy dialog box, enter 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 Sun StorEdge EBS Powersnap Module documentation.
5. Click OK.

How to Edit a Snapshot Policy

To edit a policy:

1. In the Administration window, click Configuration.
2. In the expanded left pane, select Shapshot 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.

How to Copy 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. Enter the name for the new Snapshot Policy resource in the Name attribute, edit any other attributes as appropriate, and click OK.

How to 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 Ad Hoc Backups

If a browse or retention policy is specified with an ad hoc backup from the command-prompt, the browse or retention policy takes effect for all of the save sets included in the ad hoc backup. Specify browse and retention policies with an ad hoc 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. For more information about **save** and **nsr_getdate**, refer to the *NetWorker Command Reference Guide* or the UNIX man pages.

If a browse or retention policy is *not* specified for an ad hoc backup, the policies are determined as follows:

- Browse policy

  The save sets included in an ad hoc backup adopt the browse policy of the Client resource. If there are multiple Client resources for the Sun StorEdge EBS host, the Client resource with the longest browse time is adopted. For example, if there are three Client resources for the Sun StorEdge EBS client *mars*, each with one of the following browse periods:

  – One week
  – One month
  – One year

  An ad hoc backup of *mars* adopts a browse policy of one year.

- Retention Policy

  The save sets included in an ad hoc 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

  – or the media pool retention policy.

  There are special considerations for retention policies and clone data. For more information, see "Specifying Browse and Retention Policies for Clone Data" on page 245.

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 the following 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**, the following 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. For more information about **nsrmm** and **nsr_getdate**, refer to the *NetWorker Command Reference Guide* or the UNIX man pages.



Example: Changing Browse and Retention Policies with nsrmm

These are examples that use **nsrmm** to change browse and retention policies:

- Change the retention time to midnight, January 1, 2009. Change the browse time to midnight, January 1, 2006.

  ```
  nsrmm -S 3315861249 -e "01/01/09 23:59:59"
  -w "01/01/06 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. For more information about **mminfo**, refer to the *NetWorker Command Reference Guide* or the UNIX man pages.

Chapter 6: Directives

This chapter covers the following topics:

- "Directives Overview" on page 185
- "Creating a Global Directive Resource" on page 186
- "Editing a Global Directive Resource" on page 187
- "Deleting a Global Directive Resource" on page 187
- "Copying a Global Directive Resource" on page 187
- "Applying Directives" on page 188
- "Sun StorEdge EBS User Program Local Directives" on page 190
- "Preconfigured Global Directive Resources" on page 192
- "Format of Directive Statements" on page 194
- "Application Specific Modules (ASMs)" on page 197

Directives Overview

Directives are special instructions that control how the Sun StorEdge EBS server processes files and directories during backup and recovery. As a Sun StorEdge EBS administrator, you can create directives to customize Sun StorEdge EBS to your specific needs, maximize the efficiency of backups, and apply special handling to individual files or directories.

Types of Local and Global Directives

There are three types of directives:

- *Global directives* — Administrators can create global directives by using the Sun StorEdge EBS Administrator program. These directives are stored as resources on the Sun StorEdge EBS server, and can be selectively applied to individual clients using the Directive attribute of the Client resource.
- *Sun StorEdge EBS User local directives (Windows only)* — On clients running Microsoft Windows, users with local Windows Administrator or Backup Operator privileges can create local directives by using the Sun StorEdge EBS 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

For instructions on the syntax to use when creating directives, see "Format of Directive Statements" on page 194.

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, enter a name for the new directive.
5. In the Comment attribute, enter a description of the directive.
6. In the Directive attribute, enter one or more directives.
7. Click OK.

The directive is now available to be applied to a Sun StorEdge EBS Client resource. For more information, see "Applying a Global Directive" on page 188.

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.
6. 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.
7. Click OK.

Applying Directives

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



Example: 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 StorEdge EBS client.

1. Create a global directive in the following 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
   ```

   For information about creating a global directive, see "Creating a Global Directive Resource" on page 186.

2. Apply the directive to the appropriate Sun StorEdge 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.
   f. Click OK.

When a scheduled backup is performed on the Sun StorEdge EBS client, all files that match the **.tmp* pattern in the specified directories will be skipped.



Example: Applying a 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 StorEdge EBS client.

Using the Sun StorEdge EBS User program, create a local directive. The directive is saved in the *networkr.cfg* file in the following format:

```
<<"C:\mydir">>
skip: *.tmp
```

When a scheduled backup is performed on the Sun StorEdge EBS client, all files that match the **.tmp* pattern in the *C:\mydir* directory will be skipped.

For more information on creating local directives with the Sun StorEdge EBS User program, see "Sun StorEdge EBS User Program Local Directives" on page 190.



Example: 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 StorEdge EBS client. When a scheduled backup is performed on the Sun StorEdge 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 enter the following directive in the file:

   ```
   skip: *.tmp
   ```

2. Place the *nsr.dir* file in the *C:\mydir* directory on the Sun StorEdge EBS client.

   Note: To create directive files on a client 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 enter the following directive in the file:

   ```
   skip: *.tmp
   ```

2. Place the *.nsr* file in the */mydir* directory on the Sun StorEdge 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, Sun StorEdge EBS User program local directives are enforced over local directive files (*nsr.dir* files).

Sun StorEdge EBS User Program Local Directives

On Windows, users can create local directives with the Sun StorEdge EBS User program. These directives are saved in a file named *networkr.cfg*.

When you perform an ad hoc backup from the Sun StorEdge EBS User program, only local directives that were created with the Sun StorEdge EBS 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 StorEdge EBS **save** without the **-i** option is run at the command-prompt.

Sun StorEdge EBS User program local directives are also enforced during scheduled backups and archive operations.

How to Set Up a Sun StorEdge EBS User Program Local Directive

To set up a Sun StorEdge EBS 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 Sun StorEdge EBS User Program.
3. From the Options menu, select Local Backup Directives.
4. Set the local directive for each data item. You can clear data items to exclude them from scheduled backups, and select items for password-protection, encryption, and compression, for both manual and scheduled saves.

Note: If password-protection or encryption is selected, the password must be specified first. For information about setting a password, see "Setting Up PW1 Password-Protection" on page 73.

5. From the File menu, select Save Directive to save changes.

Depending on user privileges and OS version, the Sun StorEdge EBS User program creates *networkr.cfg* in one of the following 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 one of the following directories, depending on the operating system:
  - Windows 2000, XP Professional, or Server 2003:

    %SystemDrive%\Documents and Settings\<User_name>\Application Data\Legato
  - Windows NT 4.0:

    %SystemRoot%\Profiles\<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 StorEdge EBS software comes with a number of preconfigured global Directive resources. All preconfigured Directive resources can be modified, but they cannot be deleted.

Table 22 on page 192 lists the preconfigured directives and their descriptions.

Table 22. Preconfigured Directives (Part 1 of 2)

Directive Resource	Description
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 standard	Contains a set of directives used to back up standard Mac OS clients. Applies the following ASMs: The **skip** ASM is applied to the following 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*

Table 22. Preconfigured Directives (Part 2 of 2)

Directive Resource	Description
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.
NetWare standard	Used to back up NetWare clients. By default, this resource has no directives.
NetWare with compression	Used to back up and compress NetWare clients. Applies the **compressasm** ASM to all files.
NT standard	Used to back up Windows clients. By default, this resource has no directives.
NT with compression	Used to back up and compress Windows clients. It applies the **compressasm** ASM to all files.
UNIX standard	Contains a set of directives used to back up standard UNIX clients. Applies the following 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.

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 # character is treated as a comment.

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

Directory Specifications

Directory specifications indicate the highest-level directory to which the following ASMs apply. The format for a directory specification is:

<<directory>>

To apply ASMs to a directory, use the following format:

- On Windows:

  ```
  <<"C:\BIN">>
  ```
  asm

  Note: 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 as follows:

- On UNIX:

  ```
  <</etc>>
  ```
  asm

Multiple directory specifications can be included in the same Directive resource:

- On Windows:

  ```
  <<"C:\BIN">>
  ```
  asm

  ```
  <<"C:\TEMP">>
  ```
  asm

- On UNIX:

  `<<"/etc">>`
  asm

  `<<"/tmp">>`
  asm

When multiple directory specifications are used, directives that follow a directory specification are applied to that directory until the next directory specification. You cannot use wildcards in directory specifications.

ASM Specifications

ASMs specify the action to take on one or more files. The syntax for an ASM specifications is:

`[+]`*asm*`:` *argument*

where:

- The optional plus sign (+) indicates that the directive applies to both the current directory and all subdirectories.
- The *asm* is the name of the ASM to be executed
- The *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 (" "). For a description of available ASMs and examples, see "Application Specific Modules (ASMs)" on page 197.

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 `* .?*` (ensure that a space is placed 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 23 on page 196 lists the three save environment keywords.

Table 23. Save Environment Keywords

Keyword	Instruction	Example

forget	Instructs the Sun StorEdge 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, enter:
		`<</SRC>>`
		`+skip: *.o`
		`<</SRC/SYS>>`
		`forget`
		This uses the **skip** ASM to instruct the Sun StorEdge 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 StorEdge EBS server to ignore all directives applied to the subdirectories below the current directory.	To override any local directives set in users' home directories, enter the following:
		`<</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, enter:
		`<</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 StorEdge EBS environment to perform various tasks during backup and recovery. For example, the **compressasm** program instructs the Sun StorEdge 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 24 on page 197 describes the Sun StorEdge EBS ASMs.

Table 24. Sun StorEdge EBS ASMs (Part 1 of 4)

ASM Name	Description	Example
aes	Encrypts backup data when included in a global directive that is associated with a Client resource. **Note:** This ASM works with backup data on UNIX and Windows clients. See also the *pw2* 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 may 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: .`

Table 24. Sun StorEdge EBS ASMs (Part 2 of 4)

ASM Name	Description	Example
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 StorEdge 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`
nsrindexasm	Used to recover from Sun StorEdge EBS file index backups performed by Sun StorEdge 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 StorEdge EBS directives.	Not applicable

Table 24. Sun StorEdge EBS ASMs (Part 3 of 4)

ASM Name	Description	Example
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 StorEdge 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: *.?*`
pw1	Password-protects the specified files by using an encrypted password. For more information, see "Password ASMs" on page 201. **Note:** This ASM works with Microsoft Windows and NetWare versions of Sun StorEdge EBS software only.	`pw1: *.*` `password "-o{koo{o"`
pw2	The same as the pw1 ASM, except that the files are also encrypted. For more information, see "Password ASMs" on page 201. **Note:** This ASM is supported on Windows and NetWare clients only. See also the *aes* ASM.	`pw2: *.*` `password "-o{koo{{"`

Table 24. Sun StorEdge EBS ASMs (Part 4 of 4)

ASM Name	Description	Example
rawasm	Specify this directive to back up 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. For more information, see "Precautions When Using rawasm to Back Up UNIX Raw Partitions" on page 201.	`rawasm: /dev/oracle1`
skip	Omits files and directories from the backup, and does not place the directory or filename in the online index. In the example given, all files and directories with the name *temp* will be omitted from the backup.	`+skip: temp`
swapasm	Does not back up actual file data, but re-creates a zero-filled file of the correct size on recovery. This ASM is used on systems where the swapping device is a swap file that must be recovered with the correct size, but the contents of the swap file are not important and do not need to be backed up or restored.	`swapasm: compression.doc`

Password ASMs

Password ASMs are available for Microsoft Windows and NetWare clients only. Password ASMs are used to password-protect (**pw1**) or password-protect and encrypt (**pw2**) data. Directives that use these ASMs require a local file named *networkr.cfg*, which is created when a password is specified in the Sun StorEdge EBS User program. For more information about entering a password in the Sun StorEdge EBS User program, see "Password Protecting Windows or NetWare Backup Data" on page 73.

If you create a global directive file or a local directive file that uses **pw1** or **pw2**, include the encrypted password by copying it from the *networkr.cfg* file. The format of **pw1** and **pw2** requires that the encrypted form of the password be included on the line directly below the ASM specification. For example:

```
pw1: *.*
   password "-o{koo{o"
```

If you create a **pw1** or **pw2** directive by using the Sun StorEdge EBS User program, there is no need to copy the encrypted password from the *networkr.cfg* file.

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. For more information, see "Backing Up Raw Partitions on Windows" on page 99.

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 the following two 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
```

Then 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. For more information, see "Save Environment Keywords" on page 196.

Chapter 7: Sorting Backup Data

This chapter covers the following topics:

- "Media Pools" on page 203
- "Label Templates" on page 217

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 StorEdge EBS server writes data. The server uses media pools to sort and store data. A volumes is identified with a unique label based on configurable label templates.

Media Pools

Media pools act as filters that tell the server which backup volumes should receive specific data. The Sun StorEdge EBS server uses media pools in conjunction with label templates to keep track of what data is on which specific volume. For more information on label templates, see "How the Sun StorEdge EBS Software Uses Label Templates" on page 217.

How the Sun StorEdge EBS Server Uses 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 StorEdge 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 StorEdge EBS server writes data to the volume.
- If there is no correctly labeled volume mounted on a storage device, the Sun StorEdge 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 StorEdge 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 StorEdge EBS server directs save streams there. If the Default media pools are not properly configured for backup, an error message like the following 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 StorEdge EBS Media Pool Types

Sun StorEdge EBS software contains preconfigured media pool types to keep different types of data separate. Sun StorEdge EBS software does not mix the following types of data within a media pool:

- Backup
- Backup Clone
- Archive
- Archive Clone
- Migration
- Migration Clone

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. Customized media pools can also be created. See "Create a Media Pool" on page 213.

How Sun StorEdge EBS Software Uses Media Pool Criteria to Sort Data

When you configure the Sun StorEdge 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 StorEdge EBS client
- Save sets (file or file systems)
- Backup levels (full, levels [1 – 9], incremental, ad hoc)

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.

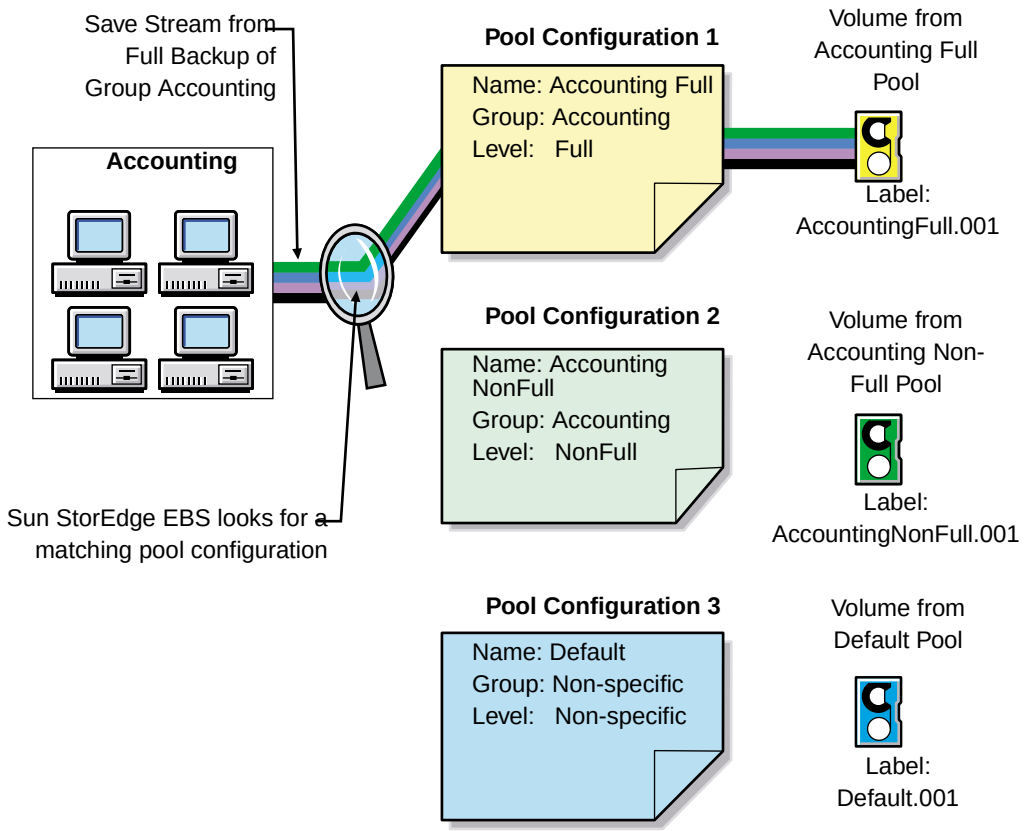
If you enter configuration criteria in both the Group attribute and Save Set attribute, only data that meets both the group criteria *and* the save set criteria is written to volumes from the specified media pool.

For more information about groups and backup levels, see "Chapter 4: Backup Groups and Schedules" on page 133.

Example: Using Media Pool Configurations to Sort Data

Figure 27 on page 206 illustrates how the Sun StorEdge 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 StorEdge EBS server looks for a media pool configuration that matches the group named Accounting and the level Full. When the Sun StorEdge 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 27. 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.

In the following example, the client file indexes are in *index*. To send the Sun StorEdge 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 25 on page 207.

Table 25. Media Pool Resource 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 StorEdge EBS server's bootstrap and index save sets are written to a separate volume labeled for the Index media pool.

How to Direct Data from a Consolidated Backup 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. For instructions, see "How to Create a Group" on page 138.
2. Configure a Media Pool resource for consolidated backups. For instructions, see "How to Create a Media Pool" on page 213.
3. In the Create Media Pool dialog box, select the name of the Group resource created in step 1 for the Groups attribute.
4. Edit the Client resource for each client to receive consolidated backups, and assign the client 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 StorEdge EBS server uses the following 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 26 on page 209 details the hierarchy that the Sun StorEdge 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 group.

Table 26. Sun StorEdge EBS Hierarchy for Resolving Media Pool Conflicts

Precedence	Group	Client	Save Set	Level
Highest	x	x	x	x
	x	x	x	
	x	x		x
	x	x		
	x		x	x
	x		x	
	x			x
	x			
		x	x	x
		x	x	
		x		x
		x		
			x	x
			x	
Lowest				x

Not Meeting the Criteria for Any Customized Media Pool

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 StorEdge 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 StorEdge 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 StorEdge 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

The following sections provide information specific to configuring the following types of media pools:

- Incremental backups
- Ad hoc 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 StorEdge EBS_install_dir*> \ *logs* directory. 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 StorEdge EBS software's hierarchy of precedence affects the way the data is stored:

- If the level value is incremental, incremental data will be routed to the associated media pool, but the corresponding changes to the client's file index will not be.
- If the level of the save is incremental, the Sun StorEdge EBS server saves all client file indexes at 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, 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. Thus, to speed the recovery, define the level value to accept both level 9 and incremental data.

By using the preconfigured NonFull media pool settings, you ensure that the client file indexes belong to the same media pool as their incremental backups. Keeping the indexes in the same media pool as their incremental backups reduces the number of volumes required for a recovery.

Ad Hoc Backups

You can create a customized media pool to receive data from an ad hoc backup by specifying manual in the Level attribute. The Sun StorEdge EBS server, however, sorts data from an ad hoc backup differently than it sorts data from a regularly scheduled backup. Since an ad hoc backup is not performed as part of a scheduled backup group, the data is not associated with any group name. Thus, when you perform an ad hoc 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 an ad hoc 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 ad hoc backups, the Sun StorEdge EBS server uses the Default media pool and looks for a mounted volume from the Default media pool on which to write data.

Note: Ad hoc backups only back up file system data, not the client file index entries that are made during scheduled backups. This makes recovering data much more difficult and time consuming. Do not substitute ad hoc backups of data for regular scheduled backups.

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. For more information about cloning, see "Chapter 9: Cloning Backups" on page 243.

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. For more information about archive feature, see "Chapter 8: Archiving" on page 225.

Using Storage Devices and Media Pool Configuration to Sort Data

Sort data by configuring media pools in conjunction with storage devices. Do one of the following:

- Use specific media to receive data
- 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. For more information about how a single save set can be written to more than one type of storage media, see "Client and Save Set Combination" on page 584.

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.

How to Create a Media Pool

Resource dialog box contents vary depending on the server. Use the steps presented here as a general guideline. However, additional attributes may be present (particularly if Diagnostic Mode is enabled). Similarly, attributes mentioned here may be absent on your server. For guidance in completing these additional attributes, click Field Help.

Before creating a media pool, do one of the following:

- Create a label template for the media pool to be created. For instructions, see "How to Create a Label Template" on page 223.
- 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 StorEdge EBS server notifies you that it will create a label template for the media pool. To have the Sun StorEdge EBS server create the label template, click OK. For a complete discussion of label templates, see " How the Sun StorEdge EBS Software Uses Label Templates" on page 217.

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, enter 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, enter 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 the following attributes: Client, Save sets, or Level. For guidelines about entering save set paths, see "Expression Matching of Save Sets to Media Pools" on page 214.
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 215.
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, the following 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, so that:

- 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 StorEdge 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
```

For information about regular expression matching, refer to the **nsr_regexp** and **nsr_pool** entries in the *NetWorker Command Reference Guide* or the UNIX man pages.

Auto Media Verification

If the Auto Media Verify attribute is enabled, the Sun StorEdge 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. If the data read matches the data written, verification succeeds.

Media is verified whenever:

- 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.

How to Edit 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 to edit.
4. From the File menu, select Properties.
5. In the Properties dialog box, make the necessary changes and click OK.

How to Copy 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 to copy.
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, enter a name for the new media pool.
6. Edit any other attributes as appropriate, and click OK.

For more information about the Media Pool attributes, click Field Help in the Properties dialog box.

How to Delete 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 to edit.
4. From the File menu, select Delete.
5. When prompted, click Yes to confirm deletion.

How to Create an Archive Media Pool

To archive data, the Sun StorEdge 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. For instructions, see "How to Create a Media Pool" on page 213.
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 StorEdge EBS server creates a unique label for each volume by applying a label template. Label templates and media pools are used to sort, store, and track data on media volumes.

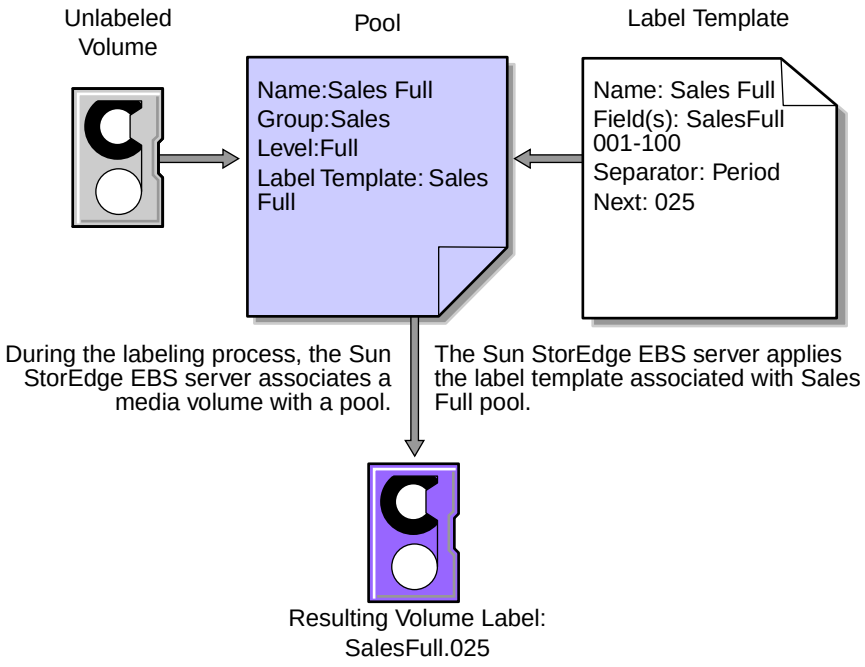
How the Sun StorEdge EBS Software Uses Label Templates

The Sun StorEdge 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 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 StorEdge EBS server uses the preconfigured Default media pool and corresponding Default label template.

Figure 28 on page 217 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 28. Labeling a Volume by Using a Label Template



Preconfigured Label Template

The following is a list of preconfigured label templates contained with the Sun StorEdge EBS server. These label templates 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 name of the Sun StorEdge EBS server and the final field contains a number to allow for expansion of the media pool, the number range starts at 001 and ends with 999.

For example:

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

How to Complete Label Template Attributes

There are certain guidelines to keep in mind when completing the attributes for a Label Template resource. Table 27 on page 219 describes how to complete the key attributes for this resource. For more information on creating a label template, see "How to Create a Label Template" on page 223.

Table 27. Key Label Template Attributes (Part 1 of 2)

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 the following characters in label template names: / \ * ? [ ] ( ) $! ^ ; ' " ' ~ < > & | { } In addition, do not use the following four characters: colon (:) • dash (-) • period (.) • underscore (_)

Table 27. Key Label Template Attributes (Part 2 of 2)

Attribute	Guidelines
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 StorEdge EBS server applies each component in a left-to-right order, starting with the first one entered. Table 28 on page 221 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 StorEdge EBS server places on a volume (according to the template). • To force a label to start the label scheme at a particular point, enter 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 StorEdge EBS server generate the first label, leave this attribute blank. When the Sun StorEdge 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 28 on page 221 lists examples of number sequences for volume labels.

Table 28. 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.



Important: When the Sun StorEdge 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.

Label Template Naming Strategies

The Sun StorEdge 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 29 on page 222 illustrates several ways you can use components to organize labels.

Table 29. 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 StorEdge 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 StorEdge 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 StorEdge EBS. For example, if you physically label a volume mars.1, its Sun StorEdge EBS name should also be *mars.1*.

How to Create a Label Template

When creating a label template, consider the labeling guidelines described in Table 27 on page 219.

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.

How to Edit 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 enter 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.

How to Copy a Label Template

To copy a Label Template:

1. In the Administration window, click Media.
2. In the 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, enter the name for the new label template.
6. Edit any other attributes as appropriate, and click OK.

How to Delete 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 template to delete.
4. From the File menu, select Delete.
5. When prompted, click Yes to confirm the deletion.

Chapter 8: Archiving

This chapter covers the following topics:

- "Overview of Archiving" on page 225
- "Permissions for Archiving" on page 228
- "About Archive Pools" on page 230
- "Archiving Data Procedures" on page 231
- "Retrieving Archived Data" on page 238
- "Archive Request Management" on page 242

The archive feature must be purchased and licensed separately from other Sun StorEdge EBS software components. For more information on licensing procedures, see "Managing Product Licenses" on page 500.

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.

Benefits of using the Sun StorEdge 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 StorEdge EBS software does not support archiving of SYSTEM or VSS SYSTEM save sets.

Archive Requirements

Before licensing and using the Sun StorEdge EBS archive feature, ensure that you have the following:

- A device, either stand-alone or in an autochanger or silo, connected to a Sun StorEdge 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.

How the Sun StorEdge EBS Server Archives Data

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

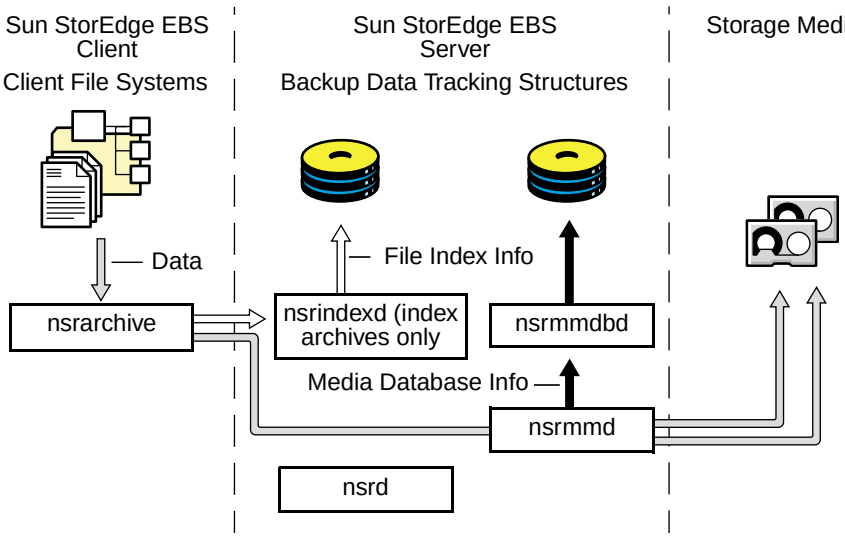
Note: Index entries are not generated when the Store Index Entries attribute in the Pool resource is set to No.

The Sun StorEdge 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 StorEdge EBS Administrator program.

Figure 29 on page 227 illustrates how the Sun StorEdge EBS software archives data.

Figure 29. Overview of Archive Operation



Indexed and Nonindexed Archiving

The Sun StorEdge 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 StorEdge EBS server, all clients of that server are enabled for the Sun StorEdge 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 StorEdge EBS server, ensure that the Archive Services attribute is enabled in the Client resource for the server. For instructions, see "Enable Archive Services for a Sun StorEdge EBS Client" on page 231.

Important: 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 StorEdge EBS client software are installed on the same computer and both back up to the same Sun StorEdge 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 StorEdge 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. For more information about the Administrators group and the Backup Local Data privilege, see "Sun StorEdge EBS User Groups" on page 533

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 StorEdge EBS server.
3. From the File menu, select Properties.
4. Select the Public Archives attribute.
5. 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 StorEdge EBS software provides three preconfigured pools to receive archived data:

- Indexed Archive pool
- PC Archive pool
- 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. For information on creating custom archive pools, see "Creating Custom Archive Pools" on page 231.

If you do not specify a pool to store archived data, the Sun StorEdge 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.



Important: 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. Therefore, you cannot retrieve individual files from the archive save set. 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 StorEdge 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).

For details and procedures about creating pools, see "Media Pools" on page 203.

Archiving Data Procedures

You can request manual archives from the client, or you can schedule archives from the server.

How to Enable Archive Services for a Sun StorEdge EBS Client

To enable archive services for a Sun StorEdge 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 resource with the same hostname.
4. In the Archive Users attribute, enter all appropriate users. For details about setting this attribute, see "Enabling or Restricting Archive Access" on page 229.
5. Make the remaining configuration choices as appropriate.

8

The computer is now an enabled archive client. However, an archive will not occur until it is requested. For instructions, see "Scheduling Data Archives" on page 235.

Note: If the Sun StorEdge EBS client is set up for encryption with the **aes** ASM, archive data will also be encrypted. For information about setting up encryption for a Sun StorEdge EBS client, see "Implementing Encryption with the AES ASM" on page 75.

Manually Archiving Data

You can manually archive data at any time. Manually archiving data is similar to performing an ad hoc backup.

How to Perform a Manual Archive from a Sun StorEdge EBS Client on Windows



Important: Manual archives 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 Sun StorEdge EBS User program are enforced. Scheduled archives, enforce all directives. For more information about scheduled archives, see "Scheduling Data Archives" on page 235.

To perform a manual archive for a Windows client:

1. In the Sun StorEdge EBS User program, click Archive to open the Archive Options dialog box.
2. Enter a comment in the Annotation attribute. This annotation is used to uniquely identify each archive save set during retrieval.

   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 the following:
   - To write a copy of each archive save set to a volume from an archive clone pool, select Clone.

     If you enable cloning, enter or select an archive clone pool for the Archive Clone Pool attribute.

- To instruct the Sun StorEdge EBS server to check the integrity of the data on the storage volume, select Verify.
- To instruct the Sun StorEdge EBS server to remove the archived files from the disk, select Groom.

5. Click OK.

   The Archive browse window appears.

6. Select Mark from the File menu to select each file or directory for archiving. When you select an item for archiving, a check mark appears next to that item.

   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.

   The Archive browse window appears.

   The Sun StorEdge EBS server displays the Archive Status window, which monitors the progress of the archive. When the Sun StorEdge EBS server is finished archiving, a message similar to the following appears in the Archive Status window:

   ```
   Archive completion time: 2-21-06 5:18p
   ```

9. If Groom was selected in step 4, the Remove Archived File dialog box prompts for confirmation before Sun StorEdge EBS software deletes archived files from the local disk.

How to Perform a Manual Archive from a Sun StorEdge EBS Client on UNIX



Important: Manual archives performed with the **nwarchive** command do not directives. However, scheduled archives do enforce global and local file (*nsr.dir*) directives. For more information about scheduled archives, see "Scheduling Data Archives" on page 235.

8

To perform a manual archive from a UNIX client:

1. Start the **nwarchive** program as follows:
   - To archive to the default Sun StorEdge EBS server, enter the following command:

     nwarchive &
   - To archive to a Sun StorEdge EBS server other than the default server, enter the following command:

     nwarchive -s *server* **&**
2. Select Mark from the Mark menu to select each file or directory for archiving.
3. For the File menu, select Start Archive.
4. Enter a comment in the Annotation attribute. This annotation is used to uniquely identify each archive save set during retrieval.

   Consider adopting a consistent naming convention so that one can easily identify archives based on the annotation name.
5. Select the appropriate archive pool from the Archive Pool attribute. Only pools with the Pool Type attribute set to Archive are listed.
6. Select the appropriate response for the following data archiving attributes:
   - Compress: To compress the data as it is archived, select Yes. Compression and decompression use more server resources and may slow data archiving and retrieval.
   - Clone: To write a copy of each archive save set to a volume from an archive clone pool, select Yes.

     If you enable cloning, enter or select an archive clone pool in the Clone Pool attribute.
   - Verify: To check the integrity of the data on the storage volume, select Yes.
   - Grooming: To remove the archived files from the disk, select Yes.
7. Click Archive.

   The Archive Status window is displayed. When the Sun StorEdge EBS server is finished archiving, a message similar to the following appears in the Archive Status window:

   ```
   Archive completion time: 2-21-03 5:18p
   ```
8. If Groom was selected in step 6, you are prompted to confirm before archived files are deleted from the local disk.

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.

How to Schedule an Archive

Note: Before you can schedule an archive request, enable the Archive Services attribute in the Client resource. For information, see "How to Enable Archive Services for a Sun StorEdge EBS Client" on page 231.

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, enter a name for the archive request.
5. In the Comment attribute, enter a description of the archive request.
6. In the Client attribute, enter the archive client's hostname.
7. In the Save Set attribute, enter the pathnames of the files and/or directories to be archived.

   To archive all of the client's save sets, enter All for the Save Set attribute.



Important: If you choose All for the Save Set attribute, 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. Enter a comment in the Annotation attribute. This annotation is used to uniquely identify each archive save set during retrieval.

   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. For more information about directives, see "Chapter 6: Directives" on page 185.
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, enter a notification for the Sun StorEdge EBS server to use after completing the archive. For more information about notifications, see "Notifications" on page 541.
14. Select the appropriate response for the following options:
    - To instruct the Sun StorEdge EBS server to remove the archive files from the disk, select Remove from the Grooming list.
    - To instruct the Sun StorEdge 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. See "View Details of a Scheduled Archive" on page 242.

How to Copy an Archive Request Resource

To copy an Archive Request resource:

1. In the Administration window, click Configuration.
2. In the left pane, click 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. Enter the name for the new archive request in the Name attribute, edit any other attributes as appropriate, and click OK.

For more information about the Archive Request attributes, click Field Help in the Properties dialog box.

How to Change the Archive Time

To change the archive time:

1. In the Administration window, click Configuration.
2. In the left pane, click 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, enter a new time in the following 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. For more information, see "Archive Request Management" on page 242.

How to Edit 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.

8

How to Delete 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 StorEdge EBS administrator or the owner of the file can retrieve the data.

For more information, see "Permissions for Archiving" on page 228.

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.

How to Retrieve Nonindexed Archives from a Sun StorEdge EBS 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 Sun StorEdge EBS 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, enter all or part of the annotation assigned to the save set when it was archived.

   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 the following if necessary:
   - To view a list of volumes required to retrieve the data from this archived save set, click Required Volumes.
   - To enter a new path for the location of the recovered data and to indicate what the Sun StorEdge 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. For information, see "Recovering by Save Set Selection" on page 379.

How to Retrieve Nonindexed Archives from a Sun StorEdge EBS 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 **nwretrieve** program. For example:

   ```
   nwretrieve &
   ```
3. To limit the archived save sets found during the query, complete one or more of the following attributes:
   - Client: The client with the archived data to retrieve.
   - Save Set Name: The name of the archived save set to retrieve.

- Search Annotation For: All or part of the annotation assigned when the save set was archived. To aid the search, use wildcard symbols.

4. Click Query to search for the save sets.

   Matching archived save sets appear in the Archives window.

   To view the details of an archived save set, such as its size, save set ID, and archive date, select the archived save set and click Details.

5. To begin the retrieval, select the archived save set to retrieve and select Start Retrieve from the File menu.

   The Retrieve Status window appears.

   To relocate the files, enter the new location in the Relocate To attribute. To overwrite existing files, enable the Overwrite Existing Files attribute.

6. Click OK.

Note: You can also recover archived save sets using save set recovery. For information, see "Recovering by Save Set Selection" on page 379.

How to Recover Indexed Archive Data from a Sun StorEdge EBS 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 Sun StorEdge EBS 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. Mark the files to be recovered by selecting the files and clicking Mark.
6. Click Start.

How to Recover Indexed Archive Data from a Sun StorEdge EBS 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.

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 **nwrecover** program:

   ```
   nwrecover &
   ```

2. From the Change menu, select Index Namespace>Archive.
3. Mark the files to be recovered by selecting the files and clicking the Mark button.
4. Click the Start on the toolbar.

Retrieving Archive Data from the Command-prompt

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

Note: To identify an archive for retrieval, use the **-A** option to specify an annotation for the archive. However, the **-A** option requires you to specify a string that uniquely identifies one annotation. For example, suppose archive A had the annotation *Accounting_Fed* and archive B had the annotation *Accounting_Local*. In this example, the following command would not to find a match:

```
nsrretrieve -A Accounting
```

However, the following command would find archive B:

```
nsrretrieve -A ting_L
```

8

Archive Request Management

This section describes how to perform tasks involving archive requests.

How to Start a Scheduled Archive at Any Time

You can start a scheduled archive immediately rather than waiting 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.

How to Stop 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.

How to Disable 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.

How to View Details of a Scheduled Archive

To open the Archive Request Details window:

1. In the Administration window, click Monitoring.
2. Right-click the archive request and select Show Details.

The Archive Request Details window provides information about the completion of an archive request:

- The Completion Time displays the time the archive finished. This is the difference between the completion and start times of the archive.
- The success of the archive request is either completed, failed, or partial.

Chapter 9: Cloning Backups

This chapter covers the following topics:

- "Overview of Cloning" on page 243
- "Volume Cloning" on page 246
- "Save Set Cloning" on page 247
- "Manual Clone Operations" on page 256
- "Recovering Cloned Data" on page 257
- "Cloning Archived Data" on page 258
- "Directing Clones to a Special Storage Node" on page 258
- "Considerations for File Type and Advanced File Type Devices" on page 260

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 StorEdge EBS administrator. Cloning can be performed on volumes and on save sets.

Cloning information can be found in the Monitoring option of the Administration window. Clone-related messages are also logged to the Sun StorEdge EBS *message* file and the *savegrp.log* file, which are located in the <*Sun StorEdge EBS_install_dir*> \ *logs* directory.

Cloning Requirements

The following are required for cloning:

- 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 StorEdge 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. The *source* volume is the original volume, and the *destination* volume is the volume to which data is cloned.
- You must be a member of the Sun StorEdge EBS Administrators group. For more information, see "Sun StorEdge EBS User Groups" on page 533.

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 StorEdge EBS server attempts to complete one group's scheduled backup before initiating a scheduled backup for another group. However, the Sun StorEdge 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. Additionally, **nsrclone** with the **-t** and **-e** options enables you to specify a date range so that only save sets that were backed up for a group within a specific period of time, such as the last 24 hours, will be cloned. For more information about the **nsrclone** command, refer to the *NetWorker Command Reference Guide* or the UNIX man pages.

Cloning Compared to Duplication

Cloning leaves traceable entries in both the client file index and the media database. Being able to track cloned data is an important distinguishing feature between cloning and a duplication 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 StorEdge EBS software does not recognize a duplicated volume, because it does not have a valid Sun StorEdge EBS label. If auto media management is enabled and a duplicated volume is left in the autochanger, the unlabeled volume is eligible for labeling and reuse during a scheduled backup. For more information about auto media management, see "Chapter 12: Media Management" on page 325.

It is possible to make an exact copy of a Sun StorEdge 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 Browse and Retention Policies for Clone Data

Clone data is given the same browse policy as the original save set. However, the retention policy for clone data can be specified independently. 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, either:

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

If a retention policy is *not* specified in either the clone pool or from the command-prompt, the retention policy is derived from the original save set.

How to 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.

   For information about editing or creating a pool, see "How to Create a Media Pool" on page 213.

2. From the Retention policy list, select a retention policy.
3. Click OK.

How to Specify a Retention Policy from the Command-prompt

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

To specify a retention policy from the command-prompt, either:

- Specify a retention policy when creating a clone save set by using the **nsrclone -y** command.
- Specify a retention policy for an existing clone save set by using the **nsrmm -e** command.



Important: For Sun StorEdge EBS 7.3, the **nsrclone -y** option has changed. Ensure that custom scripts or reporting tools that use this option are updated to reflect this change.

For more information about the **nsrclone** and **nsrmm** commands, refer to the *NetWorker Command Reference Guide* or the UNIX man pages.

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.

The Sun StorEdge EBS server allows only one clone of a particular save set per volume. Therefore, if three clones of the same save set are specified, each clone is written to a separate volume.

How to Create 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.

How to View 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.

There are two different methods for cloning save set data:

- *Manual* cloning of save sets. For more information, see "How to Manually Clone a Save Set" on page 248.
- *Automatic* cloning of save sets. For more information, see "Automatically Clone a Save Set" on page 250.

When you specify that a particular save set is to be cloned, the Sun StorEdge EBS software first determines whether that save set already has a clone. If multiple clones of the save set exist, the on volumes in an autochanger are generally selected as the source data, rather than a volume that is manually mounted. Use the command-prompt to specify the save set clone as the source.

How to View 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. For more information, see "How to Manually Clone a Save Set" on page 248.

How to Manually Clone a Save Set

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. Enter values in any of the following 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. Enter 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 enter 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 the following formats are acceptable:

- Written out completely (for example, November 1, 2006)
- Numerically as mm/dd/yy (for example, 11/01/06)
- Date and time selection from the list.

Specifying a large date range may result in a large number of save sets being selected. This can reduce response time or even require that you close and re-open the browser connection to the Sun StorEdge EBS Console.

7. Use the Status attribute to limit the search to save sets that have a particular status. Table 30 on page 249 lists the values that can be selected.

Table 30. Save Set Status Settings

Status	Description
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, and the volume is now available for automatic relabeling and overwriting.
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.
In-progress	Select if the save set is currently in the process of being backed up.
Aborted	Select if the save set was either aborted manually by the administrator during a backup, or because the computer crashed.
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.

   To sort the list of save sets, click a column heading in the table. For example, to sort the list in alphabetical order by client name, click the top of the Client column.

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.

    To view the status of a manual clone operation or to cancel a clone operation that is in progress, see "View Manual Clone History" on page 256.

How to Automatically Clone a Save Set

Save set data can be automatically cloned when a scheduled group backup is run. For general information on configuring groups, including cloning, see "Backup Groups" on page 134.

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 is specified in the Client resource's Clone Storage Nodes attribute. You can add or remove the names of storage nodes and the Sun StorEdge EBS server at any time.

Examples of Manual and Automatic Clone Operations

This section contains three different examples of save set cloning:

- "Manual Cloning of a Save Set" on page 251
- "Automatic Cloning of One Client" on page 252
- "Automatic Cloning of Several Clients" on page 253

Note: 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: Manual Cloning of a Save Set

In this example, a user has requested that the Sun StorEdge 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 state that also uses Sun StorEdge EBS software. The Sun StorEdge EBS 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 StorEdge EBS administrator must have the following information:

- Sun StorEdge EBS client name
- Name of the save set
- Date the data was backed up

To manually clone the save sets, the administrator must:

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, enter the client name.
5. In the Save Set attribute, enter the save set name.
6. For the Start Date and End Date, enter 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: Automatic Cloning of One Client

A system administrator wants to clone data automatically, on a daily basis, for the Sun StorEdge EBS client computer *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 needs to perform the following tasks:

- Create a 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*.

   For more information on creating groups, see "Task 2: Set Up a Group for Backup Clients" on page 56.

   Select the appropriate choices for the group. The required choices are:

   - Autostart = Enabled
   - Clones = Yes
   - Clone pool = Default Clone (or a custom clone pool)

2. Create the client *jupiter*.

   For more information on creating clients, see "Task 7: Create a Backup Client Resource" on page 61.

Select the appropriate backup choices for *jupiter*. The required choices are:

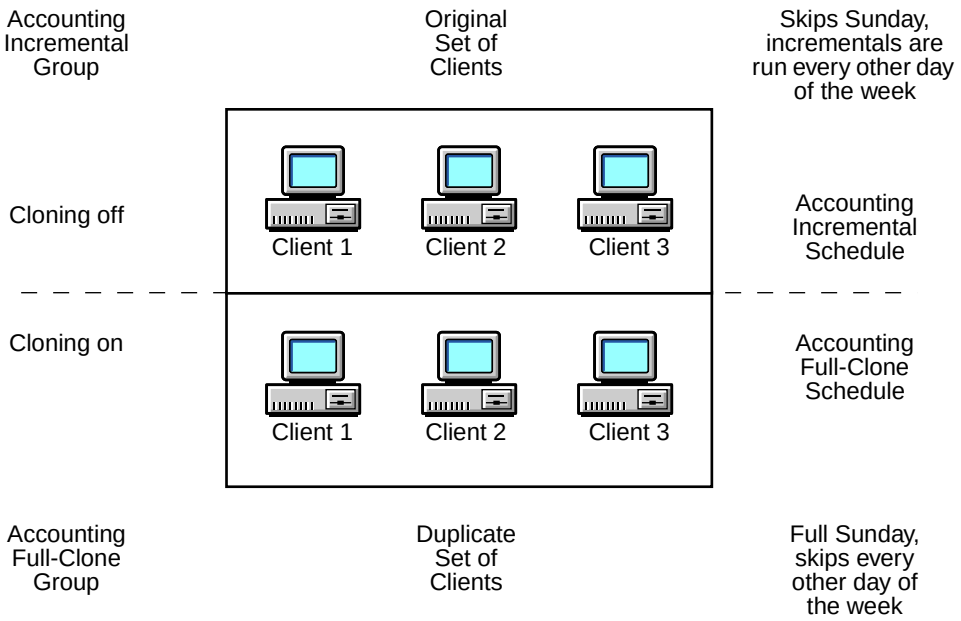
- Schedule = Full Every Friday
- Group = Documents

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: Automatic Cloning of Several Clients

Figure 30 on page 253 illustrates the level of control you can achieve with cloning. You can create an unlimited number of automatic cloning configurations using Groups, Schedules, and the Client resources.

Figure 30. Automatic Cloning Scheme



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

To do this, two different groups and schedules are created and applied to two identical sets of clients. Both groups belong to the same pool and use the browse and retention policies specified for the Client resource.

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

The system administrator must do the following:

- 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 (shown in Table 31 on page 254).

Table 31. Accounting Incremental Schedule

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

  - Create another schedule named Accounting Full-Clone (shown in Table 32 on page 254).

Table 32. 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*.

  Note: Automatic cloning is applied to one of the groups. As a result, two groups are necessary in this situation. The *Accounting Incremental* group is used for performing incremental backups, and the *Accounting Full-Clone* group is used for performing full backups and clones.

- Create *two* clients with the same name for each client computer. Two clients of the same name are needed because only one schedule can be applied to a Sun StorEdge EBS client at a time. In this example, one set of clients uses the Accounting Incremental schedule and the other (identical) set of clients uses the Accounting Full-Clone schedule.

- Select the appropriate group and schedule for each set of clients.

To automatically clone data, the system administrator must:

1. Create two schedules:
   a. Create the *Accounting Incremental* schedule to perform incremental backups every day except Sunday, when it skips the backup (see Table 31 on page 254).
   b. Create the *Accounting Full-Clone* schedule to perform full backups every Sunday and skip backups the rest of the week (see Table 32 on page 254).

      For details on creating schedules, see "Task 1: Set Up a Schedule for Backups" on page 56.
2. Create two groups:
   a. Create the *Accounting Incremental* group with the following attributes:
      – Autostart: Enabled
      – Clones: No
   b. Create the *Accounting Full-Clone* group with the following attributes:
      – Autostart: Enabled
      – Clones: Yes
      – Clone: Default Clone

      For details on creating groups, see "Task 2: Set Up a Group for Backup Clients" on page 56.
3. Create two sets of identical clients, and select the appropriate schedule and group for each client.

   For details on creating clients, see "Task 7: Create a Backup Client Resource" on page 61.
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:
   - Schedule: Accounting Incremental
   - Group: Accounting Incremental

   After making these choices *for each client*, click OK.
5. Select the second (duplicate) set of clients, one client at a time, this time selecting Accounting Full-Clone for the Schedule and Group attributes.

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

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, see "Volume Cloning" on page 246 or "Save Set Cloning" on page 247.

How to View 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 appears that shows manual clone history information.

How to Stop 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

A clone volume is used for recovery when the Sun StorEdge EBS server attempts to recover a particular save set and one of the following has occurred:

- The original save set volume has been deleted.
- The status of the original save set has been marked *suspect*.

The volumes required for a save set recovery appear in the Required Volumes window. For more information, see "How to View the Volumes Required for Data Recovery" on page 378.

To recover save set files from a clone volume, designate the original save set as suspect. If you cannot determine which save set on the volume to set to suspect, change the status of all save sets on the volume to suspect. For information about changing the status of a save set, see "Change Save Set Status within the Volume Save Sets Window" on page 363.

If a clone volume containing the cloned save set exists, it appears in the Required Volumes window. Otherwise, the software attempts to recover the data from the suspect volume.

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. For more information, see "How to Restore a Save Set Entry in the Online Indexes" on page 397.

Recovering a Cloned Save Set from the Command-prompt

Recovery of a cloned save set can be performed at the command-prompt, as in the following example:

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

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

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. For more information, see "How to Manually Clone a Save Set" on page 248.
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

The Clone Storage Nodes attribute enables you to direct clone operations to a specific storage node. When a volume is being cloned, the value of the Clone Storage Nodes attribute is checked on the storage node:

- If the Clone Storage Nodes attribute has no value, the value listed for the Sun StorEdge EBS server's Clone Storage Nodes attribute is used.
- If the Sun StorEdge EBS server's Clone Storage Nodes attribute also has no value, the value in the Sun StorEdge EBS server's Storage Node attribute is used.

Note: 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 StorEdge EBS server. This ensures that browse and retention policies are handled consistently regardless of where the clone data was directed.

Independently Directing Clones for Each Storage Node

To direct clones for each storage node independently:

1. Open for editing, the Client resource for the storage node whose clone data will be directed to another storage node.

   For more information, see "How to Edit a Client" on page 568.

2. Click the Globals (2 of 2) tab.
3. In the Clone Storage Nodes attribute, add the hostname of the storage node that is to receive the directed clones.

   The first entry in the list that has a functional, enabled device is selected to receive the cloned data from the storage node.

   Note: The Clone Storage Nodes attribute applies only to Sun StorEdge EBS servers and storage nodes. This attribute should be blank for Sun StorEdge EBS clients that are not also Sun StorEdge EBS servers or storage nodes.

Directing Clones from All Storage Nodes to a Single Destination

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

1. For each storage node Client resource, leave the Clone Storage Nodes attribute blank.
2. Open for editing, the Client resource for the Sun StorEdge EBS server.

   For more information, see "How to Edit a Client" on page 568.

3. Click the Globals (2 of 2) tab.
4. In the Clone Storage Nodes attribute, add the hostname of the storage node that is to receive the clones for all storage nodes.

Considerations for File Type and Advanced File Type Devices

This section covers considerations for 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.



Example: Manual Cloning with Advanced File Type Device

In a situation where there are three save sets:

- Save set A - 10 kilobytes
- Save set B - 10 megabytes
- Save set C - 10 gigabytes

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.

Cloning to an Advanced File Type Device

If a saveset is cloned from a file device to an advanced file type device and the following conditions apply:

- Both the original backup (file type device) and the clone volume (advanced file type device) are mounted.
- One attempts to recover the saveset by save set ID.

In this case, the data will be restored from the clone volume (advanced file type device) and not the original backup volume.

Chapter 10: Staging Backups

This chapter covers the following topics:

- "Save Set Staging" on page 261
- "Working with Staging Policies" on page 262
- "Sun StorEdge EBS DiskBackup" on page 266

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. Once the save set is staged to a tape device, it cannot be staged again. However, one could still clone the tape volume.

Staging can be driven by 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.

Note: If a file type device is on a storage node that is running Sun StorEdge EBS software 6.1 or earlier and the file type device is full, the file type volume will not be automatically marked as appendable after the staging operation.

Working with Staging Policies

This section describes how to create, edit, and delete staging policies.



Important: To prevent an advanced file type device from becoming full during backup, the staging policy must be set up so that older save sets are automatically moved to another medium to make disk space available in the advanced type device.

How to Create a Staging Policy

Before creating a staging policy, configure all appropriate devices. Otherwise, no devices will be listed in the Devices attribute.

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, enter a name for the staging policy.
5. In the Comment attribute, enter a description of the staging policy.

6. To enable staging to begin immediately or to be invoked automatically at a later time, set the Enabled attribute to Yes.

   You can enable or disable staging at any time.

7. In the Devices attribute, select the file type and adv_file type devices as the source device for staging.

   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. The other volume types can be staged to any pool.

9. In the High-Water Mark (%) attribute, enter 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 lower 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, enter 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, enter 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, enter 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, enter 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 Operation tab.
    b. In the Start Now attribute, select *one* of following 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 Apply or 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.

How to Edit a Staging Policy

Note: You *cannot* edit the name of an existing 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 resource to edit.
4. From the File menu, select Properties.
5. Make any necessary changes and click OK.

How to Copy 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. Enter the name for the new Staging resource in the Name attribute, edit any other attributes as appropriate, and click OK.

How to Delete a Staging Policy

To delete a staging policy:

1. In the Administration window, click the Configuration button.
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 StorEdge EBS console. When staging from the Sun StorEdge 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 StorEdge 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 more information about **nsrstage** or **mminfo**, refer to the *NetWorker Command Reference Guide* or the UNIX man pages.

Sun StorEdge EBS DiskBackup

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

Sun StorEdge EBS software offers two variations of this solution:

- File Type Device (Windows storage node, local path only)
- Advanced File Type Device (AFTD)

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, do one of the following:

- 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. For more information about save set staging, see "Save Set Staging" on page 261.

File Type Device

Configure the Sun StorEdge 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 conditions apply to 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 is an arbitrary value that represents a guess as to what the volume capacity is likely to be and has no impact on how much data is actually written to a volume. For file type devices, however, the Volume Default Capacity is a hard limit to the amount of data that can be written to the device. In order to avoid accidentally filling a file system, it is important to set this value 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.

  The Volume Default Capacity appears in the Configuration tab of the Properties window of a device when View>Diagnostic Mode is set. The value for this attribute must be less than 4 TB.

  Note that the adv_file type device ignores the Volume Default Capacity value.
- Save sets from file type devices can be cloned to an advanced file type device, to take advantage of advanced file type device features.



Important: 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 the following:
  - The upper limits supported by the operating system
  - The file size specified by the disk device vendor

See "Differences between File Type Devices and AFTDs" on page 281 for a table of 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 can be dynamically extended (using a volume manager) if the disk runs out of space while being accessed as an AFTD during backup. For a list of supported volume managers, refer to the *Software Compatibility Guide.*

An AFTD differs from a file type device in that a volume for an AFTD is not marked full unless the maximum number of save sets is reached. For additional differences between file type devices and AFTDs, see "Differences between File Type Devices and AFTDs" on page 281.

When an *AFTD* runs out of disk space, the current backup is suspended. The following message is displayed:

`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 StorEdge EBS server on UNIX and Linux systems, and logs a message in the *<Sun StorEdge EBS_install_path> \ logs \ media.log* file on Windows systems.

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.



Important: Managing large disk devices can involve complex workflow management once various events are triggered. This workflow can be handled by using notifications. The Sun StorEdge EBS software provides the two preconfigured notifications described previously. Users can also customize their own notifications, thus changing and expanding the way the Sun StorEdge EBS software behaves when an AFTD Full notification occurs, or even invoking custom scripts and other programs that can be used to expand the capacity of existing AFTDs, among other things. See "Notifications" on page 541 and "Configuring StorEdge EBS SNMP Notifications" on page 692 for more information about using notifications.



Example: Create a Custom Notification to Extend Disk Space

While the Sun StorEdge 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, select Configuration>Notifications, then right-click Notifications and select New.
2. For Name, enter a unique name for this custom notification, such as *First adv_full Notice*.
3. For Event, unselect / uncheck all choices except *adv_file*.

4. For Priority, unselect / uncheck all choices except *Waiting*.
5. For Action, specify the full path of the custom script to be invoked, for example: */mybin/my_first_custom_script*.
6. Select OK.



Example: Create a Custom Notification for Insufficient Disk Space

The Sun StorEdge 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, select Configuration>Notifications, then right-click Notifications and select New.
2. For Name, enter a unique name for this second custom notification, such as *Second adv_full Notice*.
3. For Event, clear or unselect all choices except *adv_file*.
4. For Priority, clear or unselect all choices except *Critical*, *Emergency*, and *Alert*.
5. For Action, specify the full path of the custom script to be invoked, for example: */mybin/my_second_custom_script*.
6. Select 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 that Dynamic Drive Sharing is not supported with the AFTD.

As part of the validation process, include the following 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



Important: Some versions of NFS or CIFS drop data blocks 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 StorEdge EBS software.

The upper limits of save set size depend on either of the following:

- The upper limits supported by the operating system
- 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.

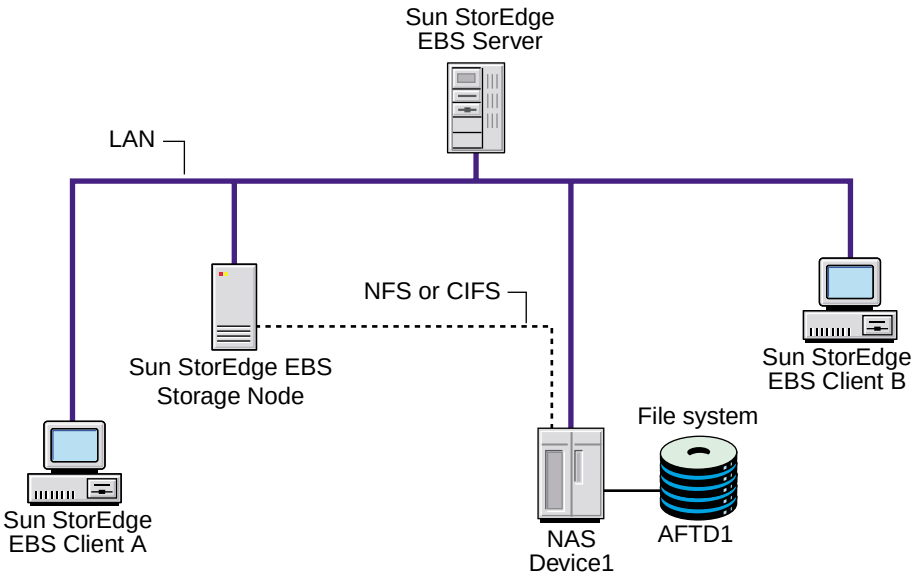
AFTD Concurrent Operations

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 at the same time. For more information, see "Differences in the Cloning Process" on page 284.

Figure 31 on page 272 shows a system configuration where both Sun StorEdge EBS client A and Sun StorEdge EBS client B can use NFS to recover files concurrently from AFTD1.

Figure 31. 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 StorEdge EBS software creates a secondary device with read-only accessibility. The Sun StorEdge EBS software then creates and automounts a volume with a *.RO* suffix 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.

See Table 36 on page 281 for information that 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, the following information would apply. See Table 33 on page 273.

Table 33. AFTD Names

Device - UNIX and Linux format	Type	Volume
/mercury/dev1	adv_file	af_1
/mercury/dev1/_AF_readonly	adv_file	af_1.RO
Device - Windows format		
d:\file_devices\dev1	adv_file	af_1
d:\file_devices\dev_AF_readonly	adv_file	af_1.RO
Device — 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

In this table, when the AFTD is labeled, the *_AF_readonly* device is created during the labeling process. The Sun StorEdge EBS software then labels and mounts a volume with a *.RO* suffix (for Read Only) on this device.

Table 33 on page 273 also illustrates that when backing up data to the device:

/mercury/dev1 (for UNIX and Linux) or

d:\file_devices\dev1 (for Windows) or

(rd=storage_node:)\\<CIFS_host>\sharepoint\dev1 (for CIFS on Windows)

it is possible to read concurrently from the device:

/mercury/dev1/_AF_readonly (for UNIX and Linux) or

d:\file_devices\dev_AF_readonly (for Windows) or

(rd=storage_node:)\\<CIFS_host>\sharepoint\dev1_AF_readonly (for CIFS on Windows)

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 34 on page 274 shows which Sun StorEdge EBS operations can be performed concurrently within a single storage node.

Table 34. Concurrent Operations within a Single Storage Node

Backups	Recoveries (Command-prompt Interface 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 the following 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.

Table 35 on page 274 shows which Sun StorEdge EBS operations can be performed concurrently when sharing an AFTD in NAS for data storage.

Table 35. Concurrent Operation with Storage Nodes Sharing an AFTD in NAS

Backups to Storage Node1 with Read/Write Volume Mounted	Recoveries from Storage Node2 with Read-Only (.RO) Volume Mounted	Manual Cloning from Storage Node2 with Read-Only (.RO) Volume Mounted	Manual or Auto Staging from Storage Node2 with Read-Only Volume Mounted
X	X		
X		X (1 clone)	
X			X (1 stage)

Limitations with Concurrent AFTD Recovery Operation

AFTD concurrent recovery currently has the following limitations:

- Not available to the Windows recover interface (**winworkr**). Use the **recover** command.
- Not available to nonfile recoveries, such as NDMP and Sun StorEdge EBS database modules.
- Recoveries in which Sun StorEdge EBS release 7.2*x* or earlier Windows **recover** clients get data from release 7.3 AFTD storage nodes are sequential.

Sharing an AFTD between Storage Nodes in a NAS or SAN Environment

In a Network Attached Storage (NAS) Environment, Sun StorEdge 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 StorEdge 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. (See "Configuring a Shared AFTD in a NAS Environment" on page 276 and "Configuring a Shared AFTD in a SAN Environment" on page 278 for more information.) Once the configuration has been done, however, no change is required in how to start the following operations:

- Backup
- Recovery
- Cloning
- Staging

Before Configuring Shared AFTDs

When preparing to share 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 StorEdge EBS software cannot change that.
- To share an AFTD between storage nodes, ensure the following:
  - 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 the following:
  - The **nsrexecd** (Sun StorEdge EBS Remote Exec) service has been started as Administrator on the storage node.
  - The CIFS-mapped AFTD was created with UNC path names that have the appropriate remote user and password specified. An example of UNC path syntax is:

    rd=sn_a:\\nas1\path\shared_aftd

  See "CIFS-Mounted AFTD on Microsoft Windows Servers and Storage Nodes" on page 292 for more information.

Configuring a Shared AFTD in a NAS Environment

When sharing an AFTD between two storage nodes in a NAS environment, the following 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 configuration steps:

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*

   See permission and password information in "Before Configuring Shared AFTDs" on page 276.
3. In the Sun StorEdge EBS software, mount the AFTD (but do not label it) on storage node B, to create the read-only volume.
4. In the Sun StorEdge 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 StorEdge EBS software, unmount the read-write volume from storage node B (if it is mounted), then in the Sun StorEdge EBS software, mount the read-write volume on storage node A.
   - Configure the Storage Nodes attribute of the Client resource (*sn_a*).
   - Ensure that the Sun StorEdge 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, in the Sun StorEdge EBS software, unmount the read-only volume from storage node A (if it is mounted), then in the Sun StorEdge EBS software, mount the read-only volume on storage node B.
   - Configure the Recover Storage Nodes attribute of the Client resource (*sn_b*).
   - Configure the Clone Storage Nodes attribute in the Client resource with the target of the clone (or stage) operation.
   - Configure the staging policy's Devices attribute to select *rd=sn_b:/shared_aftd* for autostaging from *sn_b*.

Also see "Concurrent Operation with Storage Nodes Sharing an AFTD in NAS" on page 274.

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 configuration steps:

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 StorEdge 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 Sun StorEdge EBS, mount the AFTD (but do not label it) on storage node B to create one read-only volume and one read-write volume.
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 the following limitations:

- Read-only volumes might occasionally 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, followed by correctly remounting.
- 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 StorEdge 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 StorEdge EBS software.

Changing Block Size for AFTDs

The NSR_DEV_BLOCK_SIZE_ADV_FILE environment variable can be used for customization or tuning. Block sizes can be set in kilobytes for AFTDs. The default is 128 KB. In order for any change in block size to take effect, the Sun StorEdge EBS daemon or service must be restarted after resetting the NSR_DEV_BLOCK_SIZE_ADV_FILE environment variable.

To change the AFTD block size to 64K:

1. On UNIX or Linux:

   `setenv NSR_DEV_BLOCK_SIZE_ADV_FILE 64`

2. On Windows:
   - At the command-prompt:

     `set NSR_DEV_BLOCK_SIZE_ADV_FILE=64`
   - From the Start menu, select Settings > Control Panel > System > Environment Variables

3. Restart the Sun StorEdge EBS services.

How an AFTD Displays

When the AFTD is implemented, certain display 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 from Read-Write to Read-Only

When an AFTD is labeled, an *_AF_readonly* device is automatically created and added to the Sun StorEdge EBS device list. When data is backed up, additional entries are added to the Sun StorEdge EBS media index for the data backed up to that device.

When changes are made to the user-created, AFTD (read-write device) attributes, the same changes must immediately be made to the attributes of the *_AF_readonly* device. If changes are manually made to volume (or save set) properties within the Sun StorEdge EBS indexes, those changes must also be propagated to the other read-only volumes and save sets.

To propagate the changes:

1. Access the read-write device, volume, or save set and change the appropriate attribute.
2. Access the read-only device, volume, or save set to match the attribute change made to the read-write device, volume, or save set.

Note: Always make attribute changes to the read-write device, volume, or save set first, and then propagate these changes to the read-only device, volume, or save set.

Table 36 on page 281 lists the attributes that must be propagated manually from the read-write device, volume, or save set to the read-only device, volume, or save set.

Table 36. Attributes Requiring Manual Propagation

Read-Write Device or Save set	Attribute
Device	Enabled
Save set	Mode

Note: The Volume Mount attribute need not be propagated from read-write to read-only volume.

Differences between File Type Devices and AFTDs

Table 37 on page 281 lists the functional differences between the file type device and the AFTD.

Table 37. Advanced File Type and File Type Device Differences (Part 1 of 4)

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. For further information, see the Important Note on "AFTD Environments" on page 270. Windows storage node: local path only

Table 37. Advanced File Type and File Type Device Differences (Part 2 of 4)

Function or Operation	AFTD	File Type Device
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: For more information, see "AFTD Concurrent Operations" on page 272.	Yes	No
Save set ID / clone ID: For more information, see "AFTD Concurrent Operations" on page 272.	• 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.

Table 37. Advanced File Type and File Type Device Differences (Part 3 of 4)

Function or Operation	AFTD	File Type Device
Reclaiming or recovering space	• Aborted save sets immediately removed. • The **nsrim** program removes expired save sets, once every 24 hours, after a savegroup is completed (if volume recycle is set to Auto).	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). **Note:** Once the Sun StorEdge EBS server and storage node have been upgraded to the 7.0 release, the first time **nsrim** is run at the end of a savegroup, all aborted and expired save sets are removed and can no longer be recovered.
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. For more information, see the Important Note on page 267.

Table 37. Advanced File Type and File Type Device Differences (Part 4 of 4)

Function or Operation	AFTD	File Type Device
When file system or volume is full	• Message displayed stating file system requires more space. • **nrsim** 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	Yes
Save set consolidation	Not supported	Supported
Data format in device	Save stream (**uasm**) format (uses less space)	EMC Legato 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]*

For more information about the **nsrclone** command, refer to the *NetWorker Command Reference Guide* or the UNIX man page.



Example: Manual Cloning with AFTDs

This example uses three save sets:

- Save set A - 10 kilobytes
- Save set B - 10 megabytes
- Save set C - 10 gigabytes

When save set A has completed its backup, the manual cloning process can begin immediately—while the other two, larger save sets are still being backed up.

As each save set is backed up, the cloning process for that respective save set can be launched immediately. This is essentially a time-saving feature that enables the optimization of Sun StorEdge EBS operations.

Only one save set can be cloned at a time.

Configuration Support for Sun StorEdge EBS DiskBackup

The following Sun StorEdge EBS DiskBackup configurations are supported:

- Local disk of UNIX, Linux, or Microsoft Windows versions of Sun StorEdge EBS storage node. This is supported on both file type devices and AFTDs.
- A network-attached disk device that is NFS-mounted to a UNIX version of the Sun StorEdge 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 StorEdge EBS storage node running on Microsoft Windows. This is supported only on AFTDs.

Supported Disk Vendors

The Sun StorEdge EBS DiskBackup solution supports the following disk vendors:

- Those NFS implementations that must behave as UNIX or Linux NFS.
- Those CIFS implementations that must behave as Windows CIFS.



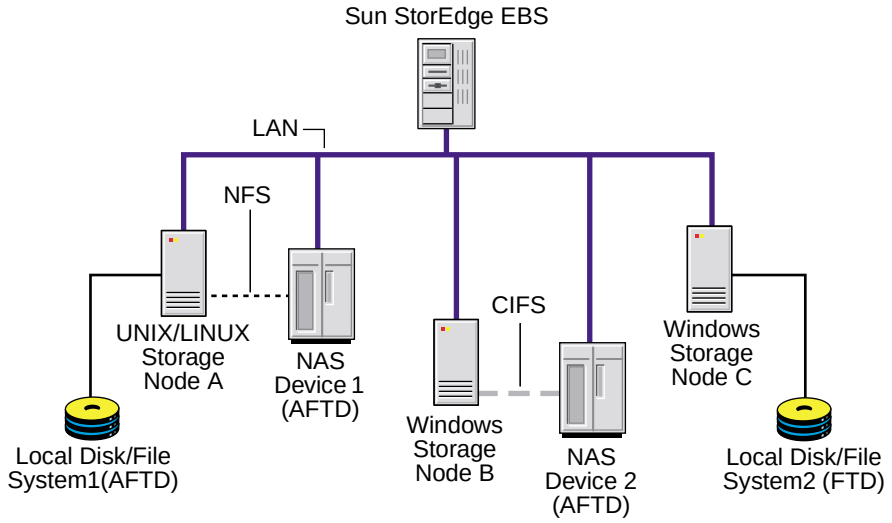
Example: Sun StorEdge EBS DiskBackup within a System Configuration

Figure 32 on page 287 illustrates how to use Sun StorEdge EBS DiskBackup within a system configuration. In this example:

- The AFTDs are created on Disk Device 1, Disk Device 2 and the Local Disk 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 StorEdge EBS DiskBackup can be tailored to the specific configuration of the system.

Figure 32. Sample Sun StorEdge EBS DiskBackup Configuration



Sun StorEdge EBS DiskBackup Configuration

The following sections explain how to configure a Sun StorEdge EBS DiskBackup Device:

- On UNIX or Linux servers and remote storage nodes with local or NFS-mounted disks. See "File Type Device or AFTD on UNIX and Linux Servers and Remote Storage Nodes" on page 288.
- On Microsoft Windows servers and remote storage nodes with local disks. See "File Type Device or AFTD on Local Disks of Microsoft Windows Servers and Storage Nodes" on page 290.
- On Microsoft Windows servers and remote storage nodes with CIFS-mounted disks (AFTD only). See "CIFS-Mounted AFTD on Microsoft Windows Servers and Storage Nodes" on page 292.

For specific instructions and restrictions on backing up to a file type device disk, see "Sun StorEdge EBS DiskBackup" on page 266.

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 StorEdge EBS software to back up data to a file type device or AFTD by creating a new Device resource for the Sun StorEdge EBS storage node just as for any other storage device.



Important: Do not modify directories or files created by file type devices or AFTDs. Changes to the device path or its contents prevent the Sun StorEdge EBS software from locating the device.

Restrictions Regarding NFS File Systems

When a file type device (but not an AFTD) is created on an NFS file system, the following restrictions apply:

- 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, the data for that file device must be staged or cloned to another device.

  Subsequently, the directory contents must be deleted and the file device must be re-created and relabeled. If these actions are not completed, the change in volume default capacity does not take effect.

How to 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 StorEdge 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 StorEdge EBS DiskBackup devices because the data could be overwritten.

2. In the server's Administration interface, click the Devices button.
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. Also in the Identity area, configure the following:
   a. Add a comment (optional) in the Comment field.
   b. In the Media Type field, select a type from the menu:
      a. *file*, for a file type device.
      b. *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 the following attributes, as necessary:
   – Target Sessions
   – Max Sessions
   – Local Backup to a dedicated storage node (leave at the default No setting)
9. Click OK, when the configuration is complete.

Label and mount theSun StorEdge EBS DiskBackup Device in the same manner as other offline media used for Sun StorEdge EBS backups.

Note: 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.



Important: Save set staging is especially useful when using either device type of the Sun StorEdge EBS DiskBackup feature. For more information, see "Chapter 10: Staging Backups" on page 261

File Type Device or AFTD on Local Disks of Microsoft Windows Servers and Storage Nodes

Configure the Sun StorEdge EBS software to back up data to a file type device or AFTD by creating a new Device resource for the Sun StorEdge EBS server just as for any other storage device.



Important: Do not use the *temp* directory for Sun StorEdge EBS DiskBackup devices. The data could be overwritten.

Note: Because AFTDs are never marked as full, there is no need to configure another device in the same pool.

How to 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 StorEdge EBS server recognizes as the device name (and the destination for the data).

2. In the server's Administration interface, click the Devices button.

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. Also in the Identity area, configure the following:

   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 the following attributes, as necessary:

   – Target Sessions

   – Max Sessions

   – Local Backup to a dedicated storage node (leave at the default No setting)

9. Click OK, when the configuration is complete.

Label and mount the Sun StorEdge EBS DiskBackup Device in the same manner as other offline media used for Sun StorEdge 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.



Important: Save set staging is effective when used with either device type of the Sun StorEdge EBS DiskBackup feature. For more information, see "Save Set Staging" on page 261.

CIFS-Mounted AFTD on Microsoft Windows Servers and Storage Nodes

This section explains how to start the Sun StorEdge EBS service and configure the AFTD on a Common Internet File System (CIFS) connection on a Microsoft Windows system. Note that the Sun StorEdge EBS service must be started as administrator before the AFTD can be configured.

How to Start the Sun StorEdge EBS Service on Microsoft Windows

To start the Sun StorEdge EBS services as administrator:

1. From the Windows Start menu, select Settings > Control Panel > Administrative Tools > Services.
2. Right-click Services.
3. Stop the Sun StorEdge EBS Remote Exec service for Windows Storage Node or Sun StorEdge EBS Backup and Recover Service for the Sun StorEdge EBS server.
4. Right-click Sun StorEdge EBS Remote Exec Service or Sun StorEdge EBS Backup and Recover Service.
5. Select Properties.
6. Select the Log On tab.
7. Select This Account radio button.
8. Replace Local System with Administrator.
9. Enter the administrator password in the appropriate field.
10. Select OK.
11. Right-click the Sun StorEdge EBS Remote Exec service for Windows Storage Node or Sun StorEdge EBS Backup and Recover Service.
12. Start the Sun StorEdge EBS Remote Exec service for Windows Storage Node or Sun StorEdge EBS Backup and Recover Service.

How to Configure a CIFS-Mapped AFTD on Microsoft 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 StorEdge EBS server recognizes as the device name (and the destination for the data).

Note: Do *not* use the *temp* directory for the AFTD, because the data could be overwritten. Also, because AFTDs are never marked as full, there is no need to configure another device in the same pool.

2. In the server's Administration interface, click the Devices button.
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, using the Universal Naming Convention (UNC):

   \\<*CIFS_host*>\<*share-point-name*>*path*

   rd=<*sn_host*>:\\<*CIFS_host*>\<*share-point-name*>*path*

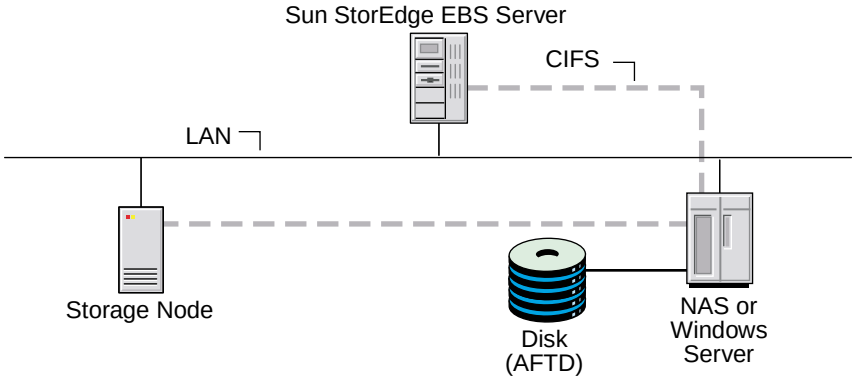
   Note: To configure a Microsoft Windows server, do not include "rd=<*sn_host*>:" in the UNC.
5. For Media Type, select *adv_file*.
6. In the Status area, make sure that the Auto Media Management feature for AFTD or file type device is *not* enabled.
7. In the Cleaning area, leave the options for cleaning at their default (disabled) settings, so that automatic cleaning is not invoked.
8. Select the Configuration tab to set the following attributes, as necessary:
   - Target Sessions
   - Max Sessions
   - Local Backup to a Dedicated Storage Node (leave at the default No setting)
9. In the Remote Host area:
   - Leave NDMP unselected.

– The following are required if the CIFS file system has been set up to require password access:
  a. Enter the CIFS file system username in the Remote User field. For example, in the configuration illustrated in Figure 33 on page 294, enter the administrator of the CIFS file system.
  b. Enter the CIFS password in the Password field. For example, in the configuration illustrated in Figure 33 on page 294, enter the password for the CIFS file system.

10. Click OK, when the configuration is complete.

Important: If a new password for the CIFS AFTD has been provided, the Sun StorEdge 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 33. CIFS-Mounted Advanced File Type Device



Note: AFTDs are labeled and mounted in the same manner as other offline media used for Sun StorEdge EBS backups. Because the AFTD automatically affixes the *.RO* suffix to a volume name, the volume name must be limited to 60 characters during labeling.



Important: Save set staging is especially useful when using either device type of the Sun StorEdge EBS DiskBackup feature. For more information, see "Chapter 10: Staging Backups" on page 261.

Deleting File Type Devices and AFTDs

For information about deleting file type and advanced file type devices, see "Deleting Libraries and Devices" on page 129.

Sun StorEdge EBS DiskBackup Licensing

Licensing for the Sun StorEdge EBS DiskBackup option is available on a per server basis.

Enablers for the Sun StorEdge EBS DiskBackup options are obtained in the same manner as for other Sun StorEdge EBS features.

Before the evaluation period is terminated, obtain the permanent license. If the enabler for the Sun StorEdge EBS server or Sun StorEdge EBS DiskBackup feature expires (for example, the evaluation period is terminated), any Sun StorEdge EBS DiskBackup devices that exist are no longer available for saves.

Once the enabler expires, it is still possible to recover the backed-up data, but it is no longer possible to create new Sun StorEdge EBS DiskBackup devices or mount any Sun StorEdge EBS DiskBackup device volumes, except as read-only. At that point, it is also no longer possible to send backups to already-mounted writable volumes on file type devices and AFTDs.

Sun StorEdge EBS DiskBackup Licensing Error Messages

One of two error messages might appear if the system is not licensed appropriately for the Sun StorEdge EBS DiskBackup Option:

- If no License Manager is configured for the Sun StorEdge EBS server:

  ```
  Server is not licensed for DiskBackup support. Install
  DiskBackup enabler.
  ```

- If License Manager is configured for the Sun StorEdge EBS server, but the DiskBackup enabler has not yet been requested:

```
DiskBackup enabler requested from EMC Legato License
Manager. Ensure enabler is available there and retry in a
few minutes.
```

These messages apply to both file and adv_file type devices, since both are components of the Sun StorEdge EBS DiskBackup Option.

Chapter 11: Device Operations

This chapter covers the following topics:

- "Compatible Block Size for UNIX and Microsoft Windows" on page 298
- "Library Device Maintenance" on page 300
- "Device Parameter Settings" on page 303
- "Common Device Interface" on page 307
- "Nonrewinding Tape Device Usage (UNIX / Linux Only)" on page 308
- "Block Size Mode (UNIX / Linux Only)" on page 309
- "Silos" on page 309
- "Miscellaneous Device Operations" on page 317
- "Tips for Using Libraries" on page 320

For operations involving storage media, see "Chapter 12: Media Management" on page 325.

One 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 StorEdge EBS servers, configured Sun StorEdge 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 StorEdge 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. Select View > Diagnostic Mode to see the full range of settings and options available in the Properties window. In the Properties window, click Field Help for a description of the various fields and attributes on a given Properties tab.

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

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.

How to 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 StorEdge EBS Administration interface, click the Devices button.
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 details table appears.
4. Right-click the drive in the details table, and select Properties.

   The Properties window appears.
5. Select the Volume tab.
6. In the Loaded Volume area, one of the displayed volume attributes is the Volume Block Size.
7. Click OK or Cancel to close the Properties window.

Solutions to Block Size Compatibility Problems

To solve problems with block size compatibility for a given device:

1. In the server's Sun StorEdge EBS Administration interface, click the Devices button.
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 to be changed.

The drive's details table appears.

4. Right-click the drive in the details 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.

To set the block size for an entire device type:

- On a UNIX system, change the block size by setting the following 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 StorEdge EBS server (also found in the Media Type attribute on the General tab of the device's Properties). Note that the syntax for *MEDIA_TYPE* is all uppercase, with underscores (_) replacing blank spaces. Therefore, a device displayed in the Sun StorEdge EBS software as "4mm 20gb" would be listed as follows:

  ```
  4MM_20GB
  ```

  – *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 StorEdge EBS server in order for changed environment variables to take effect.

Library Device Maintenance

Clean a storage device periodically to keep it working correctly. The Sun StorEdge 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 that is settable 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 the following is true:

- The last time the device was cleaned was a full cleaning interval ago.
- The Cleaning Required attribute for the device is set to Yes, either manually (by the user) or automatically (by the Sun StorEdge 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.

The cleaning attributes of a library can be set in the library's Properties window, on the General Tab. To set the value for Cleaning Delay, set View > Diagnostic Mode and then go to the Timers tab.



Important: Do not enable automated cleaning for silos in the Sun StorEdge 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.

Tape Alert

The TapeAlert feature provides diagnostic information for devices for which hardware cleaning is enabled. To use this feature, enable automatic cleaning.

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. For more information about CDI, see "Common Device Interface" on page 307.

Devices that are capable of TapeAlert perform constant self-diagnostics and communicate the diagnostic information by means of **nsrmmd** 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 38 on page 301 describes the nature of the tape alert levels.

Table 38. TapeAlert 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 indicate status information:

- A data or cleaning tape nearing its end of life
- An unsupported tape format

Note: When automatic cleaning is enabled, a diagnostic message indicating that a drive needs cleaning initiates Sun StorEdge EBS drive cleaning.

Warning messages indicate drive errors:

- Recoverable read or write errors
- Media at end of life
- Read-only tape format in the drive
- Periodic cleaning required

Critical messages are warnings that a drive might be disabled:

- Unrecoverable read or write errors
- Tape marked read-only
- Drive cleaning required immediately
- Drive prediction of hardware failure



Important: Critical messages require immediate attention to avoid data loss.

Informative and warning messages should automatically be cleared by **nsrmmd** when the reported issue is handled.

Critical messages relating to 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 StorEdge 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 *Sun StorEdge EBS 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. For more information about how to set the maximum consecutive error count, see "Media Handling Errors" on page 354.

Note: The drive must be manually reset to Enabled for the Sun StorEdge EBS software to use the device again.

To put a device in service mode, open the device's Properties window, and on the General tab, set Status Enabled to Service.

Device Parameter Settings

As of the Sun StorEdge EBS 7.3 release, device parameters can be set in the Sun StorEdge 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 StorEdge EBS server.

Environment variables can still be used to adjust several Sun StorEdge EBS device parameters, but the adjustment of environment variables should only be done by users who are intimately familiar with the server environment and performance tuning (for example, an Administrator who wants to tweak performance by changing a certain setting for all LTO devices on a particular Sun StorEdge EBS server).

The variables (and their equivalent names in the Administration interface) are described in the following sections.

Device Settings in the Sun StorEdge EBS Administration interface

To locate and change the device parameters in the Administration interface:

1. In the server's Administration interface, click the Devices button.
2. Click View, then select 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 shown:
   - Device Block Size
   - Device File Size
   - Device Load Time:
   - Device Eject Time
   - Device Poll Interval
   - Device Min Load Tries
   - Device Default Capacity
   - Device Tape Flags

When device parameters are set in this interface, it is not necessary to stop and restart the Sun StorEdge EBS server in order for the settings to take effect.

Device Settings as Environment Variables

The device-related environment variables are as follows:

- **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 StorEdge EBS server. The syntax for the device name must be all uppercase, with underscores (_) replacing blank spaces and hyphens. Thus, a device displayed in the Sun StorEdge EBS software as "8mm Mammoth-2" would be listed as:

```
8MM_MAMMOTH_2
```

Microsoft Windows users use the Control Panel's System applet to set these system environment variables. The Sun StorEdge EBS Backup and Recover Server service must be restarted for the environment variables to take effect.

UNIX/Linux users set these environment variables in the Sun StorEdge EBS startup script before the following line:

```
(echo 'starting Sun StorEdge EBS daemons:') >
/dev/console
```

and enter the variables in the following format:

ENV_VAR_NAME **=** *value*
export *ENV_VAR_NAME*

The Sun StorEdge EBS server must be restarted in order for the environment variables to take effect.

Table 39 on page 305 lists the location of Sun StorEdge EBS startup scripts on UNIX and Linux systems.

Table 39. Sun StorEdge EBS Startup Script Location on UNIX/Linux Systems

Operating System	Startup Script Location
Solaris, Linux, DYNIX/ptx, 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. 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, see "Compatible Block Size for UNIX and Microsoft Windows" on page 298.

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 StorEdge EBS software writes a filemark by closing and reopening the tape device, which takes a second or two. If this value is too small, throughput could be slowed.

On Microsoft Windows platforms, the Sun StorEdge EBS software writes asynchronous filemarks, therefore 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 second. 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 polls during load time. 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 StorEdge EBS server to send commands to tape devices.

CDI support can be set in the Sun StorEdge EBS Administration interface:

1. In the server's Sun StorEdge EBS Administration interface, click the Devices button.
2. In the menu bar, click View, then select 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 the CDI feature is enabled, it does the following:

- 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 such a situation, access the */nsr/debug* directory and create a file named *cdidisable*. Then restart the Sun StorEdge 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 the following 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 StorEdge EBS server 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.



Important: The CDI feature allows Sun StorEdge EBS software to collect better diagnostic information and facilitates when enabled. Only set or disable the CDI feature on the advice of an Sun Technical Support representative. If tape or SCSI issues occur while the CDI feature is enabled, go to *http://softwaresupport.emc.com*.

Nonrewinding Tape Device Usage (UNIX/Linux Only)

Tape drives used as storage devices must be accessed by nonrewinding device files. The Sun StorEdge 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 StorEdge 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/0mbn*:

- 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 StorEdge EBS software (for example */dev/rmt2.1*)



Important: Never change a device pathname from nonrewinding (*/dev/rmt/0cbn*) to rewinding (*/dev/rmt/0cb*). This would ensure that data could only be saved, never recovered, since 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 used with Sun StorEdge 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.

For information about setting the tape device block size in the operating system, refer to the operating system's documentation.

Silos

Silos and libraries are managed similarly by Sun StorEdge EBS software. For best results, have a good understanding of library operations before moving on to silos.

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 StorEdge 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.

For silo-specific media-management concerns, see "Media Management in a Silo" on page 355.

How Sun StorEdge EBS Software Interacts with a Silo

A Sun StorEdge EBS server acts as a client of the silo management software, which resides on the silo server. The Sun StorEdge EBS server communicates with the silo through the Silo Tape Library Interface (STLI), which must be installed on the Sun StorEdge EBS server that uses the silo.

To access the volumes and devices in a silo, the Sun StorEdge 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 StorEdge 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 StorEdge 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.

How to Install a Silo

To install a silo for use with Sun StorEdge EBS software:

1. Install the silo management software on the silo server.
2. Install the STLI on the Sun StorEdge EBS server, if necessary. For more information, refer to the documentation from the silo vendor.

   Note: For example, for a Sun StorEdge EBS server or storage node running Windows to control an STK silo, the **libattach** program must be installed.

   On UNIX and Linux systems, do not install the STLI library on the following models, because all the necessary software is installed when the Sun StorEdge 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 StorEdge EBS server is properly connected to the media devices in the silo.
4. Add the silo. See "How to Add (Configure) a Silo Device" on page 123.
5. License the Silo Support Module. See "Managing Product Licenses" on page 500.

Silo Device Naming Conventions

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 StorEdge EBS server can use. There is no way to get this information from the Sun StorEdge EBS server. To connect to more than one drive, find out 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 StorEdge 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 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 StorEdge 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.

How to Set up DAS to Work with Sun StorEdge EBS Software

To set up DAS to work with Sun StorEdge EBS software:

1. The silo administrator must configure DAS to accept commands from the Sun StorEdge EBS server or storage node computer.
2. The silo administrator must do one of the following:
    - Use the **dasadmin allocd** command to allocate one or more devices to the Sun StorEdge EBS server or storage node.
    - Configure the Sun StorEdge EBS server or storage node as an administrator, to be able to enter the **dasadmin allocd** command to allocate devices from the Sun StorEdge EBS server or storage node computer.

How to 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 StorEdge EBS software.

To find the names:

1. Set the following three environment variables:
   - DAS_SERVER, the hostname of the silo management server, which runs DAS
   - DAS_CLIENT, the hostname of the Sun StorEdge EBS server or storage node
   - ACI_MEDIA_TYPE, one of the following 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. Enter the following command to see a list of drives and the hostnames to which they are allocated:

   `dasadmin ld`

How to Release a Silo Device

When a silo device is configured for use with a Sun StorEdge EBS server, it is possible to restrict silo access only to the Sun StorEdge 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 the Devices button.
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 question window appears and asks whether to release devices.

5. Click Yes.

   The Library Operation window appears and displays the following message:

   ```
   The library operation has started.
   Please see the Monitoring->Operations screen for its status.
   ```

6. Click OK.
7. Repeat steps 1 through 6 for each device to be released.

Silo Device Cleaning

Do not enable automated cleaning for silos in the Sun StorEdge 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 40 on page 314 lists the environment variables to set.

Table 40. DAS and StorageTek Environment Variables (Part 1 of 2)

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 the following values: • 3480 • VHS • OD-Thick • 3590 • OD-Thin • CD • DECDLT • TRAVAN • 8MM • DTF • 4MM • BETACAM • D2 • AUDIOTAPE

Table 40. DAS and StorageTek Environment Variables (Part 2 of 2)

Silo Model	Environment Variables
	To specify the cartridge access port or range of ports for depositing or withdrawing volumes, use the following 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 or Linux systems:** CSI_HOSTNAME = *name_of_ACSLS_system* The following commands should also be running on the system and can be included in the Sun StorEdge 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 the Start > Programs > LibAttach menu tree. Once installed and configured, it starts on reboot.

How to Set the Environment Variables for UNIX and Linux Systems

To set environment variables for DAS and StorageTek silos for UNIX and Linux systems:

1. Locate the Sun StorEdge EBS startup script.
2. Using a text editor, open the Sun StorEdge EBS startup script and add the environment variables after the lines:

   ```
   (echo 'starting Sun StorEdge EBS daemons:') > /dev/console
   ```

   Enter the environment variables using the format:

   ENV_VAR = value
   export *ENV_VAR*
3. Restart the Sun StorEdge EBS server in order for the environment variables to take effect.

   Table 41 on page 316 lists the locations of the Sun StorEdge EBS startup scripts:

Table 41. Sun StorEdge EBS Startup Commands

Operating System	Startup Command
Solaris, Linux, IRIX	**/etc/init.d/networker**
HP-UX, HP Tru64 UNIX	**/sbin/init.d/networker**
AIX	**/etc/rc.nsr**

How to 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. Enter 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.

How to Refresh 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 StorEdge EBS servers.
3. Right-click the server, and select Refresh. (This can also be done by selecting the View menu, and then Refresh.)

How to Change 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 StorEdge EBS Libraries field, enter the appropriate time, in hours.
4. Click OK.

How to Stop a Device Operation

To stop a device operation that is in progress:

1. In the Sun StorEdge EBS Administration window, click the Monitoring button.
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.

How to Display Device Operations Messages

To display device operations messages:

1. In the Sun StorEdge EBS Administration window, click the Monitoring button.
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.

Library Sharing among Sun StorEdge EBS Hosts

The Sun StorEdge EBS software permits different Sun StorEdge EBS hosts (a Sun StorEdge 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.

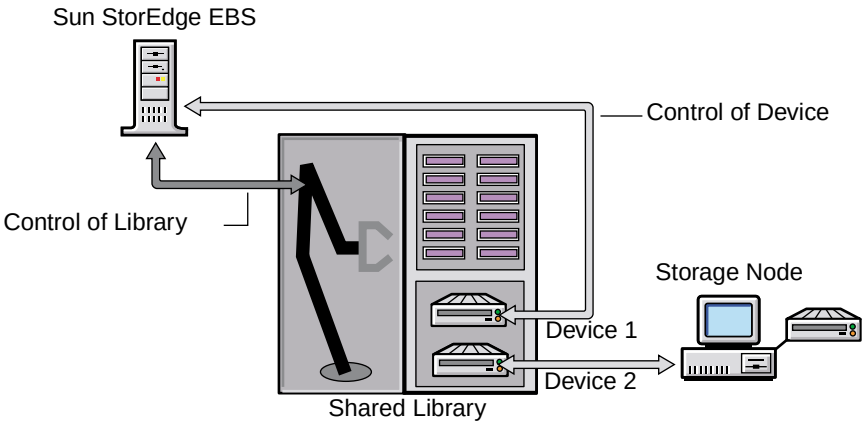
Dynamic Drive Sharing (DDS) does not support sharing libraries across Sun StorEdge EBS datazones.

How Library Sharing Works

Library sharing allows one Sun StorEdge EBS host to control the library's robotic arm, while other Sun StorEdge 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 StorEdge EBS host.

Figure 34 on page 319 shows how multiple Sun StorEdge EBS hosts can share library devices.

Figure 34. How Library Sharing Works



Sleep Periods for Library Tasks

Library device 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*. Attributes related to the sleep function are found in a library's Properties, on the Timers tab, when in View > Diagnostic Mode.

While sleeping, the library cannot receive or perform other operations. Without the sleep period, the loading or unloading of volumes might fail.

The Sun StorEdge EBS software automatically configures default sleep periods. Only change these values when troubleshooting a library's performance, or if a Sun StorEdge 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 42 on page 320.

Table 42. Library Resource Sleep Attributes

Attribute	Default Value
Load Sleep	5 seconds
Unload Sleep	5 seconds
Eject Sleep	5 seconds
Deposit Timeout	15 seconds
Withdraw Timeout	15 seconds
Cleaning Delay	60 seconds
Idle Device Timeout	0 minutes
Port Polling Period	3 seconds

Tips for Using Libraries

This section provides additional suggestions for using libraries effectively and reliably.

Library Notifications

The Sun StorEdge EBS server uses notifications to send messages about Sun StorEdge EBS events. Several preconfigured notifications, such as the following, inform 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.

For more information about notifications, see "Notifications" on page 541.

The Sun StorEdge 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 StorEdge EBS server can continue to back up or recover files. To check for messages, go to Monitoring > Alerts or Monitoring > Log.

Recycling Compared to Adding More Volumes

The Sun StorEdge EBS server saves files on volumes marked *appen* (appendable). If the volumes are marked *full*, they cannot receive backups.

If volumes are marked *full*, one of the following can be done with them:

- 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 StorEdge EBS server overwrites the data with new backups, but maintains the existing labels. For information about changing the volume mode, see "How to Change the Mode of a Volume" on page 334.

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 StorEdge EBS software become out of sync (for example, see "Example: Host Crash Requires User Intervention" on page 117. A library reset can be done using either the Administration interface or the command-prompt.

How to Reset a Library in the Administration Interface

To reset a library in the Administration interface:

1. In the Administration window, click the Devices button.
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.

   A question window appears and asks whether to reset the library.
5. Click Yes.

   The Library Operation window appears and displays the following message:

   ```
   The library operation has started.
   Please see the Monitoring->Operations screen for its status.
   ```
6. Click OK.

How to 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 StorEdge EBS software aware of these new drives, execute **nsrjb -HE** to reset the library. This must be done before the library has been configured. The **-E** option re-initializes 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 StorEdge EBS users, as the consequences of using them can be unknown. For information about these commands, refer to the *Sun StorEdge EBS 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. For more information, see "Media Pools" on page 203.

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 StorEdge EBS server when both the server and drives are shut down, restart all of the tape drives, either before or immediately after the Sun StorEdge 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 the following 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 StorEdge EBS software, causing the server to mishandle 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

Chapter 12: Media Management

This chapter covers the following topics:

- "Storage Management Operations" on page 325
- "Auto Media Management" on page 328
- "Volume Operations" on page 330
- "Media Handling Errors" on page 354
- "Media Management in a Silo" on page 355
- "Volume Save Sets" on page 361

The details describing 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 StorEdge EBS servers for which they have access permission.

Storage Management Operations

This section describes the components involved in the operation of storage volumes done through the Sun StorEdge Enterprise Backup Software server.

How the Sun StorEdge EBS Server Uses Volume Labels

A volume label is a unique internal code, applied by the Sun StorEdge EBS server, that initializes the volume for the server to use and identifies a storage volume as part of a specific pool. See "How the Sun StorEdge EBS Server Uses Media Pools" on page 203 for 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 StorEdge 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 StorEdge EBS server does the following:

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 associated with the chosen pool.

     If a recyclable volume from the same pool is relabeled, the volume label name and sequence number remain the same, but access to the original data on the volume is destroyed. The volume becomes available for new data.
   – An override volume name entered by the user.

How the Sun StorEdge EBS Server Selects a Volume

When a backup takes place, the Sun StorEdge 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 StorEdge 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 the following 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 the following 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. If multiple mounted volumes are appendable and session availability is not an issue, the volume with the oldest label time is selected.

2. 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.

   If a stand-alone device is being used and Auto Media Management is not enabled, the server sends a mount request notification.

3. If a library is in use and the above criteria have not been met, the server determines whether there is an unmounted, appendable volume. This includes newly labeled volumes. If multiple unmounted, appendable volumes are available, the volume with the oldest label time is selected.

4. If a library is in use and no unmounted, appendable volumes exist, the server determines whether there is an unmounted, recyclable volume. If Auto Media Management is enabled, the server relabels and mounts the volume.

   If Auto Media Management is not enabled, or if there are no appendable or recyclable volumes, the server sends a mount request notification.

For more information about appendable and recyclable volumes, see "Volume Operations" on page 330.

Data Recovery and Volume Selection

The Sun StorEdge 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 StorEdge EBS server displays all the volumes, in the order needed. During the recovery process, the server requests the volumes, one at a time.

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. For more information about changing the mode of a volume, see "How to Change the Mode of a Volume" on page 334.

Auto Media Management

The Auto Media Management feature gives the Sun StorEdge EBS server automatic control over media loaded in the storage device. When Auto Media Management is enabled during device configuration, the Sun StorEdge EBS server automatically does the following:

- Labels the volume (recognizes EDM labels and does not overwrite them).
- Mounts the volume.
- Overwrites volumes it considers to be unlabeled.

  The Sun StorEdge EBS server considers a volume to be unlabeled when any of the following conditions is true:

  – It has no internal label.

  – It is labeled with information other than a recognizable Sun StorEdge EBS label.

  – It is labeled with a Sun StorEdge 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.



Important: 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 StorEdge EBS volumes are shared among devices with different densities.

When the Auto Media Management feature is not enabled, the Sun StorEdge EBS server ignores unlabeled volumes and does not consider them for backup.



Important: The Sun StorEdge 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 StorEdge EBS server relabels the volume, the previously stored data is lost.

Auto Media Management with Existing Sun StorEdge EBS-Labeled Tapes

When Auto Media Management is used with Sun StorEdge EBS-labeled tapes that have not been recycled, the volumes must be removed from the media database before a utility such as **tar** is used to overwrite the labels. Also ensure that the tapes have been fully rewound before overwriting the labels. Auto Media Management can then properly relabel the tapes.

Auto Media Management with 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 if 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. At the same time, the Sun StorEdge 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 StorEdge EBS server checks to see whether the newly mounted volume is a candidate to receive data:
    a. If so, the write operation continues.
    b. If not, the Sun StorEdge EBS server continues to wait for a writable volume to continue the backup.

  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 StorEdge EBS server automatically ejects the volume after a few seconds.



Important: 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.

Auto Media Management with Libraries

Auto Media Management is not turned on for libraries during autoconfiguration. Auto Media Management for a library can be set by changing the library's properties after configuration:

1. In the server's Administration window, click the Devices button.
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 StorEdge Enterprise Backup Software server. Information about storage volumes is available for each device on a Sun StorEdge 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 the following 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 StorEdge 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 StorEdge EBS server reports on the status of volumes using values such as:

- *Volume Name*
- *Written*
- *%Used*
- *Location*
- *Mode*

Note: Many of the volume operations that can be performed with a right-click followed by a selection from a menu can also be done by using the icons on the icon toolbar.

Performing volume operations requires that the user have the correct permissions to use the Sun StorEdge EBS server and its storage nodes.

How to View Volume Status Information

To find information about a volume and its status:

1. In the server's Administration window, click the Media button.

   Media-related topics appear in the navigation tree.

2. Select Volumes.

   The Volumes detail table appears. Table 43 on page 332 lists the volume-related categories displayed in the Volumes Details Table.

Table 43. Volumes Detail Table (Part 1 of 2)

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, the following designations might appear: 1. (A) indicates an archive volume. 2. (R) indicates a volume that is considered read-only.
Barcode	Barcode label, if one exists.
Used	Indicates the exact number of bytes written to 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 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 44 on page 333 lists the Sun StorEdge EBS volume modes and their definitions. For information on changing volume modes, see "How to Change the Mode of a Volume" on page 334.
Expiration	Changing the expiration date is only possible from the command-prompt. Use the **nsrmm** command to do this.

Table 43. Volumes Detail Table (Part 2 of 2)

Category	Description
Location	Refers to an administrator-defined character field that describes 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 44. Volume Modes (Part 1 of 2)

Mode Value	Meaning	Description
appen	appendable	This volume contains empty space. Data that meets the acceptance criteria for the pool to which this volume belongs can be appended.
man	manual recycle	This volume is exempt from automatic recycling. The mode can be changed only manually. The manual recycle mode and the option to change it are available from the Volumes menu. The default recycle mode is Auto.
(R)	read-only	The save sets on this volume are considered read-only. The mode can be changed only manually. For more information, see "Read-Only Mode" on page 334.

Table 44. Volume Modes (Part 2 of 2)

Mode Value	Meaning	Description
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 *Sun StorEdge EBS Command Reference Guide*) for more information.

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.

To change the volume mode, see "How to Change the Mode of a Volume" on page 334.

How to Change the Mode of a Volume

To change the mode of a volume:

1. In the server's Sun StorEdge EBS Administration window, click the Media button.
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, offering the following choices:

   - Appendable
   - Read Only
   - Recyclable

4. Select a mode.
5. Click OK.

Volume Recycle

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.

How to Change a Volume's Recycle Policy

To change a volume's recycle policy:

1. In the server's Administration window, click the Media button.
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.

Volume Labels

The Sun StorEdge 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 StorEdge EBS software applies a label template to create a unique internal label for each volume. The Sun StorEdge 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 tape 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. For more information about label templates and pools, see "Chapter 7: Sorting Backup Data" on page 203.

Label Templates

Several preconfigured label templates are supplied with the Sun StorEdge EBS software. These preconfigured label templates cannot be deleted. For more information about label templates and preconfigured label template, see "Label Templates" on page 217.

How to Label or Relabel Library Volumes

Library volumes are labeled in the Devices task. To label a library volume:

1. In the Administration window, click the Devices button.
2. Open the Libraries folder.

   The Libraries detail table appears.
3. In the navigation pane, right-click the appropriate library and select Label.

   (It is also possible to select Label by right-clicking the library in the Libraries detail table, selecting Library Operations, then Label.)

   The selected library's details appear, including divided tables for devices and slots. The Label Library Media window also appears.

4. In the Slot Range area, enter the range of slots for the backup volumes to be labeled. To label only one backup volume, simply enter its slot number in both fields.

   The default label name of the volume (or of the first volume, if a range is selected) appears in the Label field.

5. The Default pool name appears in the Pools 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.

6. In the Operation Type area, select either Slow / Verbose (the default) or Fast / Silent.

   When Slow / Verbose is selected, the Supply Input option and icon on the Monitoring > Operations screen 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.

   See "How to Supply User Input" on page 485 for more information.

7. If the volume should not be recycled automatically, click Manual Recycle.

   If Manual Recycle 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 for manual recycle, the Sun StorEdge 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 simply by unselecting the Manual Recycle checkbox. The volume must be explicitly reset to use auto recycle. For more information, see "Volume Recycle" on page 335.

8. Click OK.

   The Library Operation window appears, stating that the library operation has started, and that its status can be tracked on the Monitoring > Operations screen.

9. Select Monitoring, then the Operations tab, to track the status of the label operation and, if Slow / Verbose was selected in the Label Library Media window, to confirm overwriting the existing volume label with a new label:

   a. Right-click the label operation in the Operations tab to confirm intent to overwrite the existing volume label with a new label, and select Supply Input.

      A Question window appears and displays the following message:

      `Label <labelname> is a valid Sun StorEdge EBS label. Overwrite it with a new label?`

   b. Click Yes to overwrite the label with a new label, or No to cancel the label operation.



Important: When a volume is relabeled, that volume is initialized and becomes available for writing again.

It takes the Sun StorEdge EBS server some time to label the volumes in a library, so consider labeling volumes before it is time to back up or recover files.

Empty Slots in Label Operations

Slots that have been intentionally left empty (such as bad slots) are skipped during labeling operations. The Sun StorEdge EBS software logs a message similar to: "Slot 5 empty, skipping."

Barcode Labels

To label a library volume with a barcode is one of the options available during automatic device configuration (see "How to Add a Library Device" on page 122), or can be set in the library's Properties tab after configuration.

Barcode labels are beneficial in two ways:

- They significantly shorten the time needed to inventory volumes.
- They provide improved labeling accuracy.

Barcode labels make volume inventory fast and efficient, because they eliminate the need to mount the volumes in a device. Instead, 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, because 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 StorEdge EBS server records and tracks the label in the media database. The Sun StorEdge EBS server uses barcode labels only to inventory volumes. A volume must have a label, but it need not have a barcode label.



Important: Libraries include hardware that reads barcode labels. The barcode information is then forwarded to the Sun StorEdge 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 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 StorEdge EBS media database must be correct or null. To view the location field, use the **mmlocate** command.

How to Configure a Library to Use Volumes with Barcodes

To select whether barcodes are used or matched after configuration:

1. In the Administration window, click the Devices button.
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 the following:
   - Bar Code Reader
   - Match Bar Code Labels
6. Click OK.

For more information about barcodes, see "Barcode Labeling Tips" on page 341.

How to 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 the Devices button.
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 StorEdge 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.

Inventory the volumes to ensure that the Sun StorEdge EBS server has the most current volume information.

Use Media > Volumes to match the correct volume labels to the barcode labels. This is the easiest way to make sure that the correct volume labels are placed onto the volumes with the existing 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 StorEdge 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 StorEdge 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 device's Properties window.

The following are tips for using barcode labels with the Sun StorEdge EBS software:

- When Sun StorEdge 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 StorEdge 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 StorEdge EBS volumes. Using identical labels defeats the purpose of using barcode labels, which is to speed up the inventory process and ensure labeling accuracy.
- Volume names must be unique on the Sun StorEdge 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 device's properties is enabled, the Sun StorEdge 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 begin the labeling process again, or disable the Match Barcode Labels attribute in the device'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 StorEdge 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_name*.001" and ending with "*server_name*.100", or as wide a range as necessary. Instructions for barcode labels should be provided with the library hardware documentation. Contact the hardware manufacturer with questions about barcode labels.

  A consistent labeling scheme helps better organize and track volumes. It also facilitates the inventory process if all of the volumes, rather than just some of them, 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.

How to Mount or Unmount a Volume in a Library

To mount a volume in a library:

1. In the Administration window, click the Devices button.
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 the following message:

      ```
      The library operation has started.
      Please see the Monitoring->Operations screen for 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 the following message:

      ```
      The library operation has started.
      Please see the Monitoring->Operations screen for its status.
      ```

   b. Click OK.

How to 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 that 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 the Devices button.
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 the following message:

   ```
   The library operation has started.
   Please see the Monitoring->Operations screen for 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.

How to Label and Mount a Volume in One Operation (Stand-Alone Tape Drive)

When more than one storage device is connected to the Sun StorEdge 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 StorEdge 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 the Devices button.
2. Manually insert an unlabeled or recyclable volume in the Sun StorEdge EBS server storage device, or ensure that a volume of this type is already present for the Sun StorEdge EBS server to access.
3. Select Devices in the navigation tree.

   The Devices details table appears.

4. Right-click the appropriate stand-alone device in the details table, and select Label.

   The Label window appears.

   a. Enter a unique label name, or accept the default name associated with the selected pool.

      If the volume is unlabeled, the Sun StorEdge 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 StorEdge 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 StorEdge 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 simply by unselecting the Manual Recycle checkbox. The volume must be explicitly reset to use auto recycle. For more information, see "Volume Recycle" on page 335.

d. The Mount After Labeling attribute is selected by default.

   The Sun StorEdge EBS server automatically labels the volume, and then mounts the volume into the device.

5. Click OK.

   If the volume is recyclable, a question window appears, warning that the named volume is about to be recycled, and asking whether to continue. Click Yes, if the volume should be relabeled and recycled.

After a volume is labeled and mounted in a device, the volume is available to receive data. Since the Sun StorEdge EBS label is internal and machine-readable, place an adhesive label on each volume that matches that internal volume label.

For information on using barcode labels, see "How to Configure a Library to Use Volumes with Barcodes" on page 339.

How to 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 unselect the Mount After Labeling attribute in the Label window.

How to Mount an Uninventoried Volume

To mount a volume that is not included in the library inventory, but which is a valid (properly labeled) Sun StorEdge EBS volume:

1. In the Administration window, click the Devices button.
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. Enter 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 Monitoring > Operations screen 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.

   For more information, see "How to Supply User Input" on page 485.

9. Click OK.

   The Library Operation window displays the following message:

   ```
   The library operation has started.
   Please see the Monitoring->Operations screen for its
   status.
   ```

   The Sun StorEdge EBS software then inventories the specified slot.

   Also see "How to Inventory Volumes in a Library" on page 350.

10. Mount the inventoried volume as described in "How to Mount or Unmount a Volume in a Library" on page 342.

Note: Unlabeled tapes may not be mounted for inventorying, because unlabeled tapes can only be mounted to be labeled. An attempt to mount an uninventoried volume 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. Select a time that is appropriate for the system configuration.

Selecting zero means that the device never times out and the tape must be ejected manually.

Using Libraries with a Volume Import and Export Capability

The Sun StorEdge 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, requiring the long operation of reinitializing the library. The Sun StorEdge 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.

How to Deposit a Volume by Using the Import/Export Feature

The following are general instructions. Specific instructions for working with a CAP can vary, depending on the library manufacturer. For specific instructions, it might be necessary to 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 the Devices button.
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.

   A Question window appears, asking whether to deposit the volume.
6. Click Yes.

   The Library Operation window displays the following message:

   ```
   The library operation has started.
   Please see the Monitoring->Operations screen for its
   status.
   ```

7. Click OK.

   After a while, a Question window appears, requesting that you load the cartridges into the ports and enter Yes to continue.

8. Click Yes.

   After a while, a Question window appears, asking whether to continue depositing volumes.

9. Click Yes to continue depositing volumes, or No, when done.

How to Withdraw 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 the Devices button.
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.

   A Question window appears, asking whether to withdraw the volume.

6. Click Yes.

   The Library Operation window displays the following message:

   ```
   The library operation has started.
   Please see the Monitoring->Operations screen for its status.
   ```

7. Click OK.
8. Check Monitoring > Log to see the result. A successful Eject / Withdraw operation ends with a Succeeded comment in the Log.

Inventorying Library Volumes

When the Sun StorEdge EBS software labels the contents of a library, it 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, Sun StorEdge 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 StorEdge 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 StorEdge EBS software must be made aware 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, because the Sun StorEdge EBS software automatically records the slot number in which each newly labeled volume is located.

The Sun StorEdge 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 are used, or if the library contents change often. For more information on using barcode labels, see "Barcode Labels" on page 338.

How to Inventory Volumes in a Library

To inventory volumes in a library:

1. In the Administration window, click the Devices button.
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. Enter 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 the following message:

   ```
   The library operation has started.
   Please see the Monitoring->Operations screen for its
   status.
   ```

8. Click OK.

If the volumes do not have barcode labels, the Sun StorEdge EBS software must mount each volume, read its label, and unmount it. In this case, the inventory process can take some time to complete.

Removal of 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 StorEdge 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 knowledge of the volume from the Sun StorEdge 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 StorEdge 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.

How to Delete Volume Data

To delete a volume:

1. In the Administration window, click the Devices button.
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, run the **nsrck** command from the command-prompt.

Note that volume deletion can also be done using the Media button, and then selecting Volumes.

For more information, refer to the **nsrmm** and **mminfo** UNIX man pages or the *Sun StorEdge EBS Command Reference Guide*.

How to Mark a Volume as Full for Off-Site Storage

When removing a volume from a library to store off-site, mark the volume as "full" so that the Sun StorEdge 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.

One could also change the volume's Location attribute to an informational message, such as "Move to off-site storage in September 2006".

To change the volume's Location attribute:

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

   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. Enter a location description.

   Note that this is descriptive information. Filling in this field does not send a volume anywhere.

5. Click OK.

Volume Cloning

For information about volume clone operations, see "Volume Cloning" on page 246.

Media Handling Errors

The architecture of device drivers can produce media handling errors. The Sun StorEdge EBS software automatically retries a failed operation such as a **mount** or **read** of a volume. The number of times the Sun StorEdge 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.

There could be several reasons why a mount or read operation might fail. Two examples follow:

- 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.

How to 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 StorEdge 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 the Devices button.

   The Devices detail table appears.
3. Right-click the drive to be reenabled, and select Properties.

   The Properties window appears.
4. On the General tab, in the Status area, set Enabled to Yes.
5. Click OK.

If the disabled drive is part of a library, it might be necessary to reset the device as follows:

1. From the command-prompt, change the path to the directory that contains the Sun StorEdge EBS binaries.
2. Enter the following 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 StorEdge EBS software from overwriting volumes used by other programs.

This section addresses the following:

- How volumes are specified for the Sun StorEdge EBS software to use.
- How volumes are mounted in the devices.
- How volumes are identified and counted in a silo.

Silo Slot Numbers

In a library, the Sun StorEdge EBS software specifies many functions by slot number. A library has a fixed number of slots, and Sun StorEdge EBS software uses the slot number to refer to a volume's physical location.

A silo uses the same idea, 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. For more information, refer to the **nsrjb** UNIX man page or the *Sun StorEdge EBS Command Reference Guide*.

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 StorEdge EBS pool.

For more information, see "Mounting and Unmounting Volumes" on page 342.

Labeling a Silo Volume

The Sun StorEdge EBS labels for volumes in a silo include both a regular Sun StorEdge 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. For instructions on how to label silo volumes, see "Volume Labels" on page 336 and "Barcode Labels" on page 338.

The use of barcodes, and matching barcode labels and Sun StorEdge EBS volume labels, are both available for a silo. The Match Barcode Labels attribute can be left unselected, but ensure that the Barcode Reader attribute is selected. When both attributes are selected, the internal volume label that Sun StorEdge EBS software writes on the media of each volume matches the barcode label on the outside of the volume. When the labels match, it is easier to track volumes, but the Sun StorEdge 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 StorEdge EBS software (or another application) associates with some of the data.

Using a Silo with Volume Import/Export Capability

Sun StorEdge 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 StorEdge EBS software does not automatically inventory the volume after a deposit.

Either Sun StorEdge 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 StorEdge EBS software cannot be used to insert volumes.

To issue deposit and withdraw commands, enter the following:

- To add and deposit volumes:

  `nsrjb -a -T` *tags* **`-d`**
- To remove and eject / withdraw volumes:

  `nsrjb -x -T` *tags* **`-w`**

where *tags* specifies the tags or barcodes of volumes in a remote silo.

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. For more information, refer to the **nsrjb** UNIX man page or the *Sun StorEdge EBS Command Reference Guide*.

Allocating (Adding) Silo Volumes

When volumes are added, the Sun StorEdge EBS server is directed to the volumes it can use.



Important: 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 StorEdge 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 StorEdge EBS software does not provide a method for setting up this sort of protection; the silo management software must configure it.

Adding a volume causes the Sun StorEdge 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 StorEdge EBS software.

How to Add a Silo Volume

1. In the Administration window, click the Devices button.
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. For example, to name four tapes A01B, A02B, A03B, and A04B, enter A0, 1-4, and B.
   - The List option allows the entry of barcode IDs, separately. For example, enter the name for each tape: A01B, A02B, A03B, and A04B.
6. Enter 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 the following message:

   ```
   The library operation has started.
   Please see the Monitoring->Operations screen for its status.
   ```

8. Click OK.

   On return to the Library's details table, the added volumes will be shown. Note that silos do not display empty slots.

If the particular silo model does not automatically deposit the volume, then place the volumes in the insert area, and select Deposit from the menu.

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*

For more information on STLIs, see "How Sun StorEdge EBS Software Interacts with a Silo" on page 310.

Deallocating (Removing) Silo Volumes

When a Sun StorEdge EBS software STL volume in a silo is no longer needed, the volume can be deallocated from the library. 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 StorEdge EBS media database remain intact. If the volume is allocated again, later, Sun StorEdge EBS software can retrieve the data from it.

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 StorEdge EBS server access the data. If the volume is to be removed from the silo for offsite storage, it must be removed with Sun StorEdge EBS software and then ejected from the silo by using the silo management software.

How to Remove a Silo Volume

To remove a silo volume:

1. Unmount the volume from the device. For instructions on unmounting volumes, see "Mounting and Unmounting Volumes" on page 342.
2. In the Administration window, click the Devices button.
3. Open the Libraries folder in the navigation tree.
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 to create a list of barcode IDs. For example, to name four tapes A01B, A02B, A03B, and A04B, enter A0, 1-4, and B.
   - The List option allows the entry of barcode IDs, separately. For example, enter the name for each tape: A01B, A02B, A03B, and A04B.
7. Enter the appropriate volume identifiers in the Barcodes field.
8. Click OK.

   The Library Operation window displays the following message:

   ```
   The library operation has started.
   See the Monitoring->Operations screen for its status.
   ```
9. Click OK.

   The removed volumes are no longer listed in the Library details table.

For more information, see "How Sun StorEdge EBS Software Interacts with a Silo" on page 310.

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 StorEdge 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. For information about inventorying a library, see "Inventorying Library Volumes" on page 350.

The Sun StorEdge EBS software examines all of the volumes in the silo and compares the new list of volumes to the Sun StorEdge EBS media database. Then the Sun StorEdge EBS software produces a message listing any volumes located in the silo that are not in the media database.

When the Sun StorEdge 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 StorEdge EBS media database, the Sun StorEdge EBS software does not need to load the volume into the device and read. The inventory proceeds rapidly. If, however, the Sun StorEdge 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.

Two places exist where the details of save sets can be viewed:

- In the Volume Save Sets window
- In the Save Sets detail table, which offers two tabs:
  - The Query Save Set tab
  - The Save Set List tab

How to View 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 the Media button.

   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.

   The Volume Save Sets window appears. Table 45 on page 362 shows the attributes and their descriptions.

4. Click OK to leave the Volume Save Sets window.

Table 45. Volume Save Sets Window (Part 1 of 2)

Category	Description
Client	Name of the Sun StorEdge 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 *ad hoc* backups, the level is blank.
Status	Type of save set. For a listing of save set values and descriptions, see Table 30 on page 249.

Table 45. Volume Save Sets Window (Part 2 of 2)

Category	Description
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 StorEdge EBS storage node.

How to Change Save Set Status within the Volume Save Sets Window

While the Volume Save Sets window is displayed:

1. Select a save set.
2. Click Change Status.

   The Change Save Set Status window appears.
3. Select one of the following:
   - 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.

How to View 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 the Media button.

   Media-related topics appear in the navigation tree.

2. Select Save Sets.

   The Save Sets detail table appears and contains two tabs for configuring save set queries and listing save set details:

   – The Query Save Set tab

   – The Save Set List tab

   Select either tab. See "The Query Save Set Tab" on page 364 and "The Save Set List Tab" on page 365 for 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 the following settable query fields:

- Query Parameters area:
  – Client Name
  – Save Set
  – Save Set ID
  – Volume
  – Pool
  – Copies
  – Save Time (a range)
- Status area:
  – Browsable
  – Scanned-In
  – In-Progress
  – Suspect
  – Recyclable
  – Recoverable
  – Aborted

- Maximum Level area:
  - Full
  - 1 through 9
  - All

If no save sets are found that match the query parameters, an error message appears when exiting 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 the following tabular information:

- Save Set
- SSID (Save Set ID)
- Level
- Status
- Volume Name
- Client
- Size
- Files
- Pool
- Time

Chapter 13: Recovering Data

This chapter covers the following topics:

- "Choosing a Recovery Method" on page 367
- "Recovering by File Selection" on page 368
- "Recovering by Save Set Selection" on page 379
- "Disaster Recovery" on page 386
- "Recovering Encrypted Data" on page 386
- "Directing Recoveries to Another Client" on page 387
- "Recovering Online Indexes for a Sun StorEdge EBS Server" on page 395
- "Recovering the Console Database" on page 399
- "Moving a Sun StorEdge EBS Client to a New Sun StorEdge EBS Server" on page 399
- "Recovering the Windows System Configuration" on page 400
- "Recovering Windows Volume Mount Points" on page 408
- "Recovering Windows DHCP and WINS Databases" on page 411
- "Restoring a Windows Content Index Server" on page 412

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. Directed recovery is not supported for save set recovery. For more information about directed recoveries, see "Directing Recoveries to Another Client" on page 387.
- 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. For more information about the **scanner** program, see "How to Restore a Save Set Entry in the Online Indexes" on page 397.
- Memory on the recovery host is an issue. Recovering by save set selection requires less memory than recovering by file selection.

13

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. For more information about browse policies, see "About Browse and Retention Policies" on page 169.

If backed-up data has exceeded its browse policy, the data may still be recovered by using save set recovery. For more information about save set recovery, see "Recovering by Save Set Selection" on page 379. Alternatively, users could 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. For more information about re-creating the client file index, see "How to Restore a Save Set Entry in the Online Indexes" on page 397.

Users can also recover files from other client computers by using directed recovery. For more information, see "Directing Recoveries to Another Client" on page 387.

How to Recover Data on UNIX

To recover data to a UNIX client:

1. Start the **nwrecover** program by entering the following command:

   ```
   nwrecover
   ```

   - Use the **-s** *server* option with **nwrecover** to specify a particular Sun StorEdge EBS server on the network.
   - Use the **-c** *client* option to specify a particular Sun StorEdge 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 372.

3. Click the Start button to begin the recovery.

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 there is a naming conflict during recovery, you are prompted to choose a workaround. For more information, see "Resolving Recovered Data Conflicts" on page 374.

If the recovery is successful, a message similar to the following appears:

```
Received 1 file(S) from NSR server server
Recover completion time: Tue Jan 21 08:33:04 2006
```

How to Recover Data on Windows

To recover data to a Windows client:



Important: If VSS is licensed and enabled, the databases must be taken offline before recovering Micorsoft Exchange Server or SQL Server data.

1. In the Sun StorEdge EBS User program, click the Recover button to open the Source Client dialog box.
2. Select the source client whose 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 file name.

     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. See "Recovering an Earlier Version of a File" on page 372.
6. Click the Start button to begin the recovery.

It takes the Sun StorEdge 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 using VSS, you must restart the recovery process.

If there is a naming conflict during recovery, you are prompted to select a workaround. For more information about naming conflicts, see "Resolving Recovered Data Conflicts" on page 374.

If the recovery is successful, a message similar to the following appears:

```
Received 1 file(S) from NSR server server
Recover completion time: Tue Jan 21 08:33:04 2006
```

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 StorEdge EBS log file.

How to Recover 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*, enter the following:

```
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 2005
recover> quit
```

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

Note: To avoid using the Windows version of **recover.exe** on Windows operating systems, do one of the following:

- Include <*Sun StorEdge EBS_install_path*> \ *bin* \ *recover.exe* at the command-prompt
- Ensure that <*Sun StorEdge 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. For more information about browse policies, see "About Browse and Retention Policies" on page 169.

If backed-up data has exceeded its browse policy, the data may still be recovered by using save set recovery. For more information about save set recovery, see "Recovering by Save Set Selection" on page 379. Alternatively, users could 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. For more information about re-creating the client file index, see "Restore a Save Set Entry in the Online Indexes" on page 397.

How to Recover an Earlier Version of a File

Note: 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 Sun StorEdge EBS User program, click the Recover button 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 entering 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 368.

   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 Browse Time from the Change menu.
2. In the Change Browse Time dialog box, select a new day within the calendar. Select Previous or Next to change from the current month.
3. In the optional Browse 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 368.

   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.

Resolving Recovered Data Conflicts

By default, the Sun StorEdge EBS server attempts to recover data to its original location. If a name conflict occurs, the Sun StorEdge EBS software can respond either automatically or on an file-by-file basis.

How to Resolve Data Conflicts on UNIX

When you start a recovery by using the **nwrecover** program, the Conflict Resolution window opens and asks whether to prompt when a conflict occurs.

- Select Yes to select a resolution option each time there is a conflict.
- Select No to select a resolution option that is used each time there is a conflict.

Table 46 on page 374 lists the available conflict resolution options.

Table 46. Naming Conflict Resolution Options

Option	Action
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.

How to Resolve Data Conflicts on Windows

By default, the Naming Conflict dialog box appears each time there is a filename conflict during a recovery. Table 46 on page 374 lists the conflict resolution options available during a recovery.

Table 47. Naming Conflict Resolution Options (Part 1 of 2)

Option	Action
Rename Recovered File	Renames the recovered files. For more information, see "Rules for Renaming File on Windows" on page 375.
Discard Recover File	Discards the recovered file and keeps the existing file.

Table 47. Naming Conflict Resolution Options (Part 2 of 2)

Option	Action
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 StorEdge EBS server uses the following 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 StorEdge EBS server fails to rename the file.
  - If the Sun StorEdge EBS User program is being used, the user is prompted to enter a name for the file.

How to Resolve Naming Conflicts Automatically on Windows

You can configure the Sun StorEdge 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 Sun StorEdge EBS User program, click the Recover button 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 47 on page 374.
6. Click OK.

Another way to avoid name conflicts is to relocate the recovered data prior to recovery. For more information, see "Relocating Recovered Data" on page 376.

Relocating Recovered Data

By default the Sun StorEdge 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 StorEdge EBS server to relocate the recovered file prior to the recovery.

How to Relocate Recovered Data on Windows

To relocate recovered data:

1. In the Sun StorEdge EBS User program, click the Recover button 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, enter the directory where the data should be relocated.

   Note: Provide the full pathname of the relocation directory.
6. Click OK.
7. Perform a data recovery as described in "Recover Data on Windows" on page 370.

How to Relocate Recovered Data on UNIX

To relocate recovered data:

1. In the **nwrecover** program, select Relocate from the Options menu.
2. In the Relocate To field, enter the directory where the data should be relocated.

   Note: Provide the full pathname of the relocation directory.

3. Click OK.
4. Perform a data recovery as described in "How to Recover Data on UNIX" on page 369.

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.

How to View Version Details on Windows

To view the version details of backed-up data:

1. In the Sun StorEdge EBS User program, click the Recover button to open the Source Client dialog box.
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.

How to 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 View menu, select Versions.

   Versions are sorted according to the backup time (the most recent backup first).
3. To display additional file details, Select File Details from the View menu.

How to Search 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 Sun StorEdge EBS User program, click the Recover button 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. Enter the name of the file or directory. Use wildcards to expand the search.

On UNIX

To search for a recoverable file:

1. In the **nwrecover** program, select Search from the View menu.
2. Enter the name of the file or directory. Use wildcards to expand the search.
3. Enter a path to begin the search.

How to View 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 Sun StorEdge EBS User program, click the Recover button.
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 and mark 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 Volumes from the View menu.

Recovering by Save Set Selection

Save set recovery allows 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. For more information about deciding whether to perform a save set recovery or a recovery by file selection, see "Choosing a Recovery Method" on page 367.

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. For information about retention policies, see "Chapter 5: Browse and Retention Policies" on page 169. For information about recovering data that has exceeded its retention policy, see "How to Restore a Save Set Entry in the Online Indexes" on page 397.



Important: 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. For more information about the relationship between full backups, backup levels 1 to 9, and incremental backups, see "Backup Levels" on page 154.

How to Recover a Save Set on Windows



Important: Only members of the Windows Administrators group have permission to perform a save set recovery on Windows. In the case of a remote Sun StorEdge 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 Sun StorEdge EBS 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.

How to Recover a Save Set on UNIX



Important: Only *root* has the permission to perform a save set recovery on UNIX. In the case of a remote Sun StorEdge 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 Save Set Recover from the Options menu.
2. In the Client text box, enter or select the appropriate client.
3. From the Save Set Name list, select the name of the appropriate save set.

4. In the Instances list box, select the particular version of the save set (if there are multiple versions).
5. Click Recover.
6. In the Save Set Recover Status dialog box, click Start to begin the recovery.

   The recovery status appears in the Save Set Recover Status dialog box.

How to Recover a Save Set from the Command-prompt

To perform a save set recovery from the command-prompt, use the recover command with the following syntax:

recover -s *server* **-S** *ssid* **-d** *directory_path*

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

How to Recover Individual Files or Directories from a Save Set on Windows

To recover specific data from a save set.

1. In the Sun StorEdge EBS 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. Enter 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.

How to Recover Individual Files or Directories from a Save Set on UNIX

To recover specific data from a save set on UNIX.

1. In the **nwrecover** program, select Save Set Recover from the Options menu.
2. In the Client text box, enter or select the appropriate client.
3. From the Save Set Name list, select the name of the appropriate save set.
4. In the Instances list box, select the particular version of the save set (if there are multiple versions).
5. Click Recover to open the Save Set Recover Status dialog box.
6. Enter the names of the files and directories (including the full case-sensitive paths) to be recovered in the Paths to Recover text box. Click Add after each entry to add the entry to the list box.

   Note: Place a forward slash after a specified directory. For example, */usr/home/test* recovers any directories named *test** found in */usr/home* (*test1*, *test2*, etc.), while */usr/home/test/* recovers only the *test* directory.
7. Click Start to begin the recovery.

   The recovery status appears in the Save Set Recover Status dialog box.

How to 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. For information about using the **scanner** program to re-create the media database, see "How to Restore a Save Set Entry in the Online Indexes" on page 397.

On Windows

To view a list of the volumes required for a save set recovery:

1. In the Sun StorEdge EBS 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 Save Set Recover from the Options menu.
2. In the Client text box, enter or select the appropriate client.
3. From the Save Set Name list, select the name of the appropriate save set.
4. From the Instances list, select the particular version of the save set (if there are multiple versions).
5. Click Recover to open the Save Set Recover Status dialog box.
6. Click Volumes to display the volumes required for the recovery.

How to Relocate Recovered Save Set Data

While in the Save Set Recover Status dialog box, you can instruct the Sun StorEdge EBS server to 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 Sun StorEdge EBS 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. Enter the path in the Relocate Recovered Data To text box or click Browse to select the destination directory.
5. Click OK.

On UNIX

To relocate recovered save set data:

1. In the **nwrecover** program, select Save Set Recover from the Options menu.
2. In the Client text box, enter or select the appropriate client.
3. From the Save Set Name list, select the name of the save set to recover.
4. From the Instances list, select the particular version of the save set (if there are multiple versions).
5. Click Recover to open the Save Set Recover Status dialog box.
6. Click Recover Options to open the Save Set Recover Options dialog box.
7. Enter the path in the Relocate Recovered Data To text box.
8. Click OK.

How to Resolve Conflicts with Recovered Save Set Data

You can configure the Sun StorEdge EBS server to resolve conflicts during the recovery.

On Windows

To resolve data recovery conflicts with save set data:

1. In the Sun StorEdge EBS 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 48 on page 385 lists the options and actions taken for the Duplicate File Resolution.

Table 48. Naming Conflict Resolution Options

Option	Action
Always Prompt	Indicates that automatic naming conflict resolution is disabled.
Rename the Recovered File	Renames the recovered files. For more information, see "Rules for Renaming File on Windows" on page 375.
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

To resolve data recovery conflicts with save set data:

1. In the **nwrecover** program, select Save Set Recover from the Options menu.
2. In the Client text box, enter or select the appropriate client.
3. From the Save Set Name list, select the save set to recover.
4. From the Instances list, select the particular version of the save set (if there are multiple versions).
5. Click Recover to open the Save Set Recover Status dialog box.
6. Click Recover Options to open the Save Set Recover Options dialog box.
7. Select the appropriate option in the Duplicate File Resolution attribute. Table 46 on page 374 lists the options and their meanings.
8. Click OK.

Disaster Recovery

Recovering data from Sun StorEdge 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 *NetWorker Disaster Recovery Guide*. For the most recent disaster recovery information, refer to the *Sun StorEdge EBS Release Supplement*.

Note: Keep a printed copy of this documentation in a safe, convenient place.



Important: 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.

Recovering Encrypted Data

This section describes considerations for recovering encrypted data. Backup data can be encrypted using the **aes** Application Specific Module (ASM) or the **pw2** ASM. For more information about encrypting backup data, see "Encrypting Backup Data" on page 74.

Recovering AES 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 StorEdge EBS server. For more information, see "How to Set the AES Pass Phrase for a Sun StorEdge EBS Server" on page 75.

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*

  nwrecover -p *pass_phrase*

  winworkr -p *pass_phrase*

You can enter multiple pass phrases with the **-p** option, for example:

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 with no data. However, if unencrypted data is also selected for recovery, it will be recovered.

Recovering PW2 Encrypted Data

To recover backup data that was encrypted with the **pw2** ASM, specify the password that was in effect when the data was backed up. For information about password-protecting data with the **pw2** ASM, see "Password Protecting Windows or NetWare Backup Data" on page 73.

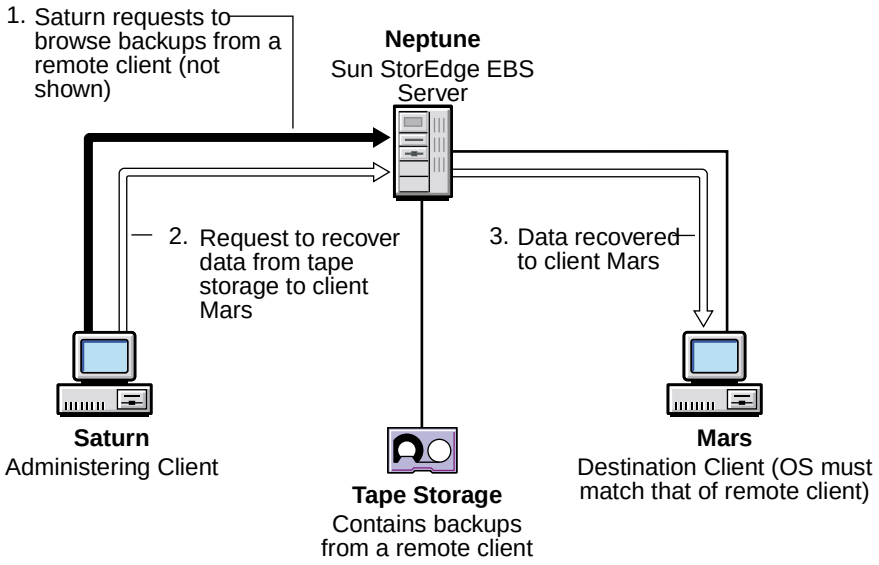
Directing Recoveries to Another Client

Directed recovery allows a Sun StorEdge EBS administrator to remotely recover backed-up data from a Sun StorEdge EBS client (the source client) and direct it to another Sun StorEdge EBS client (the destination client).

Note: Directed recoveries are supported for file selection recoveries only. Directed recoveries are not supported for save set recoveries.

Figure 35 on page 388 shows a user on client *saturn* performing a directed recovery of data from a remote client (source) to the client *mars* (destination).

Figure 35. A Directed Recovery from a Remote Client



Uses for Directed Recovery

A directed recovery is particularly useful in the following 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 StorEdge EBS client subject to the following guidelines:

- 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 StorEdge 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.
- Sun StorEdge EBS servers and clients running Microsoft Windows must be in a Microsoft Windows Domain (not a Workgroup), *or* the Sun StorEdge EBS server and the destination client must have the same username and password.
- For Sun StorEdge EBS clients running Microsoft Windows XP Professional or Microsoft Windows NT, the Sun StorEdge 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 information about enabling File and Print Sharing, refer to the Microsoft Windows documentation.



Important: Directed recovery of SYSTEM or VSS SYSTEM save sets is not supported.

How to Enable Directed Recoveries

The following tasks describe how to enable directed recoveries:

- "Task 1: Add Privileges to Members of the User Group" on page 390
- "Task 2: Set the Disable the Directed Recover Attribute" on page 392
- "Task 3: Add a Client Computer to the Remote Client's servers File" on page 392

Task 1: Add Privileges to Members of the User Group

Directed recovery is a restricted Sun StorEdge EBS function available only to users who have the Remote Access All Clients privilege.



Important: Be selective in granting the Remote Access All Clients privilege. This privilege enables 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 StorEdge EBS server and client in the enterprise prior to accepting data recovery requests. For general information about access requirements, see "Managing Server Access" on page 530.

Access requirements to perform directed recoveries are as follows:

- Users in the Administrators group on the Sun StorEdge EBS server are automatically granted the necessary privileges. For see the specific access privileges that are required if a user is not a member of the Sun StorEdge EBS Administrators group, see Table 49 on page 391.

  Notes:

  The user must be in the Administrators group on the Sun StorEdge EBS server when working with pre-release 7.0 Sun StorEdge EBS servers.

  The user must have the Remote Access All Clients privilege on the Sun StorEdge EBS server. If the user does not have Remote Access All Clients privileges, then the *user@destination_client* must be listed in the source client's Remote Access list. For information about editing a client, see "Edit a Client" on page 568.

Table 49 on page 391 shows the access permissions and authorizations required to perform various combinations of directed recoveries.

Table 49. Access Permissions Used in Directed Recoveries (Part 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.
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.

Table 49. Access Permissions Used in Directed Recoveries (Part 2 of 2)

Administering Client	Source Client	Destination Client	Who May Perform a Directed Recovery
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.

Task 2: Set the Disable the Directed Recover Attribute

Directed recovery can be turned on or off for Sun StorEdge 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 this attribute to Yes. Setting this attribute set to Yes, disables directed recoveries to the client.

For information about editing the *nsrla.res* database, see "How to Edit the NSRLA Database" on page 570.

Task 3: Add a Client Computer to the Remote Client's servers File

To add a computer name to the remote Sun StorEdge EBS client's *servers* file, see "How to Edit the servers File" on page 585.

How to Perform a Directed Recovery between Windows Clients

To perform a directed recovery:

1. Start the Sun StorEdge EBS User program, by using the **-s** option to connect to the Sun StorEdge EBS server to which the source client data is backed up:

   winworkr -s *server*

Note: Before you can perform directed recovery of a virtual cluster client, the Client resource must be configured on the Sun StorEdge EBS server. For instructions on creating a Client resource, see "Client Configuration" on page 568. Be sure that the Client resource Remote Access attribute contains *both* of the cluster node names.

2. Click the Recover button 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.

Note: Special data handling features, such as compression and PW2 encryption, are not available for the items you select.



Important: 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.

How to Perform a Directed Recovery between UNIX Clients

The **nwrecover** program does not support directed recovery. Use the **recover** command to perform a directed recovery for UNIX clients. The **recover** command has a variety of options that you can add to modify the basic browse and recover functionality. Table 50 on page 393 lists the options available with the **recover** command.

Table 50. Recover Options (Part 1 of 2)

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.

Table 50. Recover Options (Part 2 of 2)

Option	Meaning
-c *client*	Specifies the name of the source client.
-R *recover_destination*	Specifies the name of the destination client.
-i (N,Y, R)	Specifies what the Sun StorEdge 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.

For complete information about the **recover** command, refer to the *NetWorker Command Reference Guide* or the UNIX man pages.



Example: 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.

Enter the following command:

`recover -s venus -c mars -R jupiter -iY` *directory_name*

13

How to Perform a Directed Recovery of Specific Files and Directories on UNIX

To perform a directed recovery of specific files and directories:

1. Enter the **recover** command by using the following syntax:

   `recover -s` *server* **`-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 entering the following 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 entering the following command:

   ```
   recover> list
   ```

6. Start the directed recovery by entering the following command:

   ```
   recover> recover
   ```

Note: You can run the **recover** command from a Windows client to administer directed recoveries between Sun StorEdge 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. To avoid use of the Windows version of **recover.exe**, do one of the following:

- Include <*Sun StorEdge EBS_install_path*> \ *bin* \ *recover.exe* at the command-prompt
- Ensure that <*Sun StorEdge EBS_install_path*> \ *bin* is listed before *%SystemRoot%* \ *System32* in the $PATH environment variable.

Recovering Online Indexes for a Sun StorEdge EBS Server

This section includes procedures for recovering the client file index and the media database.

How to 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, enter the following:

      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 396.

Note: To completely remove a client file index, enter **nsrck -R -Y** *client_name*, then enter **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 *NetWorker Command Reference Guide* or the UNIX man pages.

To recover a client file index for a specific date, enter the following command:

nsrck -t "*time mm/dd/yyyy***" -L7** *client_name*



Example: Recovering a Client File Index from a Specific Time and Date

The following command recovers a client file index from 4 p.m., March 2, 2006:

nsrck -t "04 pm 03/02/2006" -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 *NetWorker Command Reference Guide* or the UNIX man pages.

How to Restore a Save Set Entry in the Online Indexes

The following 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.

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 StorEdge EBS server in the media database (which means that the save set's retention policy has not expired), follow these steps to recover the save set's entry back into the client file index:

1. Log in as root or Window Administrator and run the **mminfo** command at the command-prompt:

   mminfo -a -v *volume_name*

   where *volume_name* is the volume that contains the client file index.

2. From the **mminfo** output, find the *ssid* that contains the required file. Make sure it is not the bootstrap *ssid*.

3. Enter the **nsrmm** command with a new browse time:

   nsrmm -S *ssid* **-w** *browse_time*

   where *browse_time* is a time in the future.

4. Enter 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 you believe 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 the following 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 StorEdge 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. For information about browsing for a file, see "Recovering by File Selection" on page 368.

Recovering the Console Database

The Console database contains management data such as report information. For information about recovering a corrupted Console database, refer to the *NetWorker Disaster Recovery Guide*.

Moving a Sun StorEdge EBS Client to a New Sun StorEdge EBS Server

This section describes how to move a Sun StorEdge EBS client to a new Sun StorEdge EBS server without losing the ability to recover the client files that were backed up on the old Sun StorEdge EBS server.

How to Move a Sun StorEdge EBS Client to a New Sun StorEdge EBS Server

To move a client to a new Sun StorEdge EBS server:

1. Record the client ID attribute of the Sun StorEdge 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 StorEdge EBS server, create a new client.
   a. In the Name attribute, enter 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 the following format for the name of the client:

      ~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, enter the following name:

      `~jupiter-1`
   b. Click the Globals (1 of 2) tab.
   c. In the Client ID attribute, enter 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 StorEdge EBS server.

   You can now recover data that was backed up when the Sun StorEdge EBS client was set up on the old Sun StorEdge 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 following step 2. However, run **scanner -i** again after creating the Client resource to import save set information into the client file index.

13

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, the following 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: For Windows NT 4.0, only the SYSTEM STATE save set is backed up.

The following VSS SYSTEM save sets are backed up if VSS is licensed and enabled:

- VSS SYSTEM BOOT
- VSS SYSTEM FILESET
- VSS SYSTEM SERVICES
- VSS USER DATA
- VSS OTHER
- VSS ASR DISK

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 Sun StorEdge EBS User program or from the command-prompt, you must have local Windows Administrator privileges.



Important: 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, the following 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.



Important: Sun StorEdge 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.

Note: You cannot perform a *directed* recovery of a VSS SYSTEM save set from the command-prompt.

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 Sun StorEdge EBS clients running a 32-bit version of Windows, approximately 500 MB of extra disk space is usually sufficient.
- For Sun StorEdge EBS clients running a 64-bit version of Windows, approximately 1 GB of extra disk space is usually sufficient.

The default location for the restored system temporary files is a system drive where the original files reside. For the VSS SYSTEM save sets, the temporary files are placed in the *temp* directory on the system drive. For information on how to expand the available space in the *temp* directory by moving it to another large partition, refer to the Microsoft documentation.

When recovering SYSTEM or VSS SYSTEM save sets by using the Sun StorEdge EBS 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.

How to 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 the following order:
   - SYSTEM DB
   - SYSTEM FILES
   - SHAREPOINT
   - SYSTEM STATE

   The command used to recover each save set should look similar to the following:

   Sun StorEdge EBS_install_path`\bin\recover.exe -iY`
   [`-s` *Sun StorEdge 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.

How to 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 the following order:
   - VSS SYSTEM SERVICES
   - VSS SYSTEM FILESET
   - VSS USER DATA
   - VSS OTHER
   - VSS SYSTEM BOOT

   The command used to recover each save set should look similar to the following:

   Sun StorEdge EBS_install_path`\bin\recover.exe -iY`
   [`-s` *Sun StorEdge 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 StorEdge EBS log file.

How to Recover VSS SYSTEM Save Set Components from the Command-prompt

You can back up and recover individual components (writers) within the following VSS SYSTEM save sets:

- VSS SYSTEM SERVICES
- VSS USER DATA
- VSS OTHER

To recover selected components, enter the following at the command-prompt:

Sun StorEdge EBS_install_path`\bin\recover.exe -iY` [`-s` *Sun StorEdge EBS_server_name*] [`-t` *browse_time*]`-N "`*VSS_SYSTEM_SAVESET_NAME*`:\`*component_name*`"`

Place a semicolon (;) between multiple component names.

For example, to recover the Event Log Writer writer and the WMI Writer writer, enter:

```
C:\Program Files\Legato\nsr\bin\recover.exe -iY
-s jupiter -N "VSS SYSTEM SERVICES:\Event Log Writer;WMI
Writer"
```

How to Perform a Point-in-Time Recovery of the SYSTEM Save Sets and VSS SYSTEM Save Sets from the Command-prompt

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 StorEdge 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. Enter the following command on the Sun StorEdge EBS client host to restore the system state or system database to a particular point in time:

   Sun StorEdge EBS_install_path`\bin\recover.exe -iY` [`-s` *Sun StorEdge 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 StorEdge EBS log file.

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 StorEdge 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 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 Sun StorEdge EBS 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 51 on page 406 describes components of the SYSTEM STATE save set that require special preparation before being recovered.

Table 51. Preparing to Recover the SYSTEM STATE Save Set (Part 1 of 2)

Component	Recover Preparation
Active Directory (if installed)	• When the computer is restarting, "Directory Services Restore Mode" must be specified. • 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 StorEdge EBS server. For information about these settings, see "How to Prepare to Recover the System State on a Domain Controller" on page 838.
Certificate Server (if installed)	• Reinstall the Certificate Server after reinstalling the operating system. • 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. • Copy the EFS keys. For information about EFS keys, see "Encrypting File System" on page 844.
Cluster Server (if installed)	Shut down the Cluster Service on any nodes in the cluster on which the service is running. Shut down the Cluster Service by performing *one* of the following: • Enter the **net stop clussvc** command at the command-prompt. • Use the Microsoft Computer Management program.
COM+ Database	Set the *TEMP* environment variable to a valid temporary directory.
SYSVOL (if installed)	None
Internet Information Server	None

Table 51. Preparing to Recover the SYSTEM STATE Save Set (Part 2 of 2)

Component	Recover Preparation
Performance Counters	None
Registry	Do not attempt to recover a Windows NT 4.0 registry save set on a host running any other version of Windows.

Preparing to Recover the SYSTEM DB Save Set

Table 52 on page 407 describes components of the SYSTEM DB save set that require special preparation prior to recovery.

Table 52. 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	The Removable Storage service must be running. For information about Removable Storage database backup and recovery, see "Chapter 22: Support for Microsoft Removable Storage" on page 697.
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.



Important: The Sun StorEdge EBS Save Set Recovery feature does not support recovery of mount points. To recover mount points and their data, use the following special procedures. These procedures do not apply if Automated System Recovery (ASR) is used. For information about support for ASR, see "Chapter 24: Support for Microsoft Automated System Recovery" on page 731.

Recovering a Mount Point and Its Data

To recover a Windows mount point and its data:

1. Start the Sun StorEdge EBS User program.
2. Recover just the mount point, without its data. To recover a mount point on drive *D:* for example:
   a. Click the Recover button 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 unmark them before starting the recovery. A mount point may appear as a directory rather than a mount point (that is, it may have a folder icon instead of a volume icon).

f. Click Start to run the recovery.

3. Recover the mounted volume's data:

a. Click the Recover button 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.

For more information about performing data recoveries, see "Recovering by File Selection" on page 368.

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).



Important: Recover nested mount points and their data in the order shown in the following example.

The following procedure explains how to recover three nested mount points and their data in this configuration:

```
D:\
    mountpoint_name1
        data1
        mountpoint_name2
            data2
            mountpoint_name3
                data3
```

To recover this mount point tree:

1. Start the Sun StorEdge EBS User program.
2. Recover the top-level mount point*:*
   a. Click the Recover button 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.
      Notes:
      – When you mark a mount point for recovery, all files and directories beneath it are marked by default. Unselect 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 may 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, so the following procedures are not required.

How to Recover a DHCP Database

To recover the DHCP database:

1. Using the Sun StorEdge EBS User program, recover the *%SystemRoot% \System32\dhcp* directory.
2. Using Microsoft DHCP administrative tools, restore the DHCP database. For detailed instructions about Microsoft DHCP administrative tools, refer to the Microsoft documentation.

How to Recover a WINS database:

To recover a WINS database:

1. Use the Sun StorEdge EBS User program to recover the backup configured in the WINS backup procedure. For more information, see "How to Back Up a WINS Database" on page 85.
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 a 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.

Restoring CIS on Windows 2000 or above

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.



Important: 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 may 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.

How to Troubleshoot Problems with a CIS Restore

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.

Chapter 14: Reporting

This chapter covers the following topics:

- "Gathering Report Data" on page 413
- "Data Retention and Expiration Policies" on page 415
- "Report Categories" on page 418
- "Configuring Reports" on page 421
- "Viewing Reports" on page 426
- "Preconfigured Reports" on page 436
- "Customizing and Saving Reports" on page 462
- "Sharing Reports" on page 464
- "Exporting Reports" on page 465
- "Command-Line Reporting" on page 466
- "Printing Reports" on page 468

Gathering Report Data

To facilitate trend analysis, capacity planning, and problem detection, NetWorker software automatically collects data on a continual basis from the Sun StorEdge EBS enterprise. Sun StorEdge 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 415.

NetWorker software then integrates and processes this data to produce a number of reports on backup status, backup statistics, events, hosts, users, and devices. For detailed information about the various types of reports, see "Report Categories" on page 418.

For information on Retrospect® events, refer to the Dantz® Retrospect documentation.

It is possible to choose:

- Whether to collect data for the entire enterprise or for only specific Sun StorEdge EBS or Retrospect servers.
- Which types of reports to create.
- User preferences for report data, such as font, size, and whether or not to use bold. This can be useful in I18N environments.
- Which columns to display when viewing reports in a table format, and the order in which to display them.
- Whether to save customized reports for repeated use.
- How long to retain the collected data.

  Note: Only Sun StorEdge EBS administrators can modify these time periods.

- Whether to share reports with other users by giving them access to the reports.
- Whether to hide shared reports of other users when listing reports.
- Whether to run reports from the command-prompt.

  This enables the use of scheduling software so reports can be viewed at specified times.

How to 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 StorEdge 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 StorEdge EBS or Retrospect server to the enterprise. See "Add a NetWorker or Retrospect Host" on page 519.

To enable the gathering of report data:

1. From the Console window, click Enterprise.
2. Select the Sun StorEdge EBS or Retrospect server for which the collection of report information is to be enabled.
3. Right-click the Sun StorEdge EBS managed application, then select Set Features.

   The Set Features dialog box appears.
4. Select Gather Reporting Data, then click OK.

How to Disable the Gathering of Report Data

To disable the gathering of report data:

1. From the Console window, click Enterprise.
2. Select the Sun StorEdge EBS or Retrospect server for which the collection of report information is to be disabled.
3. Right-click the Sun StorEdge EBS managed application, then select Set Features.

   The Set Features dialog box appears.
4. Clear the Gather Reporting Data checkbox, then click OK.

Data Retention and Expiration Policies

NetWorker software provides separate expiration polices for retaining different types of data to meet the needs of the environment (as described in Table 53 on page 415). Only a Sun StorEdge EBS administrator can modify the policies.

Table 53. Data Retention Policies (Part 1 of 2)

Retention Policy	Determines how long NetWorker software stores the following:	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

Table 53. Data Retention Policies (Part 2 of 2)

Retention Policy	Determines how long NetWorker software stores the following:	Default
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 StorEdge EBS tasks (except License Manager tasks) performed by specified users (but only when the Sun StorEdge 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 "How to Set Expiration Policies for Data Retention" on page 417.

Note: Reports not mentioned in Table 53 on page 415 have no retention policies.

Flexibility of Retention Policies

These different policies give administrators the flexibility to retain certain types of information for less time than others.



Example: 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 StorEdge 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 StorEdge 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.

How to Set Expiration Policies for Data Retention

Note: Only an 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, enter 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.



Important: 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. For information about estimating the size of the Console database, refer to the *Sun StorEdge EBS Installation Guide.*

Report Categories

Table 54 on page 418 describes the various report categories included in NetWorker 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 54. Report Categories (Part 1 of 2)

Category of Report	Purpose
Sun StorEdge 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. For detailed information about the reports in this category, see "Sun StorEdge EBS Backup Statistics Reports" on page 436.
Sun StorEdge EBS Backup Status	Provide status information about savegroup completion and save set backups. For detailed information about the reports in this category, see "Sun StorEdge EBS Backup Status Reports" on page 444.
Events	Provide summary and detailed information about Sun StorEdge EBS events. For detailed information about the reports in this category, see "Event Reports" on page 452.
Hosts	Provide a listing of Sun StorEdge EBS servers in the Enterprise, including information about event and reporting features. For detailed information about the reports in this category, see "Host Reports" on page 457.

Table 54. Report Categories (Part 2 of 2)

Category of Report	Purpose
Users	Provide lists of defined Sun StorEdge EBS Console users, logout and login reports, audit reports, and users with restricted views. For detailed information about the reports in this category, see "User Reports" on page 458.
Devices	Provides information about the way devices are being used. For detailed information about the reports in this category, see "Device Reports" on page 460.

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 StorEdge EBS Backup Statistics, Sun StorEdge EBS or Retrospect Backup Status, Events, Hosts, and Users report categories. For a detailed description, see "Basic Reports" on page 420. Drill-down reports are available within the Sun StorEdge EBS Backup Statistics, Sun StorEdge EBS Backup Status, Events, and Devices report categories. For a detailed description, see "Drill-Down Reports" on page 420.

For information on:

- Retrospect events and error messages, refer to the Dantz Retrospect documentation.
- Managing Retrospect servers, see "Managing Sun StorEdge EBS and Retrospect Servers in the Enterprise" on page 518.

Different icons represent the different types of reports. Basic reports are represented by a red icon (), or a red icon with a person in the lower-left corner () if it is a shared report. Drill-down reports are distinguished from other reports with a downward pointing arrow (), or a downward pointing arrow and person in the lower-left corner () if it is a shared report.

For information on sharing reports, see "Sharing Reports" on page 464.

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 provides the specific focus of the needed information.

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 StorEdge 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 NetWorker software are known as *canned reports*, 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 following:

- The owner's reports (private and shared)
- The owner's reports, plus all shared custom reports

For more information, see "Customizing and Saving Reports" on page 462 and "Sharing Reports" on page 464.

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 StorEdge EBS Backup Statistics Server Summary report includes the following 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. For more information about this limitation, or for details on how to set this range, see "Date and Time Formats" on page 423.

The scope of a report can be limited by filtering out one or more parameter options, for example, as follows:

- 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 the Add button ( ).

How to Configure Reports

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 the Remove button ( ).

   To remove all of the parameters from the Selected box, click the Remove All button ( ).

   Removed parameters appear in the Available boxes.
4. To move:
   - A single parameter back to the Selected box, select it from the Available box and click the Add button ( ).
   - All unselected parameters back to the Selected box, click the Add All button ( ).

5. (Optional) If the report type includes a date-and-time-range parameter, specify the beginning and end date and time in the To and From text boxes.

   Click the down arrow to display a calendar and select a date.

   For more information about date and time formats, see "Date and Time Formats" on page 423.

6. 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 AM or PM 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. For information about exceptions to how these settings are treated, see "Locale Anomalies" on page 425.

Input Formats

Date and time input formats in the Sun StorEdge EBS software vary. Some acceptable input formats for a collection of common locales are shown in Table 55 on page 424.

Table 55. Date and Time Input Formats for Common Locales

Language	Date Formats	Time Formats
US English	• EEEE, MMMM D, YYYY (Monday, March 8, 2006) • MMMM D, YYYY (March 8, 2006) • MMM D, YYYY (Mar 8, 2006) • M/D/YY (3/8/06)	• h:mm:ss a z (11:27:30 PM PST) • h:mm:ss a (11:27:30 PM) • h:mm a (11:27 AM)
UK English	• DD MMMM YYYY 08 March 2006) • DD-MMM-YYYY (08-Mar-2006) • DD/MM/YY (08/03/06)	• 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 2006) • D MMMM YYYY (8 mars 2006) • D MMM YYYY (8 mar. 2006) • DD/MM/YY (08/03/06)	• 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 2006) • D. MMMM YYYY (7. März 2006) • DD.MM.YYYY (08.03.2006DD) • MM.YY (08.03.06)	• HH:mm:ss z (23:27:30 PST) • HH:mm:ss (23:27:30) • HH:mm (23:27)
Japanese	• YYYY/MM/DD (2004/03/06) • YY/MM/DD (04/03/06)	• HH:mm:ss z (23:27:30 PST) • HH:mm:ss (23:27:30) • HH:mm (23:27)
Simplified Chinese	• YYYY-M-D (2006-3-8) • YY-M-D (06-03-8)	• HH:mm:ss (23:27:30)

In this table, note the following:

- 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 timezone. If the "z" is present, then the output time will contain a timezone, otherwise, it will not.

Relative times can also be entered in the “From” and “To” fields. A valid relative time consists of an integer followed by a string. For example, “2 months” or “1 month”. The integer specifies how many relative time units, while the string specifies the unit of time. 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 four months. A report specifying "from 9 months to one month" includes data from nine months ago up to one month ago.

Note: For Drive Utilization reports, the time range cannot exceed eight days. That is, the date entered in the To field cannot exceed eight days from the date entered in the From field. If entering a relative time in the To field, the value cannot exceed 8 days.

Locale Anomalies

The locale settings in an operating system define factors that are specific to the local language or geography. While Microsoft Windows operating systems allow the user to set different default and current locales, the Java Runtime Environment (JRE) looks only at the default locale, and ignores any current locale settings. Consequently, NetWorker software selects the default locale settings for displaying dates and times.



Important: Understand the following in regard to Microsoft Windows operating systems:

- Changing the *default* locale puts the system into a state where much of the basic operating system service information (for example, basic messages, dates, and number formats) appears in localized form.
- If the default locale settings are changed, the system must then be rebooted. The default locale is used after rebooting.

- Changing the *current* locale does not cause the basic operating system service information to display in localized form.
- The JRE looks only at the *default* locale settings.

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. For information about printing and exporting, see "Printing Reports" on page 468 and "Exporting Reports" on page 465.

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.

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 56 on page 431.

  Note: When viewing a Drive Utilization report as a chart, it automatically displays as a Gantt chart, and 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 displayed 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 the following information:
  - Save Set Name
  - Start Time
  - End Time
  - Client Name
  - Throughput (B/Sec)

  In Drive View, the tool tip includes the following 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, the following options are available:

- Orientation (portrait or landscape)
- Table or chart format
- Size (zoom level)

Note: In Document mode, for any chart type that displays *X*-*Y* axes, two graphs are displayed. If the top graph contains excessive *Y*-axis data, the data displayed in both graphs could be truncated.

In Document mode, the columns of a tabular report cannot be sorted, rearranged, or resized. In addition, you cannot choose which columns to display, and the order in which to display them. Likewise, the chart format cannot be modified while viewing a report. NetWorker software does not maintain any customized changes made while displaying a report in interactive mode (such as sorting or rearranging the columns in a table), except for charts (in Chart Type and Chart Selector). Instead, document mode displays the report in a standard table or chart format, as specified by the internal report definition within NetWorker software.

Unlike Interactive Mode, which offers a set of parameters for chart selection that limit the data that is displayed, a report in Document mode displays all of the data. As a result, report views in Document mode often consist of several screens. For this reason, the viewing choices in Document mode include the following 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 does the following:

- 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.

Chart Types

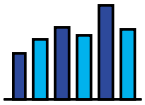
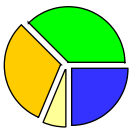
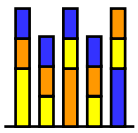
The following chart types are available in both interactive and document mode:

- 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. For more information, see "Types of Device Reports and Configuration" on page 460. The Gantt chart type is not an option for other reports.

Table 56 on page 431 shows a simplified version of chart format options.

Table 56. 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 StorEdge 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 check-boxes 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.

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.

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.

The following example shows an appropriate use of the stacking bar chart type.



Example: Appropriate Usage of the Stacking Bar Chart

To appreciate the different ways in which a stacking bar chart may be used, consider the following reports:

- A Sun StorEdge 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 StorEdge 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 NetWork 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 the following information:

- Save set name
- Start time
- End time
- Client name
- Throughput value

In Drive view, the *X* axis displays the time, and the *Y* axis displays drive data. Hovering over the chart in Drive View displays a tool tip that provides the following information:

- Drive
- Start time
- End time
- Throughput value

Chart Axis Selection

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.

How to View Reports

Before displaying a report, select the scope of the report by specifying Configuration parameters. For information on configuring reports, see "How to Configure Reports" on page 422.

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.

Restricted Views of Reports

When a Sun StorEdge EBS Console user is added or reconfigured, the user's views of Sun StorEdge 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 StorEdge 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 StorEdge 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, NetWorker software includes a variety of reports that can be viewed and output in a variety of ways.

Sun StorEdge EBS Backup Statistics Reports

The different types of reports included within the Sun StorEdge EBS Backup Statistics report category provide backup statistics for each selected Sun StorEdge EBS server within the enterprise.

Sun StorEdge EBS Backup Statistics reports may include the following information:

- Amount of data backed up.
- Number of files backed up.
- Number of save sets backed up.

Types of Sun StorEdge EBS Backup Statistics Reports and Configuration

The Sun StorEdge EBS Backup Statistics report category includes basic and drill-down reports.

The Configuration tab allows you to limit the scope of the report that was selected.

The parameters available within the Sun StorEdge EBS Backup Statistics report category are described in Table 57 on page 436. The specific parameters available depend on which Sun StorEdge EBS Backup Statistics report is selected.

Table 57. Sun StorEdge EBS Backup Statistics Parameters (Part 1 of 2)

Parameter	Description	Options
Server Name	Selects managed hosts within the enterprise.	Selected server names
Group Name	Selects one or more groups.	Selected group names
Client Name	Selects one or more clients.	Selected client names
Save Set Name	Selects one or more save sets.	Selected save set names

Table 57. Sun StorEdge EBS Backup Statistics Parameters (Part 2 of 2)

Parameter	Description	Options
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.	From and To dates and times **Note**: The date/time format available depends on the language locale of the operating system.

The parameters available for each report type in the Sun StorEdge EBS Backup Statistics report category are listed in Table 58, "Sun StorEdge EBS Backup Statistics Basic Reports," on page 438 and Table 59, "Sun StorEdge EBS Backup Statistics Drill-Down Reports," on page 442.

Save Set Data Retention Policy and Configuration

Settings for the save set retention policy impact the data that is available to the Sun StorEdge EBS Backup Statistics reports.



Example: Impact of Save Set Data Retention Policy

If a save set retention policy of six months is specified, NetWorker software cannot query the database for a time range that extends back more than six months. The report cannot display data that has expired, because that data has been removed from the database. Thus, even if a save time parameter of one year is specified, the report can display only six months of data if the limit of the save set retention policy is six months.

Backup Statistics Basic Reports

Within the Sun StorEdge EBS Backup Statistics report category, choose any of the basic reports described in Table 58 on page 438. 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. For information on selecting and removing parameters, see "How to Configure Reports" on page 422.

Note: If no time range for Save Time is specified, all available data for that report is displayed.

Table 58. Sun StorEdge EBS Backup Statistics Basic Reports (Part 1 of 4)

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 StorEdge 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

Table 58. Sun StorEdge EBS Backup Statistics Basic Reports (Part 2 of 4)

Report Name	Purpose	Configuration Parameters
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
Backup Type Summary	Indicates how many of the backup types listed in Table 57 on page 436 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

Table 58. Sun StorEdge EBS Backup Statistics Basic Reports (Part 3 of 4)

Report Name	Purpose	Configuration Parameters
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
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

Table 58. Sun StorEdge EBS Backup Statistics Basic Reports (Part 4 of 4)

Report Name	Purpose	Configuration Parameters
Save Set Details	Displays details about each save set, including backup duration. Only the Save Set Details and Save Set Details by Client reports do the following: • 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 do the following: • 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 StorEdge EBS Backup Statistics basic reports, connected in a predetermined sequence. For a list of available Backup Statistics drill-down reports, see Table 59 on page 442.

For general information about drill-down reports, see "Drill-Down Reports" on page 420.

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. For information on selecting and removing parameters, see "How to Configure Reports" on page 422.

Table 59. Sun StorEdge EBS Backup Statistics Drill-Down Reports (Part 1 of 3)

Report Name	Purpose	Sequence	Configuration Parameters
Backup Summary Over Time	Focuses on the different datazone layers. You can: • Compare statistics of the servers, groups, and clients in the enterprise. • Identify servers, groups, or clients that differ significantly from their peers.	1. Server Summary 2. Group Summary 3. Client Summary 4. Save Set Summary 5. Level Summary	• Server Name • Save Time • Backup Type • Backup Level
Monthly Group Statement	Combines a focus on the time period with a focus on savegroups. You can: • Compare statistics of servers and groups during the specified month. • Identify trends over time for groups and servers.	1. Monthly Summary 2. Group Summary by Server 3. Client Summary	• Server Name • Group Name • Client Name • Save Time • Save Set Name • Backup Type • Backup Level
Monthly Client Statement	Combines a focus on the time period with a focus on clients. You can: • Identify trends over time for clients. • Identify clients, where statistical areas differ significantly from their peers.	1. Monthly Summary 2. Group Summary by Client and Server 3. Daily Summary 4. Save Set Details	• Server Name • Group Name • Client Name • Save Time • Save Set Name • Backup Type • Backup Level

Table 59. Sun StorEdge EBS Backup Statistics Drill-Down Reports (Part 2 of 3)

Report Name	Purpose	Sequence	Configuration Parameters
Daily Save Set Statement	Combines a focus on the time period with a focus on save sets. Provides an additional layer of detail, since it includes the Save Set Details by Client report.	1. Monthly Summary 2. Daily Summary 3. Group Summary by Server 4. Save Set Details by Client	• Server Name • Group Name • Client Name • Save Time • Save Set Name • Backup Type • Backup Level
Save Set Details by Server over Time	Combines a focus on the time period with a focus on different datazone layers, starting with the server layer and moving down to save set details.	1. Server Summary 2. Group Summary 3. Monthly Summary 4. Daily Summary 5. Save Set Details by Client	• Server Name • Save Time • Backup Type • Backup Level
Level Summary by Server over Time	Combines a focus on backup levels with a focus on different datazone layers. Makes it possible to identify trends in the different levels of backups in the enterprise.	1. Level Summary 2. Server Summary 3. Group Summary 4. Monthly Summary	• Server Name • Group Name • Client Name • Save Time • Save Set Name • Backup Type • Backup Level

Table 59. Sun StorEdge EBS Backup Statistics Drill-Down Reports (Part 3 of 3)

Report Name	Purpose	Sequence	Configuration Parameters
Backup Type Summary by Server over Time	Combines a focus on backup types with a focus on the different datazone layers. Makes it possible to identify trends in backing up the different file types.	1. Backup Type Summary 2. Server Summary 3. Group Summary 4. Monthly Summary	• Server Name • Group Name • Client Name • Save Time • Save Set Name • Backup Type • Backup Level

Sun StorEdge EBS Backup Status Reports

The Sun StorEdge EBS Backup Status reports consolidate information about the success of scheduled group backups. As with the Sun StorEdge 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 StorEdge 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 StorEdge 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 run time is counted as 4 hours.

Sun StorEdge EBS Backup Status reports can include the following information:

- Total group runs
- Totals of successful, failed, and interrupted group runs
- Success ratio
- Backup duration
- Backup level

- Backup type
- Save type

Backup Type and Save Type Information

Backup type is one of the configuration parameters for both Sun StorEdge EBS Backup Statistics and Sun StorEdge 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 StorEdge 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 NetWorker software to identify the backup type and include it in the reports. The Sun StorEdge 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 Sun StorEdge EBS Module for Oracle 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—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 StorEdge EBS Backup Status Reports and Configuration

The Sun StorEdge 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 StorEdge EBS Backup Status Report category are described in the Table 60 on page 446.

Table 60. Sun StorEdge EBS Backup Status Parameters

Parameter	Description	Options
Server Name	Selects one or more Sun StorEdge 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 • INDEX • FILES • DB2 • INFORMIX • MSEXCHANGE • MSSQL • NOTES • NMEO • SAP • SYBASE
Level	Selects one or more backup levels.	• Full • Incremental • Skip • Level 1-9 (not a complete list of options)
Status	Selects status.	• Successful • Failed • Interrupted

The parameters available for each report type are listed in Table 61, "Sun StorEdge EBS Backup Status Basic Reports," on page 448 and Table 62, "Sun StorEdge EBS Backup Status Drill-Down Reports," on page 451.

Completion Data Retention and Sun StorEdge EBS Backup Status

The settings for the completion data policy impact the data that is available to the Sun StorEdge 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 StorEdge EBS Backup Status report category, choose any of the basic reports listed in Table 61 on page 448. 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. For information on selecting and removing parameters, see "How to Configure Reports" on page 422.

Note: If no time range is specified for Group Start Time, all available data for that report is displayed.

Table 61. Sun StorEdge EBS Backup Status Basic Reports (Part 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 StorEdge 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 • Group Start Time • Status
Monthly Summary	Shows status on a month-to-month basis.	• Server Name • Group Name • Status • Group Start Time
Daily Summary	Shows status on a day-to-day basis.	• Server Name • Group Name • Group Start Time • Status

Table 61. Sun StorEdge EBS Backup Status Basic Reports (Part 2 of 3)

Report Name	Purpose	Configuration Parameters
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 61. Sun StorEdge EBS Backup Status Basic Reports (Part 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 StorEdge EBS Backup Status basic reports, connected in a predetermined sequence. For a list of available Backup Status drill-down reports, see Table 59 on page 442.

For general information about drill-down reports, see "Drill-Down Reports" on page 420.

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. For information on selecting and removing parameters, see "How to Configure Reports" on page 422.

Table 62. Sun StorEdge EBS Backup Status Drill-Down Reports (Part 1 of 2)

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

Table 62. Sun StorEdge EBS Backup Status Drill-Down Reports (Part 2 of 2)

Report Name	Purpose	Sequence	Configuration Parameters
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

Event Reports

These reports provide summary information about current events on Sun StorEdge EBS, Retrospect, 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 StorEdge EBS Console displays the current events of the Sun StorEdge EBS and Retrospect servers, the Event reports provide additional advantageous features. The reports allow you to organize, export, and print the event data.

For information on Retrospect events, refer to the Dantz Retrospect documentation.

Event reports can include the following 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 63 on page 453. 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 63. Event Parameters

Configuration Parameter	Description	Options
Server Name	Selects one or more managed hosts.	Selected server names
Server Type	Selects only Sun StorEdge EBS servers, or all server types in the enterprise.	• Console • Sun StorEdge EBS
Priority	Selects only selected priority events. Priority represents the relative severity of the event. For descriptions of the priorities, see Table 71 on page 472.	• Warning • Waiting • Notice • Info • Emergency • Critical • Alert
Category	Selects only selected 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.	From and To dates and times

Event Basic Reports

Within the Events report category, select any of the basic reports listed in Table 64 on page 454. 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. For information about selecting and removing parameters, see "How to Configure Reports" on page 422.

Table 64. Event Basic Reports (Part 1 of 2)

Report Name	Purpose	Configuration Parameters
Server Summary	For all or specified Sun StorEdge 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. For a list of event priorities, see Table 71 on page 472.	• 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

Table 64. Event Basic Reports (Part 2 of 2)

Report Name	Purpose	Configuration Parameters
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 StorEdge EBS Console annotation. For information about annotations, see "Working with Annotations" on page 473. **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

Refer to these resources for more information about the contents of Event reports:

- For information on Sun StorEdge EBS Console notes, see "Working with Notes" on page 472.
- For descriptions of NetWorker software error messages and troubleshooting procedures, refer to the *Error Message Guide*.
- For information on troubleshooting NetWorker software issues, see "Chapter 26: Troubleshooting" on page 755.
- For information on Retrospect events and error messages, refer to the Dantz Retrospect documentation.
- For information on administering Retrospect through Console, see "Managing Sun StorEdge EBS and Retrospect Servers in the Enterprise" on page 518.

Event Drill-Down Reports

The drill-down reports consist of multiple Event basic reports, connected in a predetermined sequence. For a list of available Event drill-down reports, see Table 65 on page 456.

For general information about drill-down reports, see "Drill-Down Reports" on page 420.

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. For information on selecting and removing parameters, see "How to Configure Reports" on page 422.

Table 65. Event Drill-Down Reports

Report Name	Purpose	Sequence	Configuration Parameters
Events by Server	Provides a summary report on the currently active events, 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

These reports provide an overview of the servers in the Console server control zone.

Types of Host Reports and Configuration

The Hosts report category includes only basic reports, no drill-down reports. There are two basic reports, as described in Table 66 on page 457.

Table 66. Host Reports

Report Name	Purpose	Configuration Parameters	Default
Host List	Provides an overview of servers in the enterprise, including: • Whether it is a Sun StorEdge 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

For a description of the Enterprise and its folders, see "Enterprise" on page 515.

User Reports

These reports provide information on Sun StorEdge EBS Console user activity. For information about Sun StorEdge EBS Console users and creating user accounts, see "Managing Console Server Users" on page 491.

Types of User Reports and Configuration

The Users report category includes only basic reports, no drill-down reports. Select any of the basic reports listed in Table 67 on page 458.

Note: The Full Name and Description information appears in the User reports only if it was included when the user was created.

Table 67. User Reports (Part 1 of 3)

Report Name	Purpose	Configuration Parameters
User List	Lists the user IDs for all users that have been given access to the Sun StorEdge EBS Console. It includes the following 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 StorEdge EBS Console users, starting with the most recent login. Includes the following information for each user: • Last login time • Last logout time • Username • Full name (if available) • Description (if available) Also includes a listing of all logins available in the records.	Last Login Time (range)

Table 67. User Reports (Part 2 of 3)

Report Name	Purpose	Configuration Parameters
Users by Last	Lists the last logout time for all Sun StorEdge EBS Console users, starting with the most recent logout. Includes the following information for each user: • Last logout time • Last login time • Username • Full name (if available) • Description (if available) Also includes a listing of all logouts available in the records.	Last Logout Time (range)
Users Currently Logged In	Lists the usernames and login / logout times of all Sun StorEdge EBS Console users that are currently logged in. It includes the following information for each user: • Username • Full name (if available) • Description (if available) • Last login time • Last logout time	None

Table 67. User Reports (Part 3 of 3)

Report Name	Purpose	Configuration Parameters
User Audit	Allows the administrator to audit (track) a user's Sun StorEdge EBS server activities that were performed by using the Sun StorEdge 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. It includes the following information for each user: • User name • Hostname	• User Name • Hostname

Device Reports

These reports provide information about the way devices are being used. They show scheduled and ad-hoc backup activity on one or more selected devices over time. You can identify periods of heavy activity or inactivity. Device reports aid Sun StorEdge 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 68 on page 462 describes the Drive Utilization report.

The Drive Utilization Report, which is a drill-down report, supports Sun StorEdge EBS servers running Sun StorEdge 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 the following 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 the following 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 68. 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 entering 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)

Customizing and Saving Reports

Canned reports can be changed and then saved under different names, allowing users to 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 the following 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 StorEdge 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 StorEdge EBS Console. This also makes them accessible by the command-line reporting feature. For more information about command-line reporting, see "Command-Line Reporting" on page 466.

The following 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. For more information about chart axis selection, see "Chart Axis Selection" on page 434.

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 the following:

- 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 an administrator deletes from the system a user who owns reports, then the 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 Sun StorEdge EBS administrator may, however, enable it for sharing.

Only the original owner of a customized report or the Sun StorEdge EBS 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 administrator removes sharing, the report becomes private again to the original owner, the report's creator.

Any user viewing a sharable report may perform the following 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.

How to Share 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 69 on page 465 lists the formats into which reports can be exported.

Table 69. Report Export Formats

Format	Purpose
PostScript	For printing. Shows data totals.
PDF	For printing or viewing with a PDF viewer such as Adobe Acrobat. Shows data totals.
HTML	For viewing in a browser. Shows data totals.
CSV	For importing into other programs (such as spreadsheets) that accept the comma separated values (CSV) format. Does not show data totals. **Note**: For raw data only.

How to Export 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 on 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.
6. 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 the following 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.



Important: 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 character (-) followed by one or two letters and an argument, if applicable. For a complete description of the command and its options, refer to the *NetWorker Command Reference Guide* or the UNIX man pages.

System Performance

Each time the **gstclreport** command is run, it starts a separate JVM, which can use many system resources. The **gstclreport** command runs a database query and generates report output by using the results. Since this uses both CPU and memory resources on the host computer, it could affect performance of NetWorker software and of the host. Consequently, depending on the system used, it is probably not wise to run more than a few instances of the **gstclreport** command at the same time.

Security

The **gstclreport** command must contact the Console server in order to run a report. The command requires a valid username and password. A user either uses the **-P** option to enter 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 enter the **ps** command and see the commands that were used to start any currently running program.

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 1.4.2 or later to run the **gstclreport** command. The JRE must be installed before installing NetWorker software. Users provide the configuration and generation information for the reports and invoke the **gstclreport** command, which uses the JRE. For more information, refer to the *Sun StorEdge EBS Installation Guide.*

Printing Reports

All reports can be printed. This allows for the sharing of report data with users who are unable to view it online.

How to Print a Report

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.

Chapter 15: Monitoring

This chapter covers the following topics:

- "Overview of Monitoring" on page 469
- "Events" on page 469
- "Monitoring Window" on page 475

Overview of Monitoring

Sun StorEdge EBS Console provides the ability to view details of current Sun StorEdge EBS and Retrospect server activities, operations related to devices and libraries, and events that require user intervention. It makes administration of Sun StorEdge EBS and Retrospect servers more efficient by providing a centralized means of monitoring all activity throughout an enterprise. For details on adding Sun StorEdge EBS and Retrospect hosts to be monitored, see "Managing Sun StorEdge EBS and Retrospect Servers in the Enterprise" on page 518.

Events

An event signals that user intervention is required. For example, if a Sun StorEdge EBS server needs a new tape, the server alerts users to the situation by posting an event to the Console window.

NetWorker software generates an event based on various factors, including:

- The software encounters an error that requires user intervention to resolve.
- A NetWorker savegroup has failed.

15

- NetWorker software is unable to poll a host it is monitoring for events or for generating reports.
- A license or enabler code managed by the License Manager is about to expire.

Note: Some situations do not result in the generation of an event. For example, when a license managed by NetWorker Console (instead of by the License Manager) approaches its expiration date. In this situation, a message is recorded in the NetWorker logs, but an event is not generated until the expired license causes a backup to fail. Check the Administration window from time to time for important messages.

How to Enable or Disable the Option that Captures Events

The Capture Events option must be enabled for a given NetWorker or Retrospect server before NetWorker software can monitor that server for events. This option is selected by default when a host is added.

To disable the Capture Events option, or to re-enable it:

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 Set Features.

   The Set Features dialog box appears.
4. Do one of the following:
   - To enable the capturing of events, select Capture Events, then click OK.
   - To disable the capturing of events, clear the Capture Events checkbox, then click OK.

How to View Sun StorEdge EBS and Retrospect Events

To view events:

1. From the Console window, click Events.

   If any events exist, they are displayed in the Console window.

The Console window includes columns that provide specific types of information about each event. Table 70 on page 471 describes the various columns and the information they provide for Sun StorEdge EBS events.

For details on Retrospect events and error messages, refer to the Dantz Retrospect documentation..

Table 70. Events Columns

Column	Description
Priority	Represents the relative severity of the problem by displaying one of seven icons. Table 71 on page 472 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. For more information, see "Date and Time Formats" on page 423.
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. For more information, see "Working with Annotations" on page 473.
Note	Provides an editable field for making brief notes associated with an event. For more information, see "Working with Notes" on page 472.

15

Event Priorities

Each event is designated one of seven possible priorities. Table 71 on page 472 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” in between “Critical” and “Information”.

Table 71. Event Priorities

Icon	Priority	Description
	Alert	Severe error condition that demands immediate attention.
	Critical	Error condition detected by the Sun StorEdge EBS and Retrospect server that should be fixed by a qualified operator.
	Emergency	Condition exists that may cause NetWorker and Retrospect 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	Indication that the Sun StorEdge 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 the following:

- Name of the Sun StorEdge 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.

How to Add 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 enter the text of the note in the cell.
3. After entering the text, click outside the cell.

How to Edit 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.

How to Delete 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.

15

How to View an Annotation

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.

How to Add 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. Enter the text of the annotation.

   To clear the text just entered, click Reset.
4. Click OK.

How to Dismiss 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.



Important: Dismissing an event makes it disappear from the Console window for all Sun StorEdge 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.

Monitoring Window

Launch the Monitoring window from the Administration window taskbar to view details of current Sun StorEdge 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 StorEdge EBS server activities, it can be used to perform certain operations, such as starting, stopping, or restarting a group backup.

How to View the Monitoring Window

To view the Monitoring window:

1. From the Administration window, click Monitoring.
2. Click the tab corresponding to the type of information to view.

The Monitoring window includes several tabs that display specific types of information. Table 72 on page 475 describes the various tabs and the information they provide.

Table 72. Monitoring Window Tabs (Part 1 of 2)

Tab	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 tab 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.

15

Table 72. Monitoring Window Tabs (Part 2 of 2)

Tab	Information Provided
Devices	Lists devices, device status, storage nodes, libraries, volumes, pools, and related messages.
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 tab will display library data only for servers running NetWorker software release 7.3 or later. **Note:** When displaying the Show Details window from the Operations tab, 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 StorEdge 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.

In the default view, information on each tab appears in a table. Certain tabs enable you to choose either a Table view or an Icon view. You can select the type of view by right-clicking anywhere within the tab and selecting Table or Icon.

Viewing Backup Group Information

The Groups tab displays groups that are in the process of completing, or have completed, their backup. Use this tab to identify which groups backed up successfully and which groups failed.

The backup of a client group may fail for one of the following reasons:

- The Sun StorEdge EBS server failed.
- The Sun StorEdge EBS client failed.
- The network connection failed.

To find out more about a backup failure, check Group Backup details. For more information, see "How to View Group Backup Details" on page 479.

Groups Tab

The backup status is represented by an icon. Table 73 on page 477 lists and describes each of the icons.

Table 73. Groups Tab 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.

When items on the Groups tab are sorted by the Status column, they are sorted in alphabetical order based on the label of the icon.

Backup Group Operations

This section describes how to use the Monitoring window to start, stop, and restart group backups.

How to Start 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 StorEdge EBS server runs the backup at the level of the next scheduled backup, such as full, level [1 – 9], incremental, or consolidated.

To manually start a group backup:

1. From the Administration window, click Monitoring.
2. Click the Groups tab.
3. Right-click the group to start, then select Start.
4. Click Yes to confirm the start.

The Sun StorEdge 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.

How to Stop a Group

To stop a group backup:

1. From the Administration window, click Monitoring.
2. Click the Groups tab.
3. Right-click the group to stop, then select Stop.
4. Click Yes to confirm the stop.

How to Restart Groups

To restart a group backup:

1. From the Administration window, click Monitoring.
2. Click the Groups tab.
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.

How to View Group Backup Details

To view detailed information about a group backup:

1. From the Administration window, click Monitoring.
2. Click the Groups tab.
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 this button to view the messages.
5. Click OK to close the Group Backup Details dialog box.

Viewing Session Information

Use the Sessions tab to view the sessions currently running on a Sun StorEdge EBS server. You can change the view of this tab to display the following sessions:

- Save
- Recover
- Clone
- Browse
- All

The default setting for the Sessions tab is to display Save sessions. For instructions on viewing other session types, see "How to Change the Type of Sessions Displayed" on page 479.

How to Change the Type of Sessions Displayed

To change the type of sessions displayed on the Sessions tab:

1. From the Administration window, click Monitoring.
2. Click the Sessions tab.
3. From the Show drop-down list at the top of the tab, select the type of sessions to display. To display all sessions currently running on the Sun StorEdge EBS Server, regardless of type, select All Sessions.

The column headings displayed on this tab will differ depending on the type of sessions you chose to display.

15

Viewing Alerts

The Alerts tab displays alerts generated by a particular Sun StorEdge EBS server. It includes priority, category, time, and message information.

Alerts Tab

The alert priority is represented by an icon. Table 74 on page 480 lists and describes each of the icons.

Table 74. Alerts Tab Icons

Icon	Label	Description
	Alert	Severe error condition that demands immediate attention.
	Critical	Error condition detected by the Sun StorEdge EBS server that should be fixed by a qualified operator.
	Emergency	Condition exists that may cause NetWorker software to fail unless corrected immediately. This icon represents the *highest* priority.
	Information	Information about the current state of the server. This icon represents the *lowest* priority.
	Notification	Important information.
	Waiting	The Sun StorEdge 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 tab 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 tab allows you to monitor the status of all devices, including NDMP devices. If the Sun StorEdge EBS server uses shared and logical devices, the tab 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 tab. 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 tab at any time.

If the current server uses an AlphaStor library, then a Logical Name column is added to the Devices tab to accommodate logical devices.

Devices Tab

The device status is represented by an icon. Table 75 on page 481 lists and describes each of the icons.

Table 75. Devices Tab Icons (Part 1 of 2)

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.

15

Table 75. Devices Tab Icons (Part 2 of 2)

Icon	Label	Description
	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 tab are sorted by the Status column, they are sorted in alphabetical order based on the label of the icon.

Monitoring Device Operations

The Operations tab displays information about device operations. It includes the following 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. For instructions on how to deal with a user input notification, see "How to Supply User Input" on page 485.
- The origin, or source, of the operation.

  For example, the interface, **nsrjb**, or the Sun StorEdge EBS or Retrospect server.
- Time the operation started.
- Type of operation.
- Duration of the operation.
- Status messages from the operation.
- Any error messages.

  Only the last error message of the operation will appear in the Error Messages column. Move the mouse pointer over the cell to display the entire list of error messages.

Operations Tab

The operation status is represented by an icon. Table 76 on page 483 lists and describes each of the icons.

Table 76. Operations Tab Icons (Part 1 of 2)

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.

15

Table 76. Operations Tab Icons (Part 2 of 2)

Icon	Label	Description
	Running	The operation is running.
	Successful	The operation completed successfully.

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

How to View Operation Details

To view detailed information about an operation:

1. From the Administration window, click Monitoring.
2. Click the Operations tab.
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.

The operation details can be saved to a file by clicking Save in the Operation Details dialog box. When prompted, identify a name and location for the file.

How to Stop an Operation

Certain operations can be stopped from the Operations tab. To stop an operation:

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

How to 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

To view the most recent notification logs, click the Log tab in the Administration window. The Log tab of the Monitoring 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 StorEdge EBS server. The log files are located in the following directory: <*Sun StorEdge EBS_install_path*> \ *logs*.

Log Tab

The log priority is represented by an icon. Table 77 on page 485 lists and describes each of the icons.

Table 77. Log Tab Icons (Part 1 of 2)

Icon	Label	Description
	Alert	Severe error condition that demands immediate attention.
	Critical	Error condition detected by the Sun StorEdge EBS server that should be fixed by a qualified operator.

15

Table 77. Log Tab Icons (Part 2 of 2)

Icon	Label	Description
	Emergency	Condition exists that may cause NetWorker software to fail unless corrected immediately. This icon represents the *highest* priority.
	Information	Information about the current state of the server. This icon represents the *lowest* priority.
	Notification	Important information.
	Waiting	The Sun StorEdge 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 tab 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 tab displays the current status of all archive requests that are scheduled on the Sun StorEdge EBS server. Use this tab to identify which archive requests are running, completed, or failed, as well as when they were last run, and when they are scheduled to next run.

How to View 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 the Archive Requests tab.
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.

The archive request information can be saved to a file by clicking 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.

Archive Requests Tab

The archive status is represented by an icon. Table 78 on page 487 lists and describes each of the icons.

Table 78. Archive Requests Tab 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 tab are sorted by the Status column, they are sorted in alphabetical order based on the label of the icon.

15

How to Start 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. This is equivalent to selecting Start Now in the Status attribute of the Archive Request resource, described in "Scheduling Data Archives" on page 235.

To start an archive immediately from the Monitoring window:

1. From the Administration window, click Monitoring.
2. Click the Archive Request tab.
3. Right-click the appropriate archive request, then select Start.
4. Click Yes to confirm the start.

How to Stop an Archive in Progress

To stop an archive in progress:

1. From the Administration window, click Monitoring.
2. Click the Archive Requests tab.
3. Right-click the appropriate archive request, then select Stop.
4. Click Yes to confirm the stop.

How to Schedule an Archive to Start Automatically

You can also use the Monitoring window to schedule an archive to start automatically at a later time. This is equivalent to selecting Start Later in the Status attribute of the Archive Request resource and setting the Start Time, as described in "Scheduling Data Archives" on page 235.

To schedule an archive to start automatically at a later time:

1. From the Administration window, click Monitoring.
2. Click the Archive Requests tab.
3. Right-click the appropriate archive request and select Schedule Archive.
4. In the Schedule Archive Request dialog box, enter the time that the archive should start, by using the *hh*:*mm* format.
5. Click OK.

The Next Run column on the Archive Requests tab displays the entered time.

How to Disable a Scheduled Archive

If an archive request has a scheduled start time, you can disable the scheduled archiving. This procedure works whether the archive request was scheduled from within the Monitoring window, or by using the Start Later option in the Status attribute of the Archive Request resource, as described in "Scheduling Data Archives" on page 235.

To disable a scheduled archive:

1. From the Administration window, click Monitoring.
2. Click the Archive Requests tab.
3. Right-click the appropriate archive request, then select Disable Archive.
4. Click Yes to confirm the disable.

15

Chapter 16: Console Server Management

This chapter covers the following topics:

- "Managing Console Server Users" on page 491
- "Moving the Console Server Component" on page 498
- "Managing Product Licenses" on page 500
- "Setting System Options" on page 504
- "Setting Environment Variables" on page 511
- "Using the Console Configuration Wizard" on page 514
- "Displaying International Fonts in Non-US Locale Environments" on page 514

Managing Console Server Users

Sun StorEdge EBS administrators can control access to the Console server by creating user accounts and passwords. This prevents unauthorized users from gaining access to the Sun StorEdge EBS servers through the Console window. Sun StorEdge EBS administrators have special access rights.



Important: With the exception of user management (adding, editing, and deleting users, modifying system options, or restricting user views) and data retention policies (audit reporting, modifying the settings), all Console server user-account privileges complement those that the NetWorker server controls.

The Administrator User

The Sun StorEdge EBS administrator has the default password “administrator” after Sun StorEdge EBS software has been installed.



Important: To keep the Console server secure from unauthorized access, change this password at the first opportunity. See "How to Edit User Properties" on page 494.

By default, the Sun StorEdge EBS administrator has a full view of all hosts defined on the Console server. Only a standard user’s view of hosts can be restricted.

The Sun StorEdge EBS administrator cannot be deleted.

In addition to standard user access to the Console window, the Sun StorEdge EBS administrator can:

- Add, edit, and delete other users.
- Change or reset the administrator and other user passwords.
- Restrict user viewing and management of various servers.
- Modify system options. See "How to Set a System Option" on page 504.
- Modify report data retention policies. See "Data Retention and Expiration Policies" on page 415.
- Use the Console Configuration Wizard.

How to Add a User

Before a user can log in to the Console server, the Sun StorEdge EBS administrator must add the username and password. Only the Sun StorEdge EBS administrator may add users.

To add a user:

1. Log in as 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 the following characters: : < > /
   - Do not use the following ASCII characters: <=32
   - Do not begin a username with an underscore (_) character.
5. (Optional) Enter the full name of the user and a user description.
6. Enter the password.
7. Reenter the password in Confirm Password.
8. Click OK to give the user full privileges to all hosts, or go to "How to Set User Permissions to Manage Sun StorEdge EBS Servers" on page 493 to assign specific permissions.
   - To undo all changes, click Reset.
   - To dismiss the dialog box with no changes applied, click Cancel.

How to Set User Permissions to Manage Sun StorEdge EBS Servers



Important: Sun StorEdge EBS users can be restricted from interacting with Sun StorEdge EBS servers in either of the following ways:

- Set the User Authentication for Sun StorEdge EBS system option (for information about setting this option, see "Setting System Options" on page 504).
- Set the user permissions.

To set permissions for a user to view and manage Sun StorEdge EBS servers:

1. Log in as administrator.
2. From the Console window, click Setup.
3. Right-click a username, then select Permissions.

   The Edit User dialog box appears, with the Permissions tab displayed.
4. To grant permission for the user to view and manage various hosts, use the arrow keys to select the allowed hosts.

5. Click OK.
   - To undo all changes, click Reset.
   - To dismiss the dialog box, with no changes applied, click Cancel.



Important: Users can manage data only on the servers they are allowed to view. Even when users have permission to view a given Sun StorEdge EBS server, however, the management operations they can perform on it are determined by their privileges on that server.

For more information about restricting user views of various Sun StorEdge EBS servers, see "Restricting User Views" on page 497. To see how restricting views affects reporting, see "Restricted Views of Reports" on page 435.

How to Delete a User

Once a user has been deleted, that user can no longer log in to the Console server. Only the Sun StorEdge EBS administrator can delete other users.

To delete a user:

1. Log in as administrator.
2. From the Console window, click Setup.
3. Right-click a username, then select Delete.
4. 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.

How to Edit 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 StorEdge EBS administrator, however, can edit properties of other users.

To edit user properties:

1. From the Console window, click Setup.
2. Right-click a username, then select Properties.

   The Edit User dialog box appears, with the General tab displayed.

3. Change the user information (User Name, Full Name, Description, Password) as necessary. Other user information is displayed, but is not editable, including last-login information.

   Note: You can also change the administrator password by using the Console Configuration wizard. For more information, see "Using the Console Configuration Wizard" on page 514.

4. Click OK.

   – To undo all changes, click Reset.
   – To dismiss the dialog box, with no changes applied, click Cancel.

Dealing with User Groups

User groups can be created, modified, or deleted from the Administration window. For more information about user groups, see "Sun StorEdge EBS User Groups" on page 533.

Resetting the Administrator Password

If the Sun StorEdge EBS administrator password is lost or forgotten, it can be reset to the default by using the GST_RESET_PW environment variable.



Important: After the password has been reset, set the environment variable to null on the Console server. If this is not done, the administrator password will reset to the default value each time the system is restarted.

How to Reset the Administrator Password to the Default on Windows Systems

To reset the Sun StorEdge EBS administrator password on a Microsoft Windows system:

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 get to the Environment Variables dialog box differs depending on which Windows operating system you are running. Consult the documentation that accompanies your operating system for details.

2. Reboot the system.

   When the GST Service starts, the password resets.

3. Log in to the Console server as administrator, with the default password of **administrator**.

   The password can now be set to a secure password once again. For details, see "Edit User Properties" on page 494.

4. Return to the Environment Variables dialog box and delete the GST_RESET_PW variable.

5. Reboot the system.

6. Log in to the Console server as administrator by using the secure password.

How to Reset the Administrator Password to the Default on Solaris

To reset the Sun StorEdge EBS administrator password on a Solaris system:

1. Set GST_RESET_PW to a non-null value:

   set GST_RESET_PW *<a non-null value>*

2. Stop and restart the **gst** service with a reboot, or enter:

   /etc/init.d/gst **stop**

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

   The password resets.

3. Log in to the Console server as administrator, with the password of **administrator**.

   The password can now be set to a secure password once again. For details, see "How to Edit User Properties" on page 494.

4. Set GST_RESET_PW back to *null*:

   set GST_RESET_PW

   (press [Enter], without assigning a value)

5. Stop and restart the **gst** service:

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

6. Log in as administrator by using the secure password.

Restricting User Views

The Sun StorEdge EBS administrator can configure user accounts so that a user can view and manage data only on designated Sun StorEdge EBS servers. Throughout the Console window, a user's view of the Enterprise will then be limited to certain Sun StorEdge EBS servers and the data related to them.

Users can manage data only on hosts they are allowed to view. Even when they have permission to view a given Sun StorEdge EBS or Retrospect server, the management operations they can perform on it are determined by their privileges on that host.

The effects of restricting user views for various functions are as follows:

- In the Events window: The user sees only events from allowed Sun StorEdge EBS servers.

  The Retrospect events will be displayed in the Events window if 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 StorEdge EBS and Retrospect servers appear in those folders.
- In the Libraries window: The user sees only the devices controlled by allowed Sun StorEdge EBS and Retrospect servers.
- In the Reports window: The user sees report data only from allowed Sun StorEdge EBS servers, groups, clients and Retrospect servers.

For details on managing Retrospect servers, see "Managing Sun StorEdge EBS and Retrospect Servers in the Enterprise" on page 518.

- 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 Sun StorEdge EBS administrator can edit user permissions, and see and edit administrator permissions.

Important: Since each user can view and manage different sets of Sun StorEdge 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 run at the same time–if the users’ permissions include different Sun StorEdge EBS servers. This applies to all report types, whether canned or customized, private or shared.

Note: If no data is available for a given server, that server will not appear in any lists or reports, regardless of the user permissions.

For instructions on how to restrict user views, see "How to Set User Permissions to Manage Sun StorEdge EBS Servers" on page 493.

The Non-administrator User

A non-administrator (standard) user can:

- View and edit his or her own properties, and view his or her own permissions. See "Edit User Properties" on page 494.
- View properties of other users.

  Note: Only the Sun StorEdge EBS administrator can edit user permissions, and view and edit administrator permissions.
- View and manage the Sun StorEdge EBS servers to which the administrator has given him or her permissions.
- Customize and save reports.
- View his or her own reports, as well as shared reports created by other users.

Moving the Console Server Component

The Console server component of Sun StorEdge 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. For details, see "Performing an Ad Hoc Backup of the Console Database" on page 79.
2. Set up the new computer with the same operating system as the one running the current software and that it is connected to the network.
3. On the new computer, install the Sun StorEdge EBS Client software and the Sun StorEdge EBS Console server component.

   This step includes meeting the prerequisites for installing Sun StorEdge EBS software. For example, the Sun StorEdge EBS client must already be installed. If using Legato License Manager, the License Manager software must either already be installed or be available on another host. For more information, consult the *Sun StorEdge 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 more information, refer to the *License Manager Installation and Administration Guide*.
4. On the remote Sun StorEdge EBS server, set up the new target computer as a Client resource. For information about creating a Client resource, see "Task 7: Create a Backup Client Resource" on page 61.
5. For the source Console server NetWorker Client resource, 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. See "Stopping and Starting the Console Server" on page 49.
7. Stop the Console server service on the target Console server.
8. To recover the Console database, enter the following at the command-prompt:

   recoverpsm [**-f**][**-d** *recover_directory*]
   -S gst_on_S*ource_Console_server* **-s** *Sun StorEdge EBS_server*
   -c *source_Console_server*

   For a description of **recoverpsm** options, see Table 79 on page 500. Also refer to the *NetWorker Command Reference Guide* or the UNIX man pages for a complete description of the command and its options.
9. If *recover_directory* is not the same as the Console database directory, copy the database file to the directory where the Console database resides.

10. Start the Console server service on the remote Sun StorEdge EBS server.
11. Start the Console server service on the target Console server.
12. If you are using the License Manager and the License Manager host has changed, obtain a Host Transfer Affidavit, Go to *http://softwaresupport.emc.com*.

    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 79. 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 is 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.
-s *NetWorker_server*	Specifies the name of the Sun StorEdge EBS server where the Console database backup resides.
-c *source_Console_server*	Specifies the name of the source Console server, where the existing Console database resides.

Managing Product Licenses

This section provides information about managing Sun StorEdge EBS licenses.

License Manager (LLM) can be used to provide a central control point for license management of all EMC Legato products, or you can manage Sun StorEdge EBS licenses from the Administration window of Sun StorEdge EBS Console.

Note: LLM is available only from Sun.

For information about the different software licensing schemes, refer to the *Sun StorEdge EBS Installation Guide*.

For information about the License Manager, refer to the *License Manager Installation and Administrator's Guide*.

Entering a License Enabler Code

Once you have obtained your enabler code, enter it into the Sun StorEdge EBS software. You can purchase additional enablers for added features, as well as support for additional clients or autochangers. If you are entering multiple enabler codes, enter the base enabler code for the Sun StorEdge EBS software before you enter the enabler codes for additional features.

How to Enter a License Enabler Code without LLM

To enter an enabler code if you are *not* using LLM:

1. From the Administration window, click Configuration.
2. Right-click Registrations, then select New.

   The Create Registration dialog box appears.
3. In the Enabler Code attribute, enter the enabler code.
4. (Optional) In the Comment attribute, enter a description of the license.
5. Click OK.

How to Enter a License Enabler Code with LLM

To enter a license 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 a License 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.



Important: You cannot delete the enabler code that enables the basic Sun StorEdge EBS software, called the *base enabler code*. A base enabler can only be upgraded or downgraded.

How to Delete a License Enabler Code without LLM

To delete a license enabler code if you are *not* using LLM:

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.

How to Delete a License Enabler Code with LLM

To delete a license enabler code if you are using LLM:

1. From the Console window, click Setup.
2. Click Licensing.
3. Right-click the license to delete, then select Delete.
4. Click Yes to confirm the deletion.

Entering an Authorization Code

Note: Before starting this procedure, register the product and receive the authorization code. For information about registering the product, refer to the *Sun StorEdge EBS Installation Guide*.

The authorization code must be entered in order to enable the software permanently.

How to Enter an Authorization Code without LLM

To enter an authorization code if you are *not* using LLM:

1. From the Administration window, click Configuration.
2. Click Registrations.

3. Right-click the license to be authorized, then select Properties.

   The Properties dialog box appears.

4. In the Auth Code attribute, enter the authorization code for the product (the authorization code assigned to the specified permanent enabler or update enabler code).

5. Click OK.

The license is now permanently enabled.

How to Enter an Authorization Code with LLM

To enter an authorization code if you are using LLM:

1. From the Console window, click Setup.
2. Click Licensing.
3. Right-click the license to be authorized, then select Properties.

   The Properties dialog box appears.

4. In the Auth Code attribute, enter the authorization code for the product (the authorization code assigned to the specified permanent enabler or update enabler code).

5. Click OK.

The license is now permanently enabled.

Changing the License Manager Server

The License Manager server that manages Sun StorEdge EBS Console licenses can be changed at any time. For information about using the License Manager, refer to the *License Manager Installation and Administrator's Guide*.

How to Change the License Manager Server

To change which host runs the License Manager:

1. From the Console window, click Setup.
2. Right-click Licensing, then select Change LLM Server.

   The Change LLM Server dialog box appears.

3. In the LLM Server attribute, enter the hostname of the appropriate server.
4. Click OK.

You can also change the License Manager by using the Console Configuration wizard. For more information, see "Using the Console Configuration Wizard" on page 514.

Setting System Options

Sun StorEdge EBS Console includes several options that affect performance. These options enable users to fine-tune the performance of the Console server. Only NetWorker administrators can set or change system options.

How to Set 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. Change them only if you are already well acquainted with the performance of the Console server. A mistake in setting system options or environment variables can seriously degrade performance.

Sun StorEdge EBS Console gets two configuration options from the environment (GST_MAXLOGSIZE and GST_MAXLOGVERS). Those environment variables can be changed by advanced users.

Table 80 on page 505 describes the available system options.

Table 80. Console System Options (Part 1 of 4)

To do the following:	Specify this system option:
Set the level of debugging information written to the *gstd.log*. (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*, choose a larger value. • To use less disk space, choose a smaller value.	Debug level **Note**: GST_DEBUG can be set as an environment variable. This allows users to troubleshoot Sun StorEdge EBS 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.

Table 80. Console System Options (Part 2 of 4)

To do the following:	Specify this system option:
Set the minimum number of seconds that must pass before Sun StorEdge EBS Console checks a Sun StorEdge EBS server for events. • To make Sun StorEdge EBS 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 events and reporting (secs)
Set the minimum number of seconds that must pass before Sun StorEdge EBS Console checks a Sun StorEdge EBS server for updates to the Administration windows. • To make the Administration windows more responsive to changes on the Sun StorEdge EBS server, decrease the value. • To decrease the load on the CPU, increase the value. **Range:** >= 2 **Default:** 10	Polling interval for Sun StorEdge EBS activities (secs)
Set the minimum number of hours that must elapse before the Sun StorEdge EBS Console checks a Sun StorEdge EBS server for library information. • To poll the Sun StorEdge EBS server for library information more frequently, decrease the number of hours. • To poll the Sun StorEdge EBS server for library information less frequently, increase the number of hours. **Range:** Greater than or equal to 1. **Default:** 12	Polling interval for Sun StorEdge EBS libraries (hours)

Table 80. Console System Options (Part 3 of 4)

To do the following:	Specify this system option:
Determine the number of programming threads that are used when polling Sun StorEdge EBS servers. • To make Sun StorEdge EBS Console more responsive, increase the value. • To decrease the load on the CPU, decrease the value. **Range**: 1 – 10 **Default**: 5	Polling thread factor
Enable auditing of user Sun StorEdge EBS Console requests. • When enabled, this function logs all Sun StorEdge EBS Console user requests to Sun StorEdge EBS servers. • When this function is disabled, Console user requests to Sun StorEdge EBS servers are not logged. **Range**: Not applicable **Default**: Disabled	Sun StorEdge EBS user auditing

Table 80. Console System Options (Part 4 of 4)

To do the following:	Specify this system option:
Set individual user authentication for functions a user can perform on a Sun StorEdge EBS server. **Note**: Requires restarting the Console server to take effect. This option restricts user access to a given Sun StorEdge 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 Sun StorEdge EBS Console users who are allowed to access a Sun StorEdge EBS server have the same privileges on that server. • When this function is enabled, access privileges on Console servers can vary among Sun StorEdge EBS Console users. **Range:** Not applicable **Default:** Enabled	User authentication for Sun StorEdge EBS
Enable the use of user datagram protocol (UDP) when Sun StorEdge EBS Console first checks whether a Sun StorEdge 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. For more information, see "Additional Firewall Support Considerations" on page 828. **Range:** Applicable **Default:** Enabled	RPC ping via UDP when connecting to a Sun StorEdge EBS server

Individual User Authentication

NetWorker administrators can restrict or grant Console user access to Sun StorEdge EBS servers based on the Console username if the User Authentication for Sun StorEdge EBS system option has been enabled (with a subsequent restart of the Console server). This system option is enabled by default.



Important: Requests to Sun StorEdge EBS servers through the Administration window always come from the Console server, regardless of any system option settings.

When the User Authentication for Sun StorEdge EBS system option is enabled and the Console server has been restarted, the requests to Sun StorEdge EBS servers appear–from the Sun StorEdge 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 and the Console server has been restarted, a Sun StorEdge EBS server allows requests only from users who belong to the Administrators group of the NetWorker server. The username of the **gstd** process owner must be included in the Sun StorEdge EBS Administrators group on Sun StorEdge EBS servers to which NetWorker Console users have access. Sun StorEdge EBS administrators must also ensure that the usernames of NetWorker Console users who have permission to view and manage those servers are included in the Sun StorEdge EBS Administrators group. For examples, see Table 81 on page 510. For detailed information about the Administrators group, see "Administrator Privileges" on page 531.

Table 81. Usernames for Individual User Authentication

On Console server:	Operating System:	Example of username to add to the Sun StorEdge EBS Administrators group[1]:
Root or system user on the Console server–the **gstd** process owner (mandatory) **Note**: This username *must* be specified, regardless of whether individual user authentication is used.	Microsoft Windows	system@winbox.petropep.com
	Solaris	root@unixbox.petropep.com
Sun StorEdge EBS Administrator	Microsoft Windows	administrator@winbox.petropep.com
	Solaris	administrator@unixbox.petropep.com
User	Microsoft Windows	user1@winbox.petropep.com
	Solaris	user1@unixbox.petropep.com

[1] Computer examples shown are Console servers.

To include the username of the **gstd** process owner in the Sun StorEdge EBS Administrators group, see "Managing Console Server Users" on page 491.

The changes that users can make to a specific Sun StorEdge EBS server depend on the privileges granted to the Sun StorEdge EBS user group or groups to which they belong.

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. For information about firewalls, see "Appendix B: Firewall Support" on page 811.

When the User Authentication for Sun StorEdge EBS system option is set, a separate network connection is created from the Console server to a Sun StorEdge 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 StorEdge EBS system option is *not* set, there is only one network connection from the Console server to a Sun StorEdge EBS server that is being managed through Sun StorEdge EBS.

Setting Environment Variables

Table 82 on page 511 describes the three environment variables available in this Sun StorEdge EBS release. Variables are only checked 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 82. Console Environment Variables (Part 1 of 2)

To:	Change this Environment Variable:
Set the maximum size of the file *gstd.log* 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. For more information about *gstd.log*, see "The gstd Log" on page 798. **Range:** >= 1 **Default:** 1 MB	GST_MAXLOGSIZE

Table 82. Console Environment Variables (Part 2 of 2)

To:	Change this Environment Variable:
Set the maximum number of historical versions of *gstd.log* that Sun StorEdge EBS 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. For more information about *gstd.log*, see "The gstd Log" on page 798. **Range**: 1–999 **Default**: 4	GST_MAXLOGVERS
Set the level of debugging information written to the *gstd.log*. Setting this value higher increases the number of operation and status messages that Sun StorEdge EBS 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.

How to Set Environment Variables on Solaris Systems

To set environment variables on Solaris systems:

1. Open the */etc/init.d/gst* file in a text editor. (This file is a Bourne shell script.)

   Note: By default, this is a read-only file, so check the file permissions and change them, if necessary, before editing it.

2. Add the following 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.

3. Save the changes to */etc/init.d/gst*.
4. Stop and restart the Console server. For detailed instructions, see "Stopping and Starting the Console Server" on page 49.

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.

How to Set Environment Variables on Windows 2000 Systems

To set environment variables in Windows 2000 (Server or Professional):

1. View the system variables.
   a. From the Start menu, select Settings>Control Panel>System.
   b. Click the Advanced tab.
   c. Click Environment Variables.
   d. In the Environment Variables dialog box, below System Variables, click New.
2. In the New System Variable dialog box, enter the variable name and value.
3. Click Set.
4. Click OK.
5. Reboot the system to enable the software to recognize the changes.

How to Set Environment Variables on Windows XP or 2003 Systems

To set environment variables in Windows XP or 2003:

1. View the system variables.
   a. From the Start menu, select Control Panel>System.
   b. Click the Advanced tab.
   c. Click Environment Variables.
   d. In the Environment Variables dialog box, below System Variables, click New.
2. In the New System Variable dialog box, enter the variable name and value.

3. Click OK to proceed with the operation and dismiss the New System Variable dialog box.
4. Click OK to dismiss the Environment Variables dialog box.
5. Reboot the system to enable the software to recognize the changes.

Using the Console Configuration Wizard

The following tasks can be performed from the Console menu options or from the Console Configuration wizard.

Note: Only NetWorker administrators can use the Console Configuration Wizard.

- Set the NetWorker Administrator password.
- Set the License Manager server.
- Specify the Sun StorEdge EBS server that will back up the Console database.
- Add Sun StorEdge EBS servers to the Enterprise.

To access the Console Configuration wizard if you are a NetWorker administrator:

1. From the Console window, click Setup.
2. From the Setup menu, select Configuration Wizard.

Displaying International Fonts in Non-US Locale Environments

The Console server can operate in English mode on localized operating systems and display double-byte characters from Sun StorEdge EBS servers. To use or view data from a localized, non-English Sun StorEdge EBS server, ensure that the appropriate font is available to the Console server. For more information about displaying international fonts on a Console server that is operating in English mode, refer to the Console server installation section in the appropriate *Sun StorEdge EBS Installation Guide.*

Chapter 17: Sun StorEdge EBS Server Management

This chapter covers the following topics:

- "Enterprise" on page 515
- "Configuring a Sun StorEdge EBS Server" on page 527
- "Improving Backup Performance" on page 528
- "Managing Server Access" on page 530
- "Sun StorEdge EBS User Groups" on page 533
- "Server Communication Issues within Microsoft Windows" on page 540
- "Notifications" on page 541
- "Indexes" on page 552
- "Message Log Management" on page 564

Enterprise

The Enterprise is a visual representation of the Sun StorEdge EBS Console control zone. NetWorker and Retrospect servers in the Enterprise can be monitored for events, and various reports can be generated on events, backups, and user activity. For detailed information about Sun StorEdge EBS Console reports, see "Chapter 14: Reporting" on page 413.

Components of the Enterprise

The Enterpise includes the following components:

- Hosts

  A host, also known as a *managed node*, is the Sun StorEdge EBS or Retrospect 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 the Sun StorEdge EBS Servers

Use the Enterprise to organize the Sun StorEdge EBS servers by some logical or functional criteria. Examples of organizational criteria include:

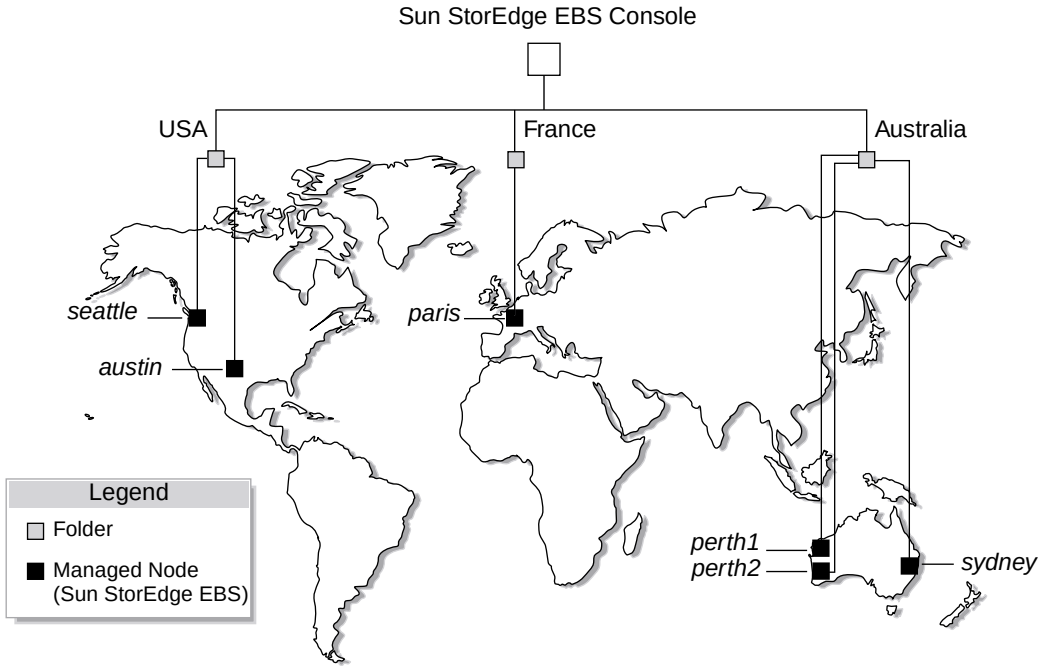
- By geography, thus putting all the hosts from the same city or country in the same folder. (See the example on page 517.)
- By function of the computers backed up by the Sun StorEdge EBS servers, such as having the servers that back up web servers in one folder, and the servers that back up mail servers 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: An Enterprise Arranged by Geographic Location

The example in Figure 36 on page 517 depicts an Enterprise arranged by geographic location. There are three folders, one for each country where the Sun StorEdge EBS servers are located: USA, France, and Australia. Each folder contains a number of hosts, that is Sun StorEdge 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 36. Sun StorEdge EBS Servers Worldwide



How to View the Enterprise

In the Console window, the organization of the Sun StorEdge 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 StorEdge EBS servers.

   The right pane displays the contents of the selected folder or host.

2. Select a view option as described in Table 83 on page 518.

Table 83. 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 Sun StorEdge EBS and Retrospect Servers in the Enterprise

Sun StorEdge EBS software allows management of Sun StorEdge EBS and Retrospect 7.5 servers within the Enterprise. You can add, delete, move, and copy Sun StorEdge EBS and Retrospect servers as needed. All of these functions can be performed through the Console window. When using Sun StorEdge EBS software to manage a large number of Sun StorEdge EBS or Retrospect servers, it might be more efficient to add or delete multiple hosts by using a single command from the command-prompt. For further information, see "Adding or Deleting Multiple Sun StorEdge EBS or Retrospect Servers by Using a Hostname File" on page 523.

The addition of Sun StorEdge EBS and Retrospect 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.

Important: This functionality is intended for existing NetWorker environments with Retrospect servers that require management in the Enterprise. It is not intended for Retrospect-only environments.

For details on:

- Monitoring Sun StorEdge EBS and Retrospect events, see "Chapter 15: Monitoring" on page 469.
- Retrospect events and error messages, refer to the Dantz Retrospect documentation.
- Gathering Sun StorEdge EBS and Retrospect report data, see "Chapter 14: Reporting" on page 413.

Requirements for Adding a Retrospect Server as a Host

The requirements for administering a Retrospect server using the Console window include:

- 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.
- The first 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.

How to Add a NetWorker or Retrospect Host

By adding a Sun StorEdge EBS or Retrospect server as a host, the Console window can display Sun StorEdge EBS or Retrospect events. This setup is required for generating Sun StorEdge EBS and Retrospect server activity-reports.

Sun StorEdge EBS Console supports hosts that use Sun StorEdge EBS server software release 6.1.4 or later, and Retrospect server software release 7.5.

Note: The host can either be a Retrospect or a Sun StorEdge EBS server, but not both.

To add a Sun StorEdge EBS or a Retrospect host 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.

   Hostnames and aliases cannot exceed 80 characters.
4. Select Sun StorEdge EBS or Retrospect, then click Next.
5. Follow the instructions on each subsequent page, then click Finish.

A host can also be added by using the Console Configuration Wizard. For more information, see "Using the Console Configuration Wizard" on page 514.

How to Delete a Host

You can delete multiple hosts or just a single host within a folder.

To delete a *single* host:

1. From the Console window, click Enterprise.
2. Right-click the host to delete, then select Delete.

   The Deleting Host dialog box appears.

   Note: If additional copies of the host exist in the Enterprise, there is an option ("Delete all existing copies of the host") to delete all instances of that same host in a single operation.
3. Click Yes to confirm deletion of the host.

   The host is deleted.

To delete *multiple* hosts, see "How to Delete a Folder" on page 522.

How to Copy a Host

Multiple copies of a host can be created for a single Sun StorEdge EBS or Retrospect 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.

How to Move 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.

Configuring Remote Access for a Retrospect Host

By configuring remote access, you can use NetWorker 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 the first port 3000 is available for Console and Retrospect to communicate. For details on how to configure the Microsoft Remote Desktop Connection or Microsoft Terminal software, refer to the appropriate Microsoft documentation.

Managing Folders in the Enterprise

Sun StorEdge 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.

How to Add 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].

How to Delete 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.

How to Copy 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.

   Note: A folder cannot be copied within the same Enterprise level.

   A copy of the folder appears in its new location.

How to Move 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.

How to Rename 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 Sun StorEdge EBS or Retrospect Servers by Using a Hostname File

For larger enterprises, it may be more convenient to add or delete multiple Sun StorEdge EBS or Retrospect 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 StorEdge EBS managed nodes with the features Capture Events and Gather Reporting Data enabled.

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. For information about this procedure, see "How to Move a Host" on page 521.

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 the following 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 the following example:



Example: Hostname File

```
#This is a hostname file for XYZ Corporation

apple
banana
grape
kiwi
mango
nectarine
pineapple
strawberry
tangerine
```

Using the gstmodconf Command

The **gstmodconf** command has the following syntax:

gstmodconf -i *file* **-f** *function* **-s** *server* **-k -p** *port* **-l** *login* **-P** *password*

For a complete description of the command and its options, refer to the *NetWorker Command Reference Guide* or the UNIX man pages.



Example: Adding Multiple Hosts with the gstmodconf Command

This example 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:

```
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

The next example shows the error that is generated if **gstmodconf** is used to add a host that already exists in the Enterprise:



Example: 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...
```

The next example 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: 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 StorEdge EBS Server

To use the Administration window to modify a Sun StorEdge EBS server, the server must meet the following requirements:

- Be included in the Enterprise.

  Note: Dantz Retrospect servers can also be included in the Enterprise.

  For information about adding hosts to the Enterprise, see "How to Add a NetWorker or Retrospect Host" on page 519.
- Not restrict users from viewing the Sun StorEdge EBS server.

The Administration window can be used to modify an existing Sun StorEdge EBS server by specifying attributes in the Properties dialog box of the Sun StorEdge EBS server. The configuration of these attributes, such as setting the parallelism level or designating administrator privileges, can affect backup performance and security.

How to Set Up the Server

Note: When setting up the Sun StorEdge EBS server, be sure to enter the Sun StorEdge EBS product serial number. The product serial number is located on the Enabler Certificate that was sent when the Sun StorEdge EBS product was ordered.

To set up the Sun StorEdge 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 StorEdge EBS Server

The Sun StorEdge EBS server can be licensed two ways:

- If you use the License Manager to license Sun StorEdge EBS servers, the Console window can be used. For instructions, see "Managing Product Licenses" on page 500. The License Manager is only available from Sun.
- If you do *not* use the License Manager to license Sun StorEdge EBS servers, the Administration window can be used. For instructions, see "How to License the Sun StorEdge EBS Server by Using the Administration Window" on page 528.

How to License the Sun StorEdge EBS Server by Using the Administration Window

To license a Sun StorEdge EBS server by using the Administration window:

1. From the Administration window, click Configuration.
2. Right-click Registrations, then select New.
3. Complete the necessary attributes, then click OK.

Improving Backup Performance

You can affect backup speed by changing the setting for multiplexing data on the Sun StorEdge EBS server. This allows the Sun StorEdge EBS server to write data from more than one save set to a single storage volume. Also, the server can distribute data from one save set across multiple volumes. You can control multiplexing through the Parallelism and Target Session attributes.

Parallelism

The *parallelism* feature keeps a steady stream of data supplied to the storage device, so that it can operate at higher speeds. Multiplexing also ensures that no particular client can monopolize the server. This means that other clients can back up data to the same server simultaneously.

Each edition of Sun StorEdge EBS software has a different default and maximum parallelism level. Each enabled storage node connected to the Sun StorEdge EBS server can increase the parallelism maximum. The maximum parallelism for any Sun StorEdge EBS server and storage node combination can vary. For more information, refer to the *Sun StorEdge EBS Release Notes*.

By optimizing the efficiency of the server, you can control how many clients can back up their data at one time. You may want to adjust the number of clients under the following circumstances:

- If the server is taking too long to back up all the clients, increase the parallelism setting.
- If the server is so busy backing up clients that it is useless for any other tasks, decrease the parallelism setting.

How to Set Parallelism

To set the Parallelism attribute:

1. From the Administration window, click Configuration.
2. Select the server name.
3. From the File menu, select Properties.

   The Properties dialog box appears.
4. Click the Setup tab and adjust the value in the Parallelism attribute.

   Note: To turn off parallelism so that the server backs up only one client at a time, set the parallelism value to 1.
5. Click OK.

Target Sessions

The *target sessions* feature enables you to set the number of backup sessions accepted by an active device. Each device connected to a Sun StorEdge EBS server or storage node (including each device in an autochanger or silo) can have a different target sessions value. For more information about adding devices, and about the Device dialog box, see "Add (Configure) Libraries" on page 121.

You can set the target sessions value to any number between 1 and 512. This value also determines which device will process the next backup session.

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

It is often more efficient for the Sun StorEdge EBS server to multiplex multiple save sets to the same device, rather than write each save set to a separate device. For this reason, the Sun StorEdge 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.

How to Set Target Sessions

To set the Target Sessions attribute:

1. From the Administration window, click Devices.
2. Click Devices in the left pane to display all the devices.
3. Right-click the device with the target session to be set, then select Properties.
4. Click the Configuration tab and adjust the value in the Target Sessions attribute.
5. Click OK.

Managing Server Access

Sun StorEdge EBS server access is managed by granting access privileges to Sun StorEdge EBS Console users.

Granting Access to Manage Sun StorEdge EBS Servers from the Console Server

If the Console server and the Sun StorEdge EBS server are installed on separate hosts, the Console administrator must be granted access to administer and monitor the target NetWorker server. To grant access, use the **nsraddadmin** command from the command-prompt on the Sun StorEdge EBS server. Grant access on each Sun StorEdge EBS server that is to be managed from the Console server. For information about granting access with the **nsraddadmin** command, refer to the Console server installation section of the appropriate *Sun StorEdge EBS Installation Guide*. For more information about the **nsraddadmin** command, refer to the *NetWorker Command Reference Guide* or the UNIX man pages.

Administrator Privileges

Use the Administration window to grant Sun StorEdge EBS administrator privileges to additional Console users. This is done by adding users to the Administrators group for the Sun StorEdge EBS server. For information, see "Sun StorEdge EBS User Groups" on page 533.

When granting privileges to users, consider the following:

- To change security settings, the user must be a member of the Sun StorEdge EBS Administrators group.
- To perform a scheduled backup of SYSTEM and VSS SYSTEM save sets, the user must be a member of the Sun StorEdge 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 StorEdge EBS privilege.

Sun StorEdge 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 StorEdge 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 StorEdge 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 usernamed *jdoe* on the host named *jupiter*, enter the following 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 StorEdge EBS from any host (this is equivalent to entering 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: Using the Sun StorEdge EBS Administrators Group

This example grants Sun StorEdge 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 StorEdge EBS Administrative Privileges

The Sun StorEdge 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 *user=user_name, domain=domain_name* or *group=group_name, domainsid=domain_id.* For example:

- **`user=joe,domain=NT-ENG`**
- **`group=Administrators,domainsid=S-1-5-32-323121-123`**

Sun StorEdge EBS User Groups

Sun StorEdge EBS includes an access control feature that is configured through the User Group resource. This feature allows Sun StorEdge 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 StorEdge 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 StorEdge EBS server. One enabler per datazone is required. This feature is available only on servers with Sun StorEdge EBS release 7.0 and later.

Preconfigured User Groups

Sun StorEdge EBS comes with two preconfigured user groups:

- Administrators
- Users

Administrators

Members of the Administrators group have permission to perform all Sun StorEdge 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. For permission requirements for a specific utility, refer to the *NetWorker Command Reference Guide* or the UNIX man pages.

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 StorEdge 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 84 on page 535 lists these preconfigured privileges and their associated permissions.

Table 84. User Group Privileges (Part 1 of 3)

Privilege	Permissions
Change Security Settings	Allows group members to edit the following 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 the Change Security Settings privilege can grant access to restricted functions to any user. User groups with this privilege must also have these privileges: Configure Sun StorEdge EBS, Operate Sun StorEdge EBS, Monitor Sun StorEdge 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 StorEdge EBS, Monitor Sun StorEdge EBS, Operate Devices and Jukeboxes, Backup Local Data, and Recover Local Data.

Table 84. User Group Privileges (Part 2 of 3)

Privilege	Permissions
Configure Sun StorEdge EBS	Allows group members to configure resources associated with the Sun StorEdge 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 via the Change Security Settings privilege. User groups with this privilege must also have these privileges: Operate Sun StorEdge EBS, Monitor Sun StorEdge EBS, Operate Devices and Jukeboxes, Backup Local Data, and Recover Local Data.
Operate Sun StorEdge EBS	Allows group members to perform Sun StorEdge 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 StorEdge EBS, Operate Devices and Jukeboxes, Backup Local Data, and Recover Local Data.
Monitor Sun StorEdge EBS or Retrospect	Allows group members to do the following: • Monitor Sun StorEdge EBS or Retrospect operations, including device status, save group status, and messages. • View media databases information. • View Sun StorEdge 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.
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.

Table 84. User Group Privileges (Part 3 of 3)

Privilege	Permissions
Recover Local Data	Allows group members to recover data from the Sun StorEdge 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. **Notes:** • 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 saveset or NDMP recovers.
Backup Local Data	Allows group members to manually back up data from their local client to the Sun StorEdge 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. **Notes:** • This privilege does not override file-based permissions. Users can only back up 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 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 Groups

This section provides procedures for creating, editing, copying, and deleting a User Group resource.

How to Create 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 531.

5. Select the privileges to assign to the user group.
6. Click OK.

How to Edit 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.

How to Copy 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.

How to Delete 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

The following sections address various client / server communication issues that occur when running Sun StorEdge 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 StorEdge 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 StorEdge EBS server name to either the local *hosts* file (located in *%SystemDrive% \ winnt \ 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 StorEdge 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 StorEdge 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 StorEdge EBS server becomes invalid. If the Sun StorEdge EBS server is a DHCP client, a static TCP / IP address for the server must be preserved. Sun StorEdge EBS clients can still use dynamic addressing.

Backup and Recover Server Service

In Microsoft Windows operating systems, the Sun StorEdge 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 StorEdge EBS environment. You can configure the events to be reported, as well as how the Sun StorEdge EBS server reports them to you. Specific programs can be executed when an event occurs, including third-party programs. By default, the Sun StorEdge EBS server sends notifications to log files located in the <*Sun StorEdge EBS_install_dir*> \ *logs* directory.

Note: Retrospect server events can be displayed in the Console window.

Preconfigured Notifications

Sun StorEdge EBS is preconfigured to provide most of the notifications required to monitor Sun StorEdge EBS and Retrospect events. Table 85 on page 541 lists these preconfigured notifications and the associated actions performed by the Sun StorEdge EBS server.

Table 85. Preconfigured Notifications (Part 1 of 4)

Notification	Action Performed by the Sun StorEdge EBS Server
Bootstrap	Sends the bootstrap information to the default printer of the Sun StorEdge EBS server.
Cleaning cartridge expired	**Windows:** Reports expired cleaning cartridge to the *media.log* file. **UNIX:** Sends the following email to root: *replace the expired cleaning cartridge*.

Table 85. Preconfigured Notifications (Part 2 of 4)

Notification	Action Performed by the Sun StorEdge EBS Server
Cleaning cartridge required	**Windows:** Reports need for cleaning cartridge to the *media.log* file. **UNIX:** Sends the following email to root: *insert a cleaning cartridge.*
Client Install	**Windows:** Reports hostname and NetWorker client software version information to the *media.log* file. **UNIX:** Sends the following email to root: *host* ***host_name*** *installed* ***product_version.*** Where ***host_name*** is the name of the NetWorker host, and ***product_version*** is the NetWorker client software release and build number.
Device cleaned	**Windows:** Reports that device has been cleaned to the *media.log* file. **UNIX:** Sends the following email to root: *device cleaning is complete*.
Device cleaning required	**Windows:** Reports that device requires cleaning to the *media.log* file. **UNIX:** Sends the following email to root: *clean the device*.
Device disabled	**Windows:** Reports that a device has been automatically disabled to the *media.log* file. **UNIX:** Sends the following 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.

Table 85. Preconfigured Notifications (Part 3 of 4)

Notification	Action Performed by the Sun StorEdge EBS Server
Filesystem full - waiting for adv_file space	**Windows:** Reports that the advanced file volume is full to the *media.log* file. **UNIX:** Sends the following 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 the following email to root: *check the size of the client file index because it will soon exceed the space available.*
Log default	**Windows:** Sends data about Sun StorEdge 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, do one of the following: • 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*.

Table 85. Preconfigured Notifications (Part 4 of 4)

Notification	Action Performed by the Sun StorEdge EBS Server
Registration	**Windows:** Sends messages about the registration status of your Sun StorEdge EBS products to the *messages* log file. **UNIX:** Sends the following email to root: *check the registration status.*
Savegroup completion	**Windows:** Reports the success of client save sets to the *savegrp.log* file. **UNIX:** Sends the following email to root: *degree of success in completing scheduled backups, cloning, and archive operations.*
SNMP notification request	Sends event notifications to a network management console (this notification only appears if the StorEdge EBS SNMP module has been purchased and enabled). For details on SNMP notifications, see "Chapter 21: SNMP Module" on page 691.
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.
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.

Customizing Notifications

Notifications require three elements:

- Events
- Actions
- Priorities

These elements are described in the following sections.

Events

Table 86 on page 545 lists the events that trigger a notification.

Table 86. Events (Part 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	The cleaning cartridge has expired and needs replacing.
Cleaning cartridge required	Mount the cleaning cartridge.
Client	NetWorker client software has been installed on a host.
Deleted media	A media device has been deleted.
Device cleaned	The device has been cleaned.
Device cleaning required	The device requires cleaning.
Device disabled	A device has been automatically disabled.
Index attention	The index needs attention.
License expiration	A license has expired.
Media	A media-related event occurred. For example, a volume may require mounting.
Media attention	The 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.

Table 86. Events (Part 2 of 2)

Event	Description
Media request	The media needs operator attention to mount backup volumes.
Registration	Product registration needs attention.
Savegroup	A backup group has completed backing up.
Server	Other server events (for example, restarting the Sun StorEdge EBS server).
Storage node	A storage node has been installed.
Write completion	A write operation is complete.

Note: For details on Retrospect events displayed in the Console window, refer to the Dantz Retrospect documentation.

Actions

Table 87 on page 546 lists the actions that the server takes after an event notification occurs.

Table 87. Actions (Part 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. For more information about using **nsrlpr**, see "How the Sun StorEdge EBS Server Uses nsrlpr to Print Notifications" on page 547.
logger	(UNIX only) Uses the UNIX **syslog** facility (*/usr/bin/logger*) to log information or send messages.
lp	(UNIX only) Prints the notification.

Table 87. Actions (Part 2 of 2)

Action	Description
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 Administrators group, or users with the Change Security Settings privilege, can change the Action attribute of an existing notification. For more information about privileges, see "Managing Console Server Users" on page 491.

Note: For details on Retrospect Actions, refer to the Dantz Retrospect documentation.

How the Sun StorEdge EBS Server Uses nsrlpr to Print Notifications

Sun StorEdge EBS Server, on Windows systems only, looks at the printer named in each of the following 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 using the **-P** option of the **nsrlpr** command.

How to Designate 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, enter the following:

   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 the following 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.

How to 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, enter the following:

   nsrlpr -P %printer

How to Test the nsrlpr Program

To test the **nsrlpr** program, enter the following 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.

If you receive the following error message:

```
Error: print server did not accept request. Job aborted.
```

then the printer name is incorrect. Correct the name and try to print again. Once you can print from the command-prompt, change the Action attribute to the print command, as follows:

nsrlpr -P *printer_name*



Important: Print jobs sent by the Sun StorEdge EBS Backup and Recover Server service run in the Local System context. Under certain conditions, it may not have access to network print queues. For more information, refer to Microsoft Knowledge Base articles 132679 and 143138 on the Microsoft website.

Priorities

Each Sun StorEdge 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 StorEdge 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 88 on page 549 lists the priorities upon which notifications are based.

Table 88. 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 StorEdge 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 StorEdge EBS to fail unless corrected immediately.

Note: Event priorities are sorted alphabetically, rather than by severity.

Notification Event Logging

Sun StorEdge EBS keeps two general notification log files. By default, these files are located in the following locations:

- *<Sun StorEdge EBS_install_dir> \ logs \ messages.log* — The data in the *messages.log* file is generated by **nsrlog**, which is part of the Sun StorEdge EBS event notification mechanism. The **nsrlog** program is triggered by a notification, and it prints the file to *messages.log*.
- *<Sun StorEdge EBS_install_dir> \ logs \ daemon.log* — The **nsrd**, **nsrexecd**, and their subordinate processes redirect their output to the *daemon.log* file.

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 allows you to look selectively at the messages that interest you by filtering messages based on the categories listed in Table 89 on page 550.

Table 89. Event Viewer Messages

Event Viewer Category	Displayed Information
Source	Events from Sun StorEdge EBS software always designate Sun StorEdge EBS as the source.
Category	Mapped from Sun StorEdge EBS notification event type (savegroup, server, registration, and so on).
Severity	Mapped from Sun StorEdge 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 StorEdge EBS software always designate the numeral 1 for the ID.

How to Create a Custom Notification

Sun StorEdge EBS also provides preconfigured notifications. For a complete list of preconfigured notifications, see "Preconfigured Notifications" on page 541.

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. For command options, see Table 90 on page 555.
7. Click OK.

How to Edit 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.

How to Copy 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.

How to Delete a Custom Notification

Note: You cannot delete any of the preconfigured notifications.

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.

   When prompted, click Yes to confirm the deletion.

Indexes

The Sun StorEdge 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 the following:
  - 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. For information about managing the size of these indexes, see "Managing the Size of the Online Indexes" on page 560.

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 StorEdge 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 the following 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.



Important: If the file system that contains the indexes gets full, the Sun StorEdge 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 StorEdge EBS uses browse and retention policies to manage the lifecycle of data, and to automatically control the size of the client file index. For more information on policies, see "Data Retention and Expiration Policies" on page 415.

Automated Index Activities

The Sun StorEdge EBS server performs the following online index activities:

- Inserts entries in the client file index for each file saved during a backup. For each new backup, the Sun StorEdge 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. For more information, see "How to Remove the Oldest Save Set Cycles" on page 561.

Checking Online Indexes

Each time the Sun StorEdge 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 StorEdge EBS Server" on page 395.

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

How to View 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 90 on page 555 identifies the index information displayed for each client.

Table 90. Indexes Window Information

Column	Description
Client Name	Names of the Sun StorEdge 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.

How to View Information about Client Save Sets

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 91 on page 555 identifies the information in the Save Sets dialog box for each save set.

Table 91. Index Save Sets Dialog Box Information (Part 1 of 2)

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.

Table 91. Index Save Sets Dialog Box Information (Part 2 of 2)

Column	Description
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 to 9)

For information about reducing the size of the client file indexes by using the Remove Oldest Cycle button, see "Reducing Client File Index Size" on page 560

How to Query 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.

   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. For a complete description of the **mminfo** command and its options, refer to the *NetWorker Command Reference Guide* or the UNIX man pages.

How to Cross-Check Client File Indexes

By cross-checking, you can verify the consistency between the client file index and the media database. If the Sun StorEdge 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 StorEdge 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 click Cross Check Index.

   The Sun StorEdge EBS server displays the following prompt:

   `Cross-checking may take considerable time. Would you like to cross-check` *client_name*?
4. Click Yes to continue.

   The Sun StorEdge EBS server displays a status box until the cross-checking is complete.

How to Refresh Index Information

Occasionally refresh the information displayed 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.

Location of the Client File Index

During the initial client setup, Sun StorEdge EBS normally designates a default location for the client file index on the Sun StorEdge EBS server. This default location is as follows:

- For UNIX: */nsr/index/client_name*
- For Windows: <*Sun StorEdge 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. The following sections address these needs.

How to Designate 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. For instructions, see "Setting Up a Scheduled Backup" on page 55
6. Click OK.

How to Change the Client File Index Location for an Existing Client



Important: 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. For more information, see "How to Move a Client File Index" on page 559.

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*, enter the following:

   ```
   nsrck -L6 jupiter
   ```

   The resulting output will be similar to the following:

   ```
   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.

How to Move a Client File Index

You might want to 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 StorEdge EBS server.
2. Copy the client file index from its current location to the new location. For example:
   - For Windows, use Windows Explorer to copy <*Sun StorEdge EBS_install_dir*> \ *index* \ *jupiter* to *C:* \ *Sun StorEdge EBS* \ *indexes* \ *jupiter*.

- For UNIX, enter the following at a command-prompt:

  `cp -rp /nsr/index/jupiter /Sun_StorEdge_EBS/indexes`

3. Update the Index Path attribute of the Client resource to point to the new location of the index. For instructions, see "How to Change the Client File Index Location for an Existing Client" on page 558.

Managing the Size of the Online Indexes

Over time, the size of the online indexes on the Sun StorEdge EBS server can become prohibitively large. Reduce the size of these indexes by using the solutions suggested in the following sections.

Reducing Client File Index Size

You can reduce the size of the client file indexes on the Sun StorEdge EBS server by using one or more of the following methods:

- Remove save sets that comprise the oldest backup cycle from the client file index. For more information, see "How to Remove the Oldest Save Set Cycles" on page 561.
- Delete volume-based entries from the client file index. For more information, see "How to Delete Volume-Based Online Index Entries" on page 563.
- Adjust the Browse Policy and Retention Policy attributes of clients backing up to the Sun StorEdge 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. For more information, see "Modifying the Browse and Retention Policy on a Save Set" on page 183.

If the size of the client file index for a client is still too large, consider moving the location of the index. For more information, see "How to Move a Client File Index" on page 559.

Reducing Media Database Size

Reduce the size of media database on the Sun StorEdge EBS server by using one or more of the following methods:

- Remove volumes that contain recyclable save sets from the Sun StorEdge EBS inventory. For more information, see "Removing Volume-Based Entries from the Online Indexes" on page 562.

  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. For more information, see "Change the Mode of a Volume" on page 334.

  When a volume is recycled, the Sun StorEdge EBS server performs the following 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. For more information, see "Modifying the Browse and Retention Policy on a Save Set" on page 183.

- Compress the media database. For more information, see "How to Compress the Media Database" on page 564.

How to Remove 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 StorEdge 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 the following:

```
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, enter the following:

```
mminfo -v -c client_name
```

For more information, refer to the **nsrmm** and **mminfo** entries in the *NetWorker Command Reference Guide* or the UNIX man pages.

When save sets are marked recoverable, users may no longer browse and recover these files by using the Sun StorEdge EBS 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. For information on this procedure, see "Recovering by Save Set Selection" on page 379.

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 StorEdge 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. For more information, see "Recovering Online Indexes for a Sun StorEdge EBS Server" on page 395.

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 StorEdge 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.

How to Delete 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 the following options to determine how volume entries will be removed:
   - File and Media Index Entries.

     For details about this option, see "Removing Client File Index Entries" on page 562.
   - File Index Entries Only.

     For details about this option, see "Removing Client File Index and Media Database Entries" on page 562.
5. Click OK.

The Sun StorEdge 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.

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 the following at the command-prompt:

nsrmm -d -S *ssid*

To remove all information related to a particular volume, enter the following at the command-prompt:

nsrmm -d *volume_name*

How to Compress 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 file as follows:
   - On Windows: <*Sun StorEdge EBS_install_dir*> \ *mm* \ *.cmprssd*
   - On UNIX: */nsr/mm/.cmprssd*
2. Enter the following at the command-prompt:

   nsrim

Message Log Management

The Sun StorEdge EBS server stores the messages generated by the Sun StorEdge EBS server daemons in a message log file in the following location:

- On Windows: <*Sun StorEdge 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.

How to Set the UNIX Startup Script to Trim Log Files

To modify the way that Sun StorEdge EBS services manage the Sun StorEdge EBS log files, set the following environment variables in the Sun StorEdge EBS startup script (*/etc/init.d/networker*) before the line:

```
(echo 'starting Sun StorEdge EBS daemons:') > /dev/console
```

Table 92 on page 565 lists the location of Sun StorEdge EBS startup scripts:

Table 92. Sun StorEdge EBS Startup Script Location

Operating System	Startup Script Location
Solaris, DYNIX/ptx, and IRIX	*/etc/init.d/networker*
HP-UX	*/sbin/init.d/networker*
HP Tru64	*/sbin/init.d/NSRstartstop*
Linux	*/etc/init.d/networker*
AIX	*/etc/rc.nsr*

Enter the environment variables by using the format:

ENV_VAR_NAME=*value*
export *ENV_VAR_NAME*

The following environment variables can be used to trim log files at startup:

- To change the maximum size of log files, change the NSR_MAXLOGSIZE 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 StorEdge EBS server for the environment variables to take effect.

Every time the Sun StorEdge 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 1,024 KB, the following occurs:

1. The *daemon.log* file is renamed *daemon.001* and a new empty *daemon.log* is created.
2. If the *daemon.log* file fills again, the names of each existing file shift so that the *daemon.001* file is renamed *daemon.002*, *daemon.log* is renamed *daemon.001*, and a new empty *daemon.log* file is created.
3. This process is repeated until the value in NSR_MAXLOGVERS is reached, at which point the highest numbered log is removed.



Important: The trimming mechanism functions only when you start **nsrd** on UNIX or the Sun StorEdge 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 StorEdge 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 StorEdge EBS daemons, and then use the Sun StorEdge EBS startup script to restart the daemons.
- On Windows, stop and restart the Sun StorEdge EBS Backup and Recover Server service.

For information about the **nsr_shutdown** command and its options, refer to the *NetWorker Command Reference Guide* or the UNIX man pages.

Chapter 18: Sun StorEdge EBS Client Management

This chapter covers the following topics:

- "Sun StorEdge EBS Client Overview" on page 567
- "Client Configuration" on page 568
- "Sun StorEdge EBS Authentication" on page 571
- "Multiple Clients from the Same Computer" on page 583
- "Adding or Changing the Sun StorEdge EBS Servers for a Client" on page 585
- "Client Priority" on page 586

Sun StorEdge EBS Client Overview

A Sun StorEdge EBS client is both a physical computer with Sun StorEdge EBS client software installed on it, and a Sun StorEdge EBS server *resource* that specifies a set of files and directories to be included in a scheduled backup. As such, a single Sun StorEdge EBS client computer can have several Client resources specified that can back up to the same or even different Sun StorEdge EBS servers.

The concept of a client computer with multiple Sun StorEdge EBS Client resources is further explained in "How to Define a Client and Save Set Combination" on page 584.

Sun StorEdge 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 StorEdge EBS server. For example, you can back up a Sun StorEdge EBS client on a Microsoft Windows computer to a Sun StorEdge EBS server on a Solaris computer.

Client Configuration

Before a client can be backed up by a Sun StorEdge EBS server, the client must have the appropriate Sun StorEdge EBS client software installed. For more information, refer to the *Sun StorEdge EBS Installation Guide*.

How to Create a Client

Sun StorEdge EBS client creation is part of the process of creating a scheduled backup. For instructions on how to create a client, see "Setting Up a Scheduled Backup" on page 55.

How to Edit a Client

Use this procedure to change Client resource attributes. You cannot change the name of a client with this procedure. To change a client name, see "How to Change a Client Name" on page 569.

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.

How to Copy a Client

Use this procedure to copy Client resource attributes.

To copy a client:

1. From the Administration interface 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.

How to Change a Client Name

The only way to change the name of a client is to re-create the client.

To change the name of a client:

1. From the Administration interface 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 "How to Delete a Client" on page 569.
7. 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. For instructions on how to create a client, see "Setting Up a Scheduled Backup" on page 55.

How to Delete a Client

When a client is deleted, the Sun StorEdge 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 interface 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 StorEdge EBS server, the previously backed-up data from the client is still accessible and recoverable directly from the volume containing the data by using the *scanner* command.

How to Recover 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. For instructions on how to create a client, see "Task 7: Create a Backup Client Resource" on page 61.

The Sun StorEdge 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 StorEdge 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: For information about permissions for directed recoveries and the Disable Directed Recover attribute in the NSRLA resource, see "Enable Directed Recoveries" on page 390.

How to Edit the NSRLA Database

To edit the NSRLA database:

1. Log in as root or as Windows administrator on the Sun StorEdge EBS client.
2. Enter the following at the command-prompt:

   ```
   nsradmin -p nsrexec
   ```

   The **nsradmin** prompt appears.
3. To determine the current settings for attributes in the NSRLA resource, enter the following at the **nsradmin** prompt:

   ```
   print type:NSRLA
   ```
4. To change the value of attributes in the NSRLA resource, enter the following at the **nsradmin** prompt:

   `update` *attribute*`:`*value*`;`

   For example, to update the Disable Directed Recover attribute, enter the following:

   ```
   update disable directed recover:Yes
   ```

5. Enter 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.

Sun StorEdge EBS Authentication

This section describes how to configure authentication between Sun StorEdge EBS hosts. Special considerations for authentication are also included in this section.

Overview

Two types of authentication are supported for Sun StorEdge EBS hosts:

- **nsrauth**
- **oldauth**

Strong Authentication (nsrauth)

The **nsrauth** authentication mechanism is strong authentication that is based on the Secure Sockets Layer (SSL) protocol, which is provided by the OpenSSL library. Sun StorEdge EBS hosts and Sun StorEdge EBS user permissions are authenticated by using **nsrauth**. The **nsrauth** authentication mechanism is available for hosts running Sun StorEdge EBS software release 7.3 or later.

Each Sun StorEdge 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 StorEdge EBS permissions

Authentication for Backwards Compatibility (oldauth)

For compatibility with earlier Sun StorEdge EBS releases, **oldauth** authentication is supported and enabled by default. If two hosts cannot authenticate by using strong authentication (**nsrauth**), they will attempt to authenticate by using **oldauth**. However, one can specify the minimum authentication strength that is allowed for any host relationship. For more information, see " Specifying the Minimum Authentication Strength between Hosts" on page 573.

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 the following resources on each Sun StorEdge EBS host whose authentication information will be updated:

- NSRLA
- NSR peer information
- NSR system port ranges

To grant authentication update privileges to the Console server:

1. Log in as root or as Windows administrator on the Sun StorEdge EBS host whose authentication information will be updated.
2. Enter the following at the command-prompt:

   `nsradmin -p  nsrexec`

   The **nsradmin** prompt appears.
3. For each resource:
   - 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. Enter 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 StorEdge EBS release 7.3 or later server be allowed to authenticate with legacy Sun StorEdge EBS clients by using **oldauth.**
- That all other Sun StorEdge EBS release 7.3 or later clients must use only **nsrauth** Strong authentication.

Note: The Console server user must have permission to update authentication resources. For more information, see "Access Privileges for Authentication Configuration" on page 572.

To specify the authentication strength between Sun StorEdge EBS hosts:

1. From the Administration window, click Configuration.
2. Open the Local Hosts folder.
3. From the Local Hosts list, right-click the host with the authentication relationships to be configured and select Configure Local Agent.
4. Select the Advanced tab.
5. In the Auth Methods attribute, at the top of the list enter the minimum allowable authentication strength for the Sun StorEdge EBS hosts that will connect to this host. Enter values in the following 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 StorEdge EBS client name.
- *10.102.0.0/255.255.0.0*
  A subnet IP address representing all Sun StorEdge EBS clients on the subnet. Alternatively, you could enter this value as *10.102.0.0/16*.
- *0.0.0.0/0*
  This value represents all clients in the domain.

where *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** / **oldauth**, which means that all clients can authenticate using either **nsrauth** or **oldauth**. For more information, see "Example: Enforcing Strong Authentication on Selected Hosts" on page 575.

6. Click OK.
7. Restart the Sun StorEdge EBS services for the host with the authentication relationships to be configured. For information about restarting Sun StorEdge EBS services, see "Stopping and Starting a Sun StorEdge EBS Server, Client, or Storage Node" on page 51.



Example: Enforcing Strong Authentication on Selected Hosts

In this example, a Sun StorEdge 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 StorEdge EBS server and select Configure Local Agent.
4. Select the Advanced tab.
5. In the Auth Methods attribute, enter the two legacy hostnames, for example:

   `jupiter.company.com, oldauth`

   `pluto.company.com, 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, enter the following 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, oldauth
   jupiter.company.com, 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 StorEdge EBS services for the Sun StorEdge EBS server. For information about restarting Sun StorEdge EBS services, see "Stopping and Starting a Sun StorEdge EBS Server, Client, or Storage Node" on page 51.

Maintaining Sun StorEdge EBS Local Host Authentication Credentials

Each Sun StorEdge EBS host that uses **nsrauth** strong authentication has unique credentials that are used to identify itself to other Sun StorEdge EBS hosts during the **nsrauth** strong authentication process. These credentials are known as local host authentication credentials.

Additionally, each Sun StorEdge EBS host maintains a copy of the local host credentials that belong to every Sun StorEdge 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. The following 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. For more information, see "Access Privileges for Authentication Configuration" on page 572.



Example 1: Securing the Initial Peer Host Authentication Process

When two Sun StorEdge 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 small 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 StorEdge EBS server *jupiter* will authenticate with the Sun StorEdge 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.

   For information about exporting authentication credentials, see "How to Export Local Host Credentials" on page 579.
2. Create a peer resource for *pluto* on *jupiter*:
   a. Open the credentials file for *pluto*. Refer to this file in the following 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.

      In the following sample credential file, the values to enter are italicized:

      ```
      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*. For more information, see "How to Create a Local Host Peer Resource Certificate" on page 582.
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.
   d. Click OK.
5. Repeat step 2 to step 4 to create a peer resource for *jupiter* on *pluto*.



Example 2: Local Host Credentials Are Exported

In most cases, local host credentials are retained. For example, when the Sun StorEdge 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 StorEdge 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. For more information, see "How to Export Local Host Credentials" on page 579.
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. For more information, see "How to Import Local Host Credentials" on page 580.

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 3: Local Host Credentials Are Not Exported

This scenario is similar to "Example 2: Local Host Credentials Are Exported" on page 578, 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 StorEdge EBS host. For more information, see "How to Create New Local Host Certificate Keys" on page 580.
2. Delete all peer certificates that are set up for the Sun StorEdge EBS host whose credentials were created in step 1. For more information, see "How to Delete Sun StorEdge EBS Local Host Peer Resources" on page 581.

How to Export 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 StorEdge 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 will be saved to the location specified.

How to Import 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 interface, click Configuration.
2. Open the Local Hosts folder.
3. From the Local Hosts list, right-click the Sun StorEdge 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.

How to Create New Local Host Certificate Keys

To create new local host certificate keys:

1. From the Administration interface, click Configuration.
2. Open the Local Hosts folder.
3. From the Local Hosts list, right-click the Sun StorEdge 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 StorEdge EBS host.
7. On all other local host resources, delete the peer resource that was set up for the Sun StorEdge EBS host whose certificate was changed in this procedure. For more information, see "How to Delete Sun StorEdge EBS Local Host Peer Resources" on page 581.

Note: Consider exporting the new certificate information to a credentials file. For information about the benefits of doing so, see "Example 2: Local Host Credentials Are Exported" on page 578.

Maintaining Local Host Peer Resources

Each Sun StorEdge EBS host has unique authentication credentials that are used during the **nsrauth** strong authentication process. These credentials include certificate information.

Each Sun StorEdge EBS host retains a record of the certificate of all Sun StorEdge EBS hosts to which it has successfully authenticated. These records are maintained as part of each Sun StorEdge EBS host's *peer* resource database.

How to Delete Sun StorEdge EBS Local Host Peer Resources

When the credentials for a Sun StorEdge EBS host change, so does its certificate. All hosts that have previously authenticated with the changed Sun StorEdge 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 interface, click Configuration.
2. Open the Local Hosts folder.
3. From the Local Hosts list, click the Sun StorEdge EBS host whose peer resource must be deleted.
4. Right-click the peer resource that corresponds to the Sun StorEdge EBS host whose credentials were changed, then select Delete.

How to Create 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 StorEdge EBS host whose credentials have changed, if this has not already been done. For more information, see "How to Export Local Host Credentials" on page 579.
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 the following 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 StorEdge EBS host whose credentials have changed. For more information, see "Example 1: Securing the Initial Peer Host Authentication Process" on page 576.

Multiple Clients from the Same Computer

The Sun StorEdge 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 the following 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 StorEdge 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 may want to 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. For more information on schedules, see "Schedules" on page 144.

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.

How to Define a Client and Save Set Combination

To define a client and save set combination:

1. Create or edit a Sun StorEdge 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. Do one of the following:
   - 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 more than one save set on a client, enter each save set (partition, file system, or file) on a separate line in the Save Set attribute.

Adding or Changing the Sun StorEdge EBS Servers for a Client

Sun StorEdge EBS uses the contents of the */nsr/res/servers* (UNIX), or the <*Sun StorEdge EBS_install_path*> \ *res* \ *servers* (Windows) file on each Sun StorEdge EBS client to control who has client-tasking rights (the right to request a program to be executed on another client). This tasking might be any of the following:

- Server performing an archive request
- Scheduled backup
- Another client requesting a directed recover

During a Sun StorEdge EBS software installation, you can add the names of Sun StorEdge EBS servers to this file. To add additional hosts at a later date, a text editor must be used to add the host name to the file.

- So that the client with the tasking rights can back up to other Sun StorEdge EBS servers, the names of the additional Sun StorEdge 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 StorEdge EBS host has tasking rights. This is a potential security concern.

How to Edit the servers File

To give clients tasking rights to a Sun StorEdge EBS client:

1. Stop the Sun StorEdge EBS host whose tasking rights are being updated. For more information, see "Stopping and Starting a Sun StorEdge EBS Server, Client, or Storage Node" on page 51.
2. Open the *servers* file in a text editor.

   The default installation location for this file is:

   - */nsr/res/servers* (UNIX)
   - <*Sun StorEdge EBS_install_path*> \ *res* \ *servers* (Windows)
3. Enter one host name per line.

4. Save the changes and exit the text editor.
5. Restart the Sun StorEdge 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 StorEdge EBS server. If a value is not specified in the Priority attribute, the back up 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 the following 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 hangs and times out, it 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. For more information, see "Task 2: Set Up a Group for Backup Clients" on page 56.

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.

Chapter 19: Cluster Support

This chapter covers the following topics:

- "Cluster Terminology" on page 587
- "Cluster-Aware and Cluster-Unaware Sun StorEdge EBS Servers" on page 588
- "Backing Up Data on UNIX Cluster Environments" on page 589
- "Configuring a Virtual Client to Back Up to a Local Storage Node" on page 592"Configuring an External Client to a Virtual Server" on page 593"Recovering Data in a UNIX Cluster Environment" on page 594
- "Tracking Scheduled Saves" on page 595
- "Using a Physical Storage Node with a Virtual Server on UNIX Cluster Platforms" on page 599
- "Using Autochangers and Stand-Alone Tape Devices" on page 600
- "Sun StorEdge EBS Cluster Licensing" on page 614
- "Client Configuration Wizard Support Requirements" on page 616

Cluster Terminology

This section defines terms specific to Sun StorEdge 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 StorEdge EBS server* — A Sun StorEdge EBS server that can recognize cluster clients.
- *Cluster client* — A Sun StorEdge EBS client within a cluster. This can be either a virtual client, or a Sun StorEdge EBS Client resource that backs up the private data belonging to a one of the physical nodes.
- *Virtual client* — A Sun StorEdge EBS Client resource that backs up data belonging 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 StorEdge EBS server running within a cluster, but *not* configured as a highly available application. A stand-alone server does not have failover capability.
- *Virtual server* — A Sun StorEdge EBS server configured to run as a highly available application. A virtual server can fail over from one cluster node to another.

19

Cluster-Aware and Cluster-Unaware Sun StorEdge EBS Servers

To perform scheduled backups within a cluster, you must have a cluster-aware Sun StorEdge EBS server and Sun StorEdge EBS client. Not all Sun StorEdge EBS servers and clients are cluster-aware. For more information about which versions are cluster aware, refer to the *EMC Legato Software Compatibility Guides*.

The following issues apply if either the server or client are not cluster-aware:

- If a Sun StorEdge EBS *server* does not have cluster-client support, it cannot recognize cluster clients. When a scheduled backup is performed, the Sun StorEdge 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 StorEdge 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.

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 an ad hoc backup of data from a virtual cluster client. For more information, see the following topics:

- "How to Configure a Scheduled Save from a Physical Cluster Client" on page 589
- "How to Configure a Scheduled Save from a Virtual Cluster Client" on page 590
- "Performing an Ad Hoc Backup from a Virtual Cluster Client" on page 591

How to Configure 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. For detailed instructions on how to back up data, see "Chapter 2: Backing Up Data" on page 53.

To prepare for regular backups, install the Sun StorEdge EBS client software on all physical clients within the cluster. For detailed installation and configuration instructions, refer to the appropriate *Sun StorEdge EBS Installation Guide*.

How to Configure a Scheduled Save from a Virtual Cluster Client

Note: Instructions for how to configure a scheduled save from a virtual cluster client are the same whether the Sun StorEdge EBS server is inside the cluster or outside the cluster.

To configure a scheduled save from a virtual cluster client under the Sun StorEdge EBS server:

1. (Optional) Define save groups as required.



Important: 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 StorEdge EBS server failover.

2. Make each physical client within the cluster a Sun StorEdge EBS client of the Sun StorEdge EBS server.

   For each physical client in the cluster:

   a. Create a new Sun StorEdge EBS client.
   b. For the Name attribute, enter the name of the physical client.

3. Make each virtual client within the cluster a Sun StorEdge EBS client of the Sun StorEdge EBS server.

   For each virtual client in the cluster:

   a. Create a new Sun StorEdge 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 StorEdge EBS server that resides in a cluster, the first time the Sun StorEdge EBS server runs, it creates the Client resource for the virtual Sun StorEdge 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 StorEdge EBS server resides, enter the following command:

   savegrp -pv -c *virtual_client group_name*

Important: If the test probe does not display the appropriate scheduled backups and index, see "Tracking Scheduled Saves" on page 595.

Performing an Ad Hoc 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 *saveset*

or

save -c *client saveset*

How to Perform an Ad Hoc Backup by Using Nonroot Administrative Privileges

For all cluster platforms except HP MC/ServiceGuard, AutoStart, and VERITAS Cluster Server, performing an ad hoc backup by using nonroot privileges is similar to backing up data from noncluster clients by using root privileges. For example:

save *saveset*

or

save -c *client saveset*

For more information on performing ad hoc backups by using nonroot privileges on the following the following platforms, see:

- "HP MC/ServiceGuard" on page 591
- "AutoStart for AIX, HP-UX, Linux, and Solaris" on page 592
- "VERITAS Cluster Server" on page 592

HP MC/ServiceGuard

To perform an ad hoc backup of data from a virtual cluster client with nonroot privileges on HP MC/ServiceGuard, do one of the following:

- 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 the following 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 an ad hoc backup of data from a virtual cluster client with nonroot privileges on AutoStart for AIX, HP-UX, Linux, and Solaris use the **ftconsole** to add nonroot users to the security access list for the FT domain.

VERITAS Cluster Server

To perform an ad hoc backup of cluster data that belongs to a physical or a virtual cluster client with nonroot privileges using VERITAS Cluster Server, set the VERITAS Cluster Server attribute AllowNativeCliUsers to 1, so that the nonroot user is not prompted for a password when executing VERITAS Cluster Server commands. This attribute instructs the VERITAS Cluster Server to use the user's operating system username and password, rather than VERITAS Cluster Server username and password, for authentication. The VERITAS Cluster Server default Cluster Guest permissions are sufficient to use Sun StorEdge EBS client software. For more information, refer to the VERITAS Cluster Server documentation.

Configuring a Virtual Client to Back Up to a Local Storage Node

Typically, the Sun StorEdge 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. This list is prioritized by the order of the nodes.

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. This is done by entering "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 StorEdge 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 StorEdge 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.

How to Direct 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 StorEdge 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 “curphyhost” in the list based on the required priority.

Configuring an External Client to a Virtual Server

Before a Sun StorEdge EBS server can back up a client, the client must grant the server access. Granting access is controlled by the following files:

On UNIX: */nsr/res/servers*

On Windows: <*Sun StorEdge EBS_install_path*> \ *res* \ *servers*

Only the specific servers listed in this file can back up the client. For more information about the *servers* file, see "Adding or Changing the Sun StorEdge EBS Servers for a Client" on page 585.

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. For more information, see the following topics:

- "How to Recover Data from a Private Disk Belonging to a Physical Client" on page 594
- "How to Recover Data from a Shared Disk Belonging to a Virtual Client" on page 594

How to Recover Data from a Private Disk Belonging to a Physical Client

Recovering data backed up from a *private disk* belonging to a physical client in a cluster is similar to recovering data on a computer that is not part of a cluster.

For detailed instructions, see "Chapter 13: Recovering Data" on page 367.

How to Recover Data from a Shared Disk Belonging to a Virtual Client

To recover data backed up from a *shared disk* belonging to a virtual client:

1. Make sure that you have properly configured the virtual client for remote access from the node that you are currently on.
   a. Start the Sun StorEdge EBS software.
   b. From the Administration window, click Configuration.
   c. In the left pane, select Clients.
   d. From the File menu, select Properties and then select the select the Globals 2 of 2 tab.
   e. 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 belonging 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 recover, 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*

For more information, refer to the **recover** man page.

How to Recover the Console Server Database

The Console database contains management data such as report information. For information about recovering a corrupted Console database, refer to the *NetWorker Disaster Recovery Guide*.

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 StorEdge 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 the following 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.



Important: Misconfiguration of the cluster resources might cause scheduled save sets to be dropped from the backup. For more information about configuring clusters for scheduled backups, see "Path-Ownership Rules" on page 597.

To override scheduled save rules, create an empty file named *pathownerignore* on each node in the cluster by running the following 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.



Important: 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 StorEdge 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. Start the Sun StorEdge EBS software.
2. From the Administration window, click Configuration.
3. In the left pane, select clients.
4. In the right pane, select the client that is using the incorrect index.
5. From the File menu, select Properties and then select the Apps & Module tab.
6. For the Backup Command attribute, enter the name of a backup script that contains **save -c** *client_name*.

For detailed information about the Backup Command attribute, see "Using the save Command with a Customized Backup Program" on page 88.

Path-Ownership Rules

In a cluster environment, the Sun StorEdge EBS software must distinguish between file systems associated with a physical client and 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 StorEdge EBS software might not back up scheduled save sets due to conflicts with path ownership. This prevents a virtual Sun StorEdge EBS client from writing to multiple client file indexes. The Sun StorEdge 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.

The following 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.



Important: A user must have Administrative privileges to query the cluster database. If a non-administrative user performs a save of a file system belonging 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 StorEdge EBS software to save the files under the virtual node's client file index, use the **-c** option. For more information, see "Forcing Save Sets to Use a Specific Index" on page 596.

How to Check Path-Ownership Rules

To check the Sun StorEdge EBS path-ownership rules:

1. Enter the following command at the command-prompt on the Sun StorEdge 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. For detailed instructions, refer to the appropriate *Sun StorEdge EBS Installation Guide*.

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 StorEdge EBS default path-ownership rules, you can create the *pathownerignore* file in the following directory:

/usr/sbin/nsr/bin

This file causes the Sun StorEdge 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.

If the scheduled save uses the wrong index, see "Forcing Save Sets to Use a Specific Index" on page 596.

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 an Ad hoc Backup

To override the default path-ownership rules for an ad hoc backup, enter 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. Start the Sun StorEdge EBS software.
2. From the Administration window, click Configuration.
3. In the left pane, select Clients.
4. In the right pane, select the client that is using the incorrect index.
5. From the File menu, select Properties and then select the Apps & Module tab.
6. For the Backup Command attribute, enter the name of a backup script that contains **save -c** *client_name*.

For detailed information about the Backup Command attribute, see "Using the save Command with a Customized Backup Program" on page 88.

Using a Physical Storage Node with a Virtual Server on UNIX Cluster Platforms

This section describes how to access a physical storage node with a virtual Sun StorEdge EBS server.

A Sun StorEdge EBS storage node can be installed and configured on nodes in the cluster. However, the Sun StorEdge EBS storage node is not a highly available application. The following instructions apply in the cases where you are using storage nodes in the cluster and storage nodes that are outside of the cluster.



Important: Sun StorEdge EBS release 6.0 and later does not officially support remote virtual storage nodes.

To access a physical Sun StorEdge EBS storage node from a virtual server (rd=*physicalhostname:devxxx*):

1. Shut down the virtual Sun StorEdge EBS server using the cluster management software.
2. Modify the Sun StorEdge EBS virtual cluster **start** script to recycle the remote Sun StorEdge EBS storage node.

   On each node to which the virtual Sun StorEdge EBS server can fail over:

   a. Open the Sun StorEdge EBS virtual cluster script, *ebs_bin_dir*, for editing.

   b. Before the line that starts the **nsrd** daemons, add the following:

   `rsh remote_hostname "`*ebs_bin_dir*`/nsr_shutdown -q && `***ebs_bin_dir*`/nsrexecd"`**

   Note: Replace the *ebs_bin_dir* reference in the script edit with the appropriate path where the Sun StorEdge EBS binaries are installed on the Sun StorEdge EBS storage node.

   c. If necessary, resynchronize the virtual Sun StorEdge EBS server **startup** script with the cluster management system.
3. Restart the Sun StorEdge 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 StorEdge 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 600.
- Shared autochanger devices, but *not* shared stand-alone tape devices, use "Configuration Option 2" on page 602.
- Nonshared devices, or any configuration, use "Configuration Option 3" on page 604, which does not require device-sharing infrastructure.

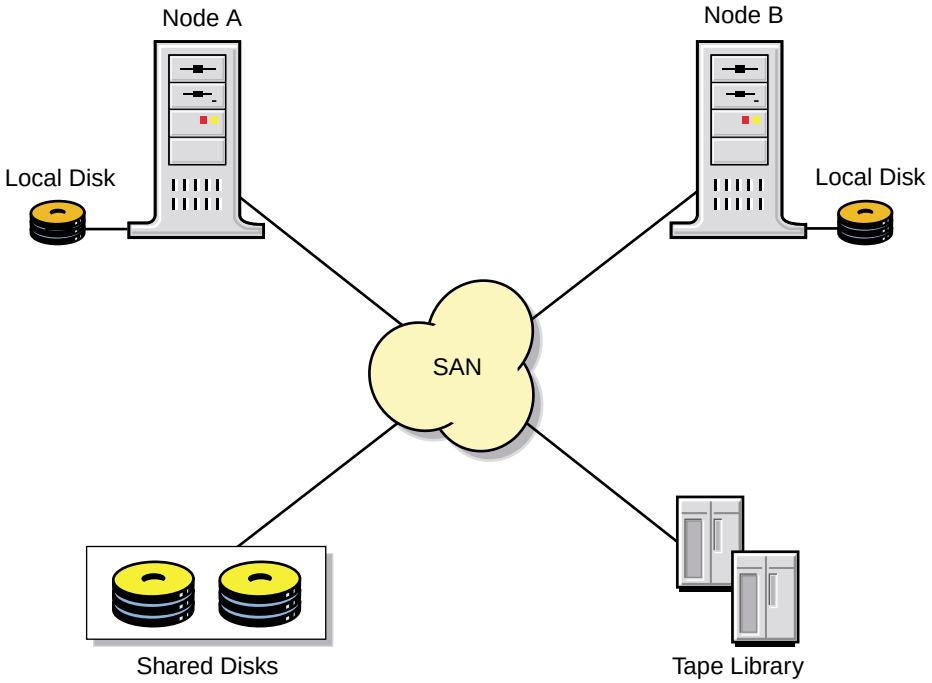
Note: Make sure that the device-sharing infrastructure is certified by the cluster software vendor.

Configuration Option 1

This configuration option offers a seamless backup solution that is 100% hands-free. It allows you to access shared tape devices and shared autochanger devices with a highly available Sun StorEdge EBS server. Figure 37 on page 601 provides a graphical view of this configuration option.

Note: Only one autochanger is required, although you can add more.

Figure 37. Configuration Option 1



To configure shared tape devices and shared autochanger devices to work with a highly available Sun StorEdge EBS server:

1. Interface all the autochanger devices to all the nodes in the cluster to which the Sun StorEdge 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 StorEdge EBS server currently running. The path must be protected from stray bus signals and unauthorized session access from the other nodes.

Important: If processes on nodes other than the one that is running on the Sun StorEdge EBS server can access the tape devices, data corruption might occur. The Sun StorEdge EBS software might not detect the data corruption.

3. Modify the virtual Sun StorEdge 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.



Important: 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**, **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 back the data from tape and comparing it to the original data.

      If the verification check indicates:

      – Failure — the switching infrastructure does not properly support path isolation. Try Option 2 or 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 back the data from tape and comparing it to the original data.
      – If the verification fails, try Option 2 or Option 3, or contact the device-sharing infrastructure vendor.
5. Configure the autochanger and tape devices.

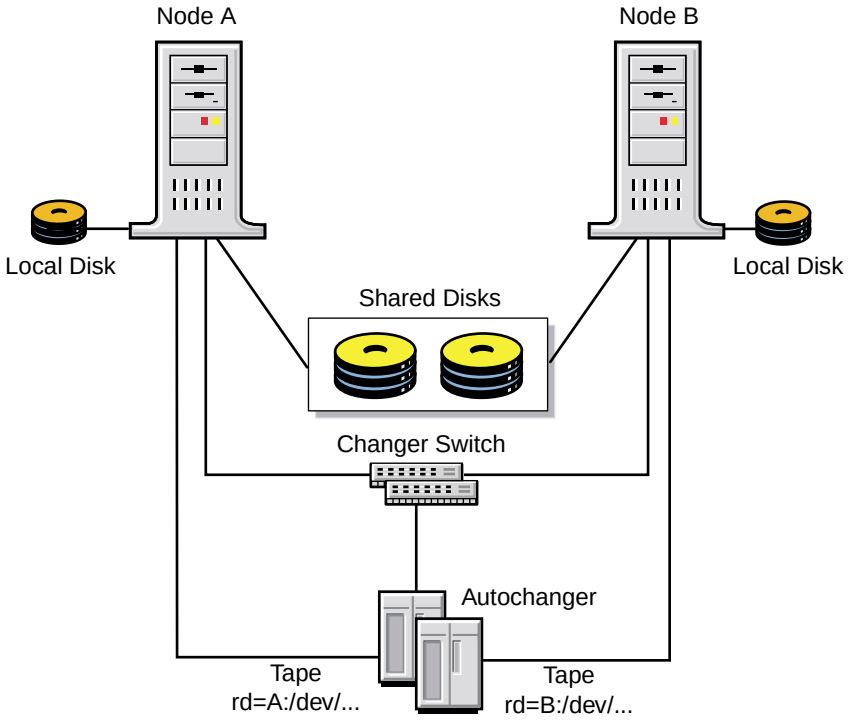
Configuration Option 2

This option offers a seamless backup solution for accessing nonshared tape devices and shared autochanger devices from a highly available Sun StorEdge EBS server. However, some intervention might be required during the data recovery process. Figure 38 on page 603 provides a graphical view of this configuration option.

Notes:

- If *N* nodes are used, *N-1* storage node licenses are required. *N* indicates the number of failover nodes for the virtual Sun StorEdge EBS server.
- Only one autochanger is required, although you can add more.

Figure 38. Configuration Option 2



To configure the devices to work with a virtual Sun StorEdge EBS server:

1. Interface the autochanger device to the Sun StorEdge 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.

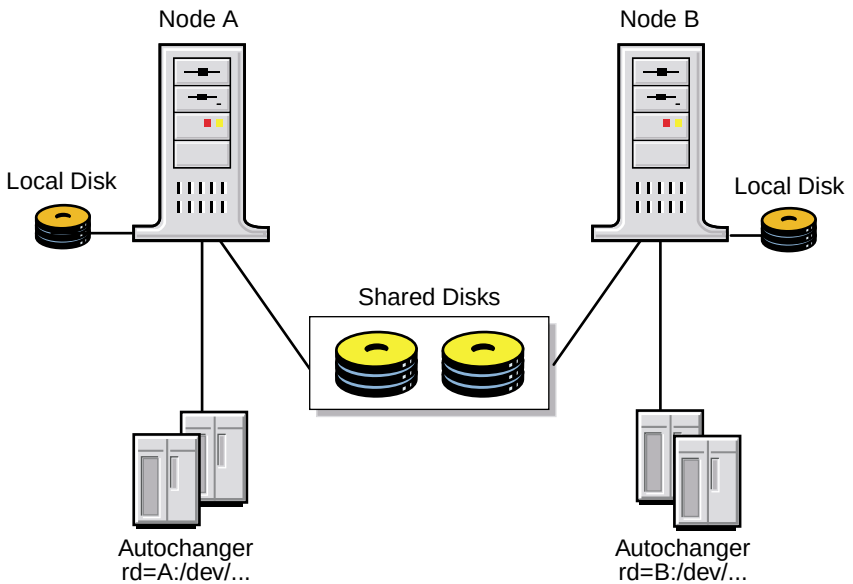
Configuration Option 3

This option offers a seamless backup solution for accessing nonshared tape devices and nonshared autochanger devices from a highly available Sun StorEdge EBS server. However, some intervention might be required during the data recovery process. Figure 39 on page 604 provides a graphical view of this configuration option.

Note: *N* indicates the number of failover nodes for the virtual Sun StorEdge 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 39. Configuration Option 3



To configure tape autochanger devices to work with a virtual Sun StorEdge EBS server:

1. Attach the autochanger exclusively to the Sun StorEdge EBS server.
2. Use the following 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 StorEdge 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 606
- "Controlling the Robotics from a Stand-Alone Host" on page 611.

For information explaining how to activate the DDS feature, see "Dynamic Drive Sharing" on page 108.

The following sections explain the supported cluster platforms and the two methods for using DDS within a cluster.

High Availability and DDS

The Sun StorEdge 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 StorEdge EBS server and savegroups.

The nsrjb Program High Availability Limitations

If the Sun StorEdge 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: Host Crash Requires User Intervention

You have the following 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 reset command **nsrjb -H** (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 StorEdge 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 StorEdge EBS server. This virtual Sun StorEdge 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 StorEdge EBS server to install the autochanger.

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.

Note: 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.

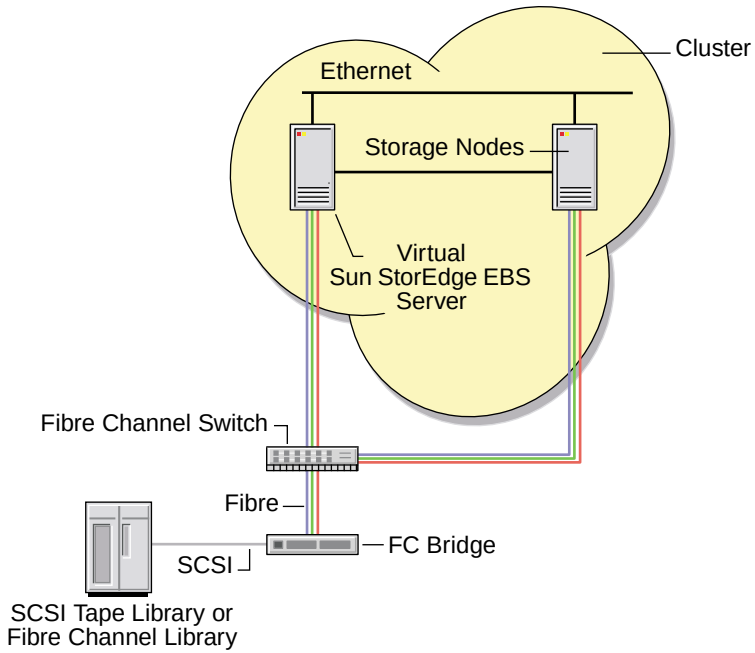
If the special device files are not bundled in with the operating system, refer to the appropriate *Sun StorEdge EBS Installation Guide*.

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 StorEdge 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 40 on page 607 illustrates how DDS can be enabled within a cluster.

Figure 40. 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: On Sun Cluster with No Special Device File" on page 609.

The following 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: On Sun Cluster with No Special Device File

Using *scsidev@2.3.0.* for the device, you would 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
```

Note: From this point forward, the **jbconfig** program is essentially configured the same way for both clusters with or without vendor-supplied drivers.

```
Pathname of the control port for the jukebox device?
scsidev@2.3.0

Turn NetWorker auto-cleaning on (yes / no) [yes]? no

The jukebox being configured is unable to locate its
drives.

You will need to provide this information.

Is (any path of) any drive intended for NDMP use? (yes /
no) [no] no

Is any drive going to have more than one path defined?
(yes / no) [no] yes

You will be prompted for multiple paths for each drive.

Pressing <Enter> on a null default advances to the next
drive.

Please enter the device path information in one of the
following formats:

   /dev/rmt/1cbn --for local path or
   host:device-path --for remote node or
   host:drive-letter:directory path --for Windows disk
   file

Drive  1, element 1

Device path 1 ? bigstar.legato.com:/dev/rmt/0cbn
Device path 2 ? [] littleman.legato.com:/dev/rmt/0cbn
Device path 3 ? []
```

```
Drive  2, element 2

Device path 1 ? bigstar.legato.com:/dev/rmt/1cbn
Device path 2 ? [] littleman.legato.com:/dev/rmt/1cbn
Device path 3 ? []

Please select the appropriate drive type number:

1) 3480
2) 3570
3) 3590
4) 4890
5) 4mm
.
.
.
40) tkz90
41) travan10
42) tz85
43) tz86
44) tz87
45) tz88
46) tz89

Enter the drive type of drive 1? 46

Are all the drives the same model? (yes / no) [yes]

Jukebox has been added successfully

The following configuration options have been set:

> Autocleaning off.

> At least one drive was defined with multiple paths.
All such drives are defined with a hardware
identification as well as a path value to avoid
confusion by uniquely identifying the drive.

The hardware identification for all drives which have
one is always 'autochanger_name - Drive #' where
"autochanger_name" is the name you gave to the
autochanger that was just defined, and the # symbol is
the drive number.

> Barcode reading to on.

> Volume labels that match the barcodes.

> Slot intended to hold cleaning cartridge to 10. Please
insure that a cleaning cartridge is in that slot >
```

```
Number of times we will use a new cleaning cartridge to
20.You can review and change the characteristics of the
autochanger and its associated devices using nwadmin.

Would you like to configure another jukebox?
(yes/no)[no] no
```

Controlling the Robotics from a Stand-Alone Host

You can set up a stand-alone physical host outside the cluster to control the robotics. This procedure is useful:

- 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 StorEdge 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 the following binaries are installed:

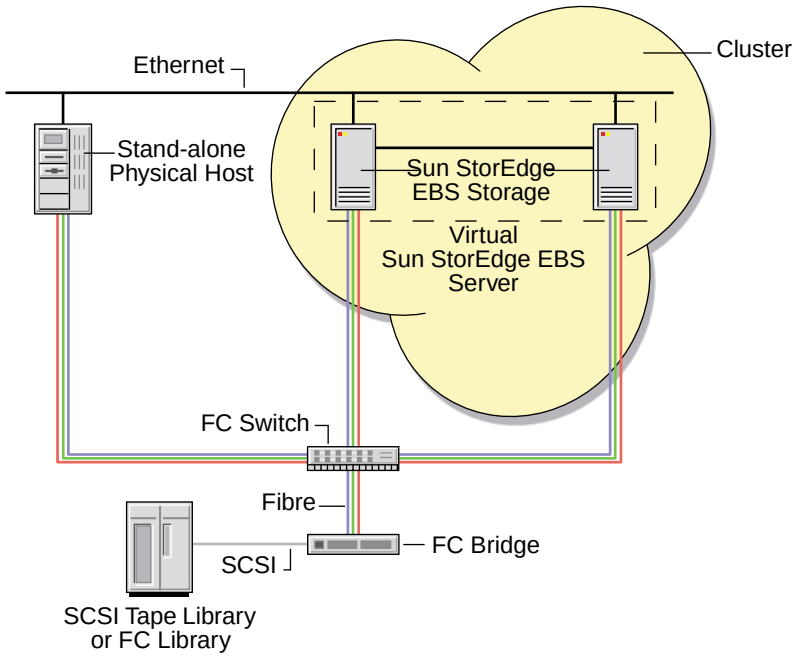
- Client software
- Storage node software
- Driver software

You would configure the **jbconfig** program on the stand-alone physical host to control the robotics.

Figure 41 on page 612 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 41. DDS Enabled with the Robotics Controlled Outside the Cluster

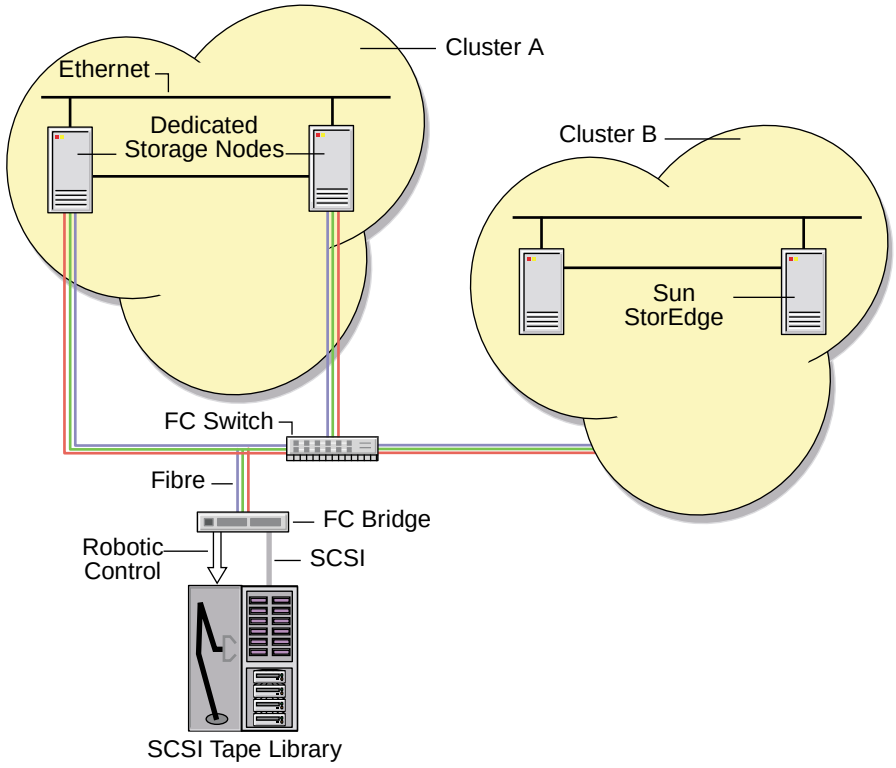


19

Controlling Robotics from a Sun StorEdge EBS Server on Another Cluster

A Sun StorEdge 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 42 on page 613. The virtual Sun StorEdge EBS server of cluster B controls the robotics for the physical storage nodes in cluster A.

Figure 42. Robotics Control between Clusters



Sun StorEdge EBS Cluster Licensing

This section discusses the following topics:

- "Enabling the Highly Available Sun StorEdge EBS Server" on page 614
- "Enabling Cluster and Virtual Clients" on page 614
- "Licensing Requirements for Backing Up Cluster Clients" on page 615

Enabling the Highly Available Sun StorEdge EBS Server

A Cluster Client Connection license is required for each physical node in the cluster on which you intend to run the highly available Sun StorEdge EBS server.

In addition, a storage node that is configured locally—using a hostname that matches the physical hostname of the node running the Sun StorEdge 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 StorEdge EBS virtual server was running on *Node A*. An example of this configuration is provided in "Configuration Option 3" on page 604.

Enabling Cluster and Virtual Clients

Sun StorEdge EBS client licensing differentiates between stand-alone computers and computers that participate in a cluster. A Sun StorEdge EBS server backing up a computer that is a member of a cluster requires a Cluster Client Connection license rather than a Client Connection license.

Cluster Client Connection licenses are distributed on a first-come-first-served basis. Each time a cluster client (physical or virtual) is backed up for the first time, a Cluster Client Connection license is allocated to the physical node where the virtual client is running. Cluster Client Connection licenses are bound to physical nodes. Therefore, once a Cluster Client Connection is allocated, any virtual clients running on that physical node can be backed up.

Cluster Clients

To back up a cluster client, the Sun StorEdge EBS server requires:

- A Client resource defined for that member of the cluster
- A Cluster Client Connection license per physical host base.

Once the Cluster Client Connection license count is exhausted, you must do one of the following:

- Purchase more Cluster Client Connection licenses
- Remove Sun StorEdge EBS cluster Client resources with a Cluster Client Connection license

If a remote cluster client does *not* use the same platform as the Sun StorEdge EBS server, you need the appropriate ClientPak enabler. For example, if the virtual Sun StorEdge EBS server is on a Sun Cluster 3.0 system and the cluster client is a member of an HACMP for AIX environment, you need a Cluster Client Connection license and a UNIX ClientPak enabler.

Virtual Clients

To back up virtual clients, the cluster node that is running the virtual client must be licensed within the Sun StorEdge EBS server with a Cluster Client Connection license, which is physical host based. The virtual client itself does not require a Cluster Client Connection license, but the physical host it resides on should be licensed with a Cluster Client Connection license.

Licensing Requirements for Backing Up Cluster Clients

Backing up cluster clients requires a cluster-aware Sun StorEdge EBS server. A Sun StorEdge EBS server does *not* have to be on a cluster to back up cluster clients.

Cluster-aware Sun StorEdge EBS servers can run on stand-alone computers or computers that participate in a cluster. All Sun StorEdge EBS release 6.0.1 and later servers are cluster-aware.

For a list of cluster-aware servers earlier than Sun StorEdge EBS release 6.1, refer to the *EMC Legato Software Compatibility Guides*.

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.

To back up a NetWorker Module 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, Cluster Client Connection licenses are required for both of the physical nodes.

Client Configuration Wizard Support Requirements

This section contains the requirements to use the Client Configuration Wizard in a cluster environment.

- On a highly available Sun StorEdge 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 StorEdge 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.

For additional general considerations that apply to the Client Configuration Wizard, see "Using the Sun StorEdge EBS Client Configuration Wizard" on page 62.

Chapter 20: Sun StorEdge EBS Server Support for NDMP

This chapter covers the following topics:

- "Overview of NDMP" on page 617
- "Supported NDMP Configurations" on page 619
- "Requirements for and Limitations of NDMP Support" on page 625
- "Configuring the Sun StorEdge EBS Server for NDMP Operations" on page 633
- "Performing NDMP Backups" on page 647
- "Performing NDMP Recoveries" on page 655
- "NDMP Cloning" on page 665
- "Ancillary Product Information" on page 667

Overview of NDMP

Network Data Management Protocol (*NDMP*) is a TCP / IP-based protocol that specifies how network components talk to each other for the purpose of moving data across the network for backup and recovery.

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 Sun StorEdge EBS for backups, cloning, recoveries, and updates or queries to resource files. All communications are through NDMP.

20

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 StorEdge 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.

NDMP Connection Support

With the Sun StorEdge 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 StorEdge EBS server) do not need to have the Sun StorEdge 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 StorEdge EBS program.

To install an NDMP data module or enable NDMP support on an application server, refer to the respective vendor's documentation.

Licensing

The Sun StorEdge EBS NDMP interface is an optional feature that must be licensed separately from the Sun StorEdge EBS base product according to a tiered licensing structure. One NDMP Client Connection license must be applied to each *NDMP host*, although a single host can have multiple Client resources.

Use of the Sun StorEdge EBS storage node software with NDMP support requires a software license (where the Sun StorEdge EBS NDMP interface is not used).

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.

Supported NDMP Configurations

The Sun StorEdge 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 StorEdge EBS Storage Node — The backup is performed from an NDMP host to a Sun StorEdge EBS tape device, disk, optical, and other media type attached to a Sun StorEdge EBS storage node.

  Sun StorEdge EBS storage nodes use a data server agent (DSA). DSA acts as an agent in the performance of backups between a Sun StorEdge EBS server and any client without Sun StorEdge EBS software. An example of a DSA is an NDMP host that generates proprietary save data and sends that data to a Sun StorEdge EBS storage device to have a save set associated to it. Similarly, a DSA will act as an agent in the recovery between a NetWorker server and any client without NetWorker software.

  Note: Any backup performed with the DSA feature has a **D** flag.
- Three-party using 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 43 on page 620 as file history) traverse the network between the backup server and the NDMP tape and data server. The Sun StorEdge EBS program uses the file history information to maintain appropriate indexes and media databases.

Figure 43 on page 620 illustrates a local backup type configuration. The NDMP host is both an NDMP data service provider (DSP) and *tape service*.

The Sun StorEdge 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.

20

The DSP performs these tasks:

- Receives the backup requests.
- Reads the backup data from disk.
- Produces a data stream for backup.

As a result, these actions occur:

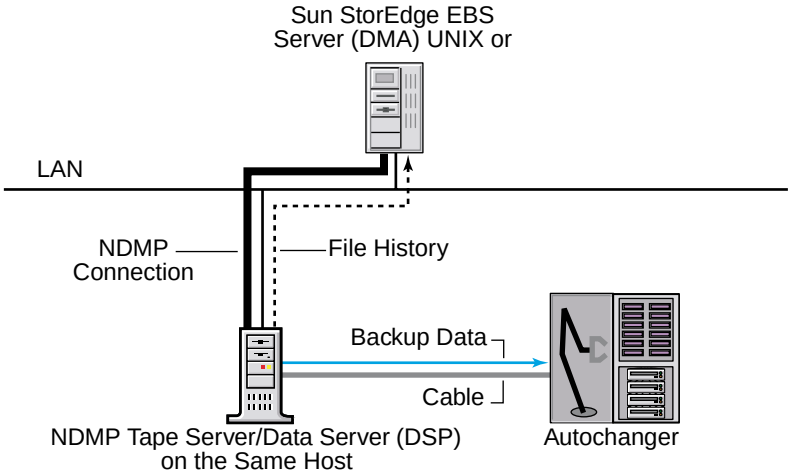
- During a backup, the tape service reads a data stream from the DSP and writes it to tape.
- During a recovery, the tape service reads data from tape and generates an NDMP data stream.

  The following then occurs:

  a. The data stream is sent to the DSP.

  b. The DSP reads from the data stream, and writes the data stream to disk.

Figure 43. NDMP Local Configuration



Backup by Using Sun StorEdge EBS Storage Node Support

This section features the supported NDMP client backup configurations which use non-NDMP based Sun StorEdge EBS devices (non-NDMP devices), such as a tape, disk, or optical.

Figure 44 on page 622 illustrates a Sun StorEdge EBS storage device attached directly to the Sun StorEdge EBS server. The NDMP backup that uses DSA, is initiated from the Sun StorEdge 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 StorEdge EBS server (**nsrmmd**).

Table 93 on page 621 shows how DSA differs in use from NDMP tape server.

Table 93. Distinctions between NDMP Tape Server and DSA

NDMP Tape Server	DSA
Supports only NDMP type of tape device	Supports any type of Sun StorEdge 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.

For more information on storage technologies under Sun StorEdge EBS storage node support for NDMP clients, see "Data Write Technologies" on page 652.

20

Figure 44. Backup Initiated from a Sun StorEdge EBS Server with an Attached Storage Device

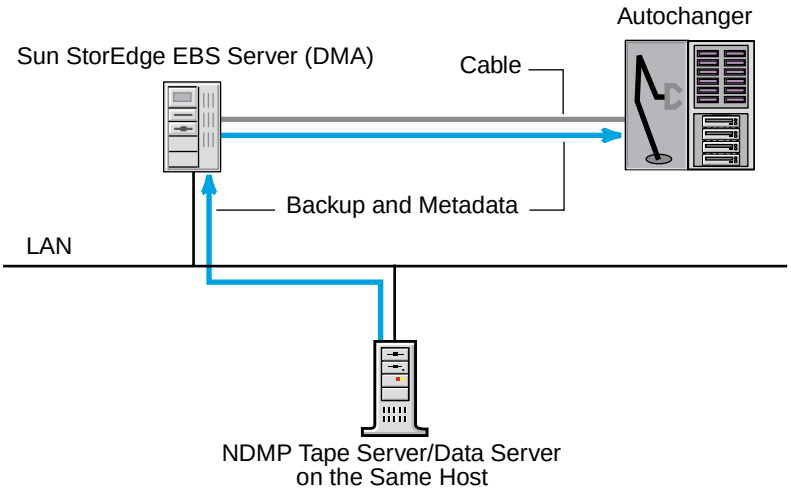
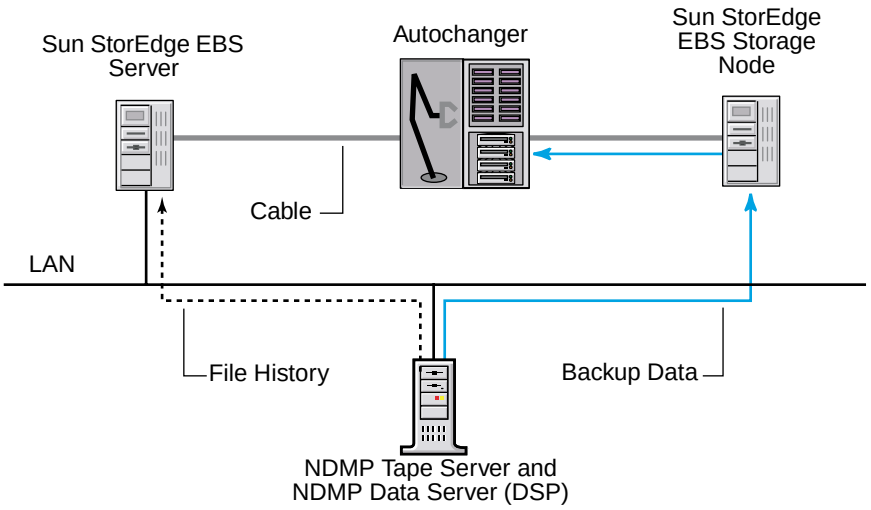


Figure 45 on page 623 illustrates a Sun StorEdge EBS configuration in which the file history is backed up to the Sun StorEdge EBS server, and the data is backed up to the Sun StorEdge EBS client with an attached storage node. The backup command, **nsrndmp_save**, is run directly on the Sun StorEdge EBS client that has a Sun StorEdge EBS storage node locally attached.

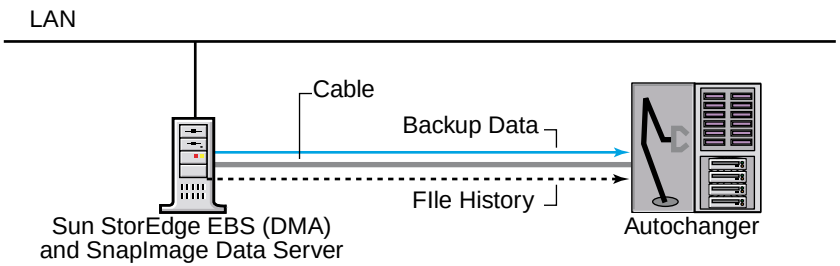
Data is processed and saved to the Sun StorEdge EBS storage node, while the file history is processed on the designated Sun StorEdge EBS server (**nsrmmd**). After the data is backed up and sessions with the NDMP server and the Sun StorEdge EBS server are closed, the file history is committed to the Sun StorEdge EBS server.

Figure 45. File History and Data to Sun StorEdge EBS Storage Node



In Figure 46 on page 623, the file history and data are backed up to the Sun StorEdge EBS server, which is also a 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 44 on page 622.

Figure 46. Backup Using 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 EMC, NetApp Filer, Auspex, or SnapImage Module.

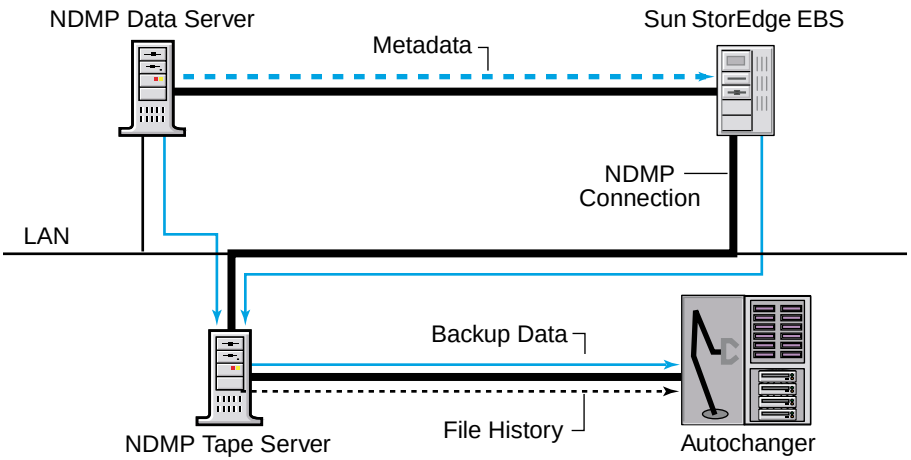
Important: The two NDMP servers must be of the same type. For example, backups between an EMC NAS device and a NetApp server are not supported. A NetApp server must back up from, and recover to, another NetApp server.

Figure 47 on page 624 demonstrates a three-party configuration, which enables backup and recovery to a Sun StorEdge 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 StorEdge EBS server (DMA).

This configuration is similar to the flow of data between a Sun StorEdge EBS client and a Sun StorEdge 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.

Figure 47. NDMP Tape Server



Supported Tape Servers

The Sun StorEdge EBS software with NDMP supports these tape servers:

- Sun StorEdge EBS NDMP TapeServer 2.0 software installed on the Sun StorEdge EBS server or another server.
- NetWorker SnapImage Module software installed on the Sun StorEdge 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.

Requirements for and Limitations of NDMP Support

These sections contain information about requirements, support, and limitations of the Sun StorEdge EBS software for NDMP, including "Sun StorEdge EBS Storage Node Support for NDMP Clients" on page 630.

Requirements for Sun StorEdge EBS Software with NDMP Support

The following requirements apply when running the Sun StorEdge EBS software with NDMP support:

- To back up the bootstrap file, a *non-NDMP device* must be attached to the Sun StorEdge EBS server. The bootstrap includes the media database and configuration files needed to recover the Sun StorEdge EBS server after a disk crash.

  This requirement does not apply when the backup is to a Sun StorEdge 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 the following:
  - 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 Japanese data in a localized Japanese NDMP environment, the Index Recover status window shows the process in English, rather than Japanese.

Limitations of the Sun StorEdge EBS Software with NDMP Support

The addition of NDMP support to a Sun StorEdge EBS datazone environment renders the following limitations:

- Browse other NDMP hosts for save set recoveries through use of the command-prompt, the Sun StorEdge EBS User (Windows), or **nwrecover** (UNIX) program. NDMP hosts cannot be browsed for backups.
- These 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 StorEdge EBS server and any client without Sun StorEdge EBS software that generates data and sends it to a Sun StorEdge 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

    To verify that the Sun StorEdge EBS server can see the library on the NDMP server, either reset or inventory the library from the Sun StorEdge EBS Console:

    1. From the Console window, click Enterprise.
    2. Click the NetWorker server hostname.
    3. Double-click NetWorker.

       The NetWorker Administration window opens.

4. Click Libraries.
5. Double-click the library to be reset or inventoried.
6. Right-click the Device name.
7. Select Inventory or Reset.

– SnapImage, DSA, and Sun StorEdge EBS Storage Node Support for NDMP Clients

  Use of DSA for backups and recoveries with a non-NDMP Sun StorEdge EBS storage node that includes the NetWorker SnapImage Module for UNIX (Release 1.6.1) and for Microsoft Windows (Release 2.0.2) is not supported.

- The **scanner** program generates media database entries for backed-up data. The **scanner** program cannot read the actual data. The utility can re-create the volume and save set information in the media database, if needed. Index recoveries are not possible. If the scanned volumes belong to a library, perform one of these actions before recovering the generated volumes and save sets.
  – An inventory on the library to synchronize the slots and volumes for automount/unmount
  – A ad hoc mount operation

  If the corresponding index save set is available, indexes can be brought back by using **nsrck L7**. The client must be configured after the save set is scanned back in from the tape.

  Note: After running the **scanner** program, inventory the volumes. In addition, reset the hardware so that the Sun StorEdge EBS server updates the location of the volume being scanned.

  To reset the hardware:

  1. From the Console window, click Enterprise.
  2. Click the NetWorker server hostname.
  3. Double-click NetWorker.

     The NetWorker Administration window opens.
  4. Click Libraries.
  5. Double-click the library to be reset or inventoried.
  6. Right-click the Device name.
  7. Select Inventory or Reset.

The **scanner –m** command scans the volume (that has DSA save sets) in the same way it does for non-NDMP save sets. DSA save sets are save sets of an NDMP client that are backed up to non-NDMP tape device by using DSA.

The **scanner -i** command fails for DSA save sets and skips them. The **scanner –i** command is not supported for NDMP backups.

- Sun StorEdge 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 StorEdge EBS storage nodes for backup and recovery of data.
- The Sun StorEdge 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 StorEdge EBS software does not support incremental backup settings for NDMP backups. If you select the incremental setting, the Sun StorEdge 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.
- NDMP and non-NDMP save sets on a single volume are not supported, unless the data is backed up to a Sun StorEdge EBS storage node.

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 higher.
- 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 StorEdge EBS software.

Index Database

During an NDMP backup, the NDMP data server sends the file history metadata information to the Sun StorEdge 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).

Note: Especially during failover of a clustered Sun StorEdge EBS environment, if NDMP indexing is interrupted, these 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: 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 this appears.

```
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
error in reading the FH entries from save through stdin
```

Sun StorEdge EBS Storage Node Support for NDMP Clients

Sun StorEdge EBS storage node functionality includes NDMP save streams. Sun StorEdge EBS storage node software can be used for both NDMP and non-NDMP based backups in a unified fashion. For NDMP backups, Sun StorEdge EBS uses the DSA feature.

Clients that generate data have backup history in the form of manageable Sun StorEdge 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 StorEdge EBS Devices

With this extension, Sun StorEdge EBS supports backup of an NDMP client to a non-NDMP based Sun StorEdge EBS device such as:

- Tape
- Optical
- Disk
- Other media types

Features with Sun StorEdge EBS Storage Node Support for NDMP

Sun StorEdge EBS storage node support for NDMP backups include these features:

- Multiplexing
- Staging
- Coexistence of NDMP save sets with non-NDMP based Sun StorEdge EBS save sets
- Firewall support
- Availability on wide variety of platforms
- Auto media verification
- Backup to disk
- True Sun StorEdge 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 StorEdge EBS storage node support for NDMP clients uses different flags with this program. Table 94 on page 631 lists the save set flags displayed by **mminfo** output and their meanings.

Table 94. 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 StorEdge EBS storage node by using DSA.
Ns	The backup was performed by an NDMP client to a Sun StorEdge EBS storage node by using DSA.

Note: The DSA save set flags displayed by **mminfo** output are not recognized by Sun StorEdge EBS client software earlier than release 7.2.

Performance

An increase in NDMP backup performance can be seen in several instances.

The NDMP tape server is limited to use 60KB 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 to update the Sun StorEdge EBS Console, Sun StorEdge 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 StorEdge EBS storage node.

Vendor Specific Information

For information regarding the limitations, library configuration, and application information variables of each vendor, see Table 95 on page 632.

Table 95. Vendor-Specific Information

Vendor	Refer to These Links
Network Appliance	"Network Appliance" on page 668
EMC Celerra	"EMC Celerra" on page 674
EMC IP4700	"EMC IP4700" on page 678
Auspex	"Auspex" on page 680
Procom NetFORCE	"Procom NetFORCE" on page 683
Mirapoint	"Mirapoint" on page 684
BlueArc	"BlueArc" on page 685

Table 95. Vendor-Specific Information

Vendor	Refer to These Links
DinoStor	"DinoStor TapeServer" on page 686
SnapImage	"SnapImage or NDMP TapeServer Module" on page 688

Configuring the Sun StorEdge EBS Server for NDMP Operations

Before performing an NDMP backup, configure the Sun StorEdge EBS server for NDMP operations. Use the Sun StorEdge EBS Console Administration window to configure the resources listed in Table 96 on page 633 on the Sun StorEdge EBS server before backing up data. Note that certain attributes of the resources are configured differently than for non-NDMP clients. For other configuration issues, see "Additional Considerations for NDMP Operations" on page 646.

Note: Resources not listed in Table 96 on page 633 should be configured the same as for non-NDMP operations.

Table 96. Configuring Resources for NDMP (Part 1 of 2)

Resource	Configuration
Autochanger	Configure a library with the Sun StorEdge EBS Console Administration window. For instructions, see "Configuring Library Support for NDMP" on page 634. For vendor-specific requirements, see "Ancillary Product Information" on page 667.
Device	To use a stand-alone device, create a Device resource for NDMP backups. For instructions, see "How to Configure an NDMP Device Resource" on page 640. To use a library, see "Configuring Library Support for NDMP" on page 634.
Advanced File Type Device	Configure an advanced file device with the Administration window of the Sun StorEdge EBS Console.

Table 96. Configuring Resources for NDMP (Part 2 of 2)

Resource	Configuration
Client	Create a Client resource with specific NDMP attributes. For instructions, see "How to Configure an NDMP Client Resource" on page 642. For vendor-specific requirements, see "Ancillary Product Information" on page 667. 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. **I**f the backup contains more than five million files, set the Inactivity Timeout attribute to zero (0). For additional issues, see "Additional Considerations for NDMP Operations" on page 646.
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.
Pool and Label Template	Specify separate pools for NDMP and non-NDMP volumes. Additionally, create a label template for NDMP operations.

Configuring Library Support for NDMP

Before performing an NDMP backup or recovery, configure the libraries by using the server's Sun StorEdge 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: For instructions on creating a Device resource for a stand-alone device, see "How to Configure an NDMP Device Resource" on page 640.

Obtaining Information about NDMP Autochangers

Before the library can be configured for an NDMP backup and recovery, obtain the pathname of the media device. For instructions on obtaining the device pathname for a particular vendor, see "Ancillary Product Information" on page 667.

Using the NetWorker Administration Window with NDMP Devices

Use the NetWorker Administrator window to search for available devices on the NDMP tape server:

1. From the Console window, click Enterprise.
2. Click the NetWorker server hostname.
3. Double-click NetWorker.

   The NetWorker Administration window opens.
4. Click Devices.
5. Right-mouse 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`
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.

For more information on the **inquire** command, refer to the *NetWorker Command Reference Guide* or the UNIX man page.

How to Configure an NDMP Autochanger

The Sun StorEdge EBS server's Administration interface and the Sun StorEdge 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.

To configure a library with the NetWorker server's Administration interface, see "Chapter 3: Configuring Libraries and Devices" on page 101.

To configure a library with the **jbconfig** tool:

1. Determine the device pathname information for the NDMP library. To obtain this information, see "Ancillary Product Information" on page 667.
2. Open a command-prompt window.
3. Change to the <*Sun StorEdge EBS_install_path*> \ *bin* directory.
4. As root or as a member of the Windows Administrator's group, type this command:

   `jbconfig`

   The use of **jbconfig** requires an NDMP password for the computer with the attached tape device.

5. 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 StorEdge 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: 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 StorEdge 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 StorEdge EBS Console window.

For more information on the **jbconfig** command, refer to the *NetWorker Command Reference Guide* or the UNIX man page. For more information on the NetWorker server's Administration interface, see "Chapter 3: Configuring Libraries and Devices" on page 101.

Sharing Libraries among NDMP Clients

The Sun StorEdge EBS software permits the control and use of individual devices within a library to different Sun StorEdge 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 NetWorker server's Administration interface or the **jbconfig** program.

The **jbconfig** program must be run from the Sun StorEdge 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. Consider these options when using the **jbconfig** program.

- If the Sun StorEdge 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 StorEdge 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**, type this when prompted.

- The Sun StorEdge EBS hostname and device path of the server that controls the first library device.
- The controlling host for *each* device in the library.



Example: Robotic Arm Control

In a three-drive library shared among two NDMP hosts and a Sun StorEdge EBS server, Host1 controls the arm and the first drive.

- When prompted for the owner of Host1, specify the hostname:device path for Host1 as the host controlling the library device.

- When prompted for the owner of Drive2, specify the hostname:device path for Host2 for the owner of Drive 2.
- When prompted for the owner of Drive 3, specify the Sun StorEdge EBS server hostname:device path.

For non-NDMP information about sharing libraries, see "Chapter 3: Configuring Libraries and Devices" on page 101.

DDS Support with NDMP

The Sun StorEdge EBS software includes support for DDS, which provides the capability to recognize shared drives. DDS enables Sun StorEdge 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 StorEdge EBS software, refer to the *EMC Legato Software Compatibility Guide.*

The DDS feature may be enabled when configuring a library through the **jbconfig** program. For more information about DDS Support with NDMP, see "Chapter 11: Device Operations" on page 297.

How to 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 StorEdge 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.

For additional information about standard DDS, see "Chapter 3: Configuring Libraries and Devices" on page 101.

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 -S 1
```

Silo Support

A Sun StorEdge 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 StorEdge EBS server.

How to Configure an NDMP Device Resource

For details on how to configure the Device resource in the standard, non-NDMP Sun StorEdge EBS software, see "Chapter 3: Configuring Libraries and Devices" on page 101. 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 StorEdge EBS Console:

1. Using the Devices tab, under Devices, complete these attributes in the Create Device window.
   a. For the Name attribute, specify the NDMP device by using the following syntax:

      rd=*NDMP_computer:tape_device_pathname*

   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.
2. In the Configuration tab, complete these new device attributes.
   a. For the NDMP attribute, select the checkbox.

      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.
3. Complete any other attributes in the same manner as a non-NDMP Sun StorEdge 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**. For information about the Dedicated Storage Node, see "Tips for Troubleshooting Storage Nodes" on page 108.

How to Configure an NDMP Client Resource

For information on creating a Client resource for non-NDMP operations, see "Task 7: Create a Backup Client Resource" on page 61. 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. Log in as root or as a Microsoft Windows Administrator.
2. From the Administration window, select Clients.
3. Under the General tab, complete these 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 be backed up.

      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. In 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. In the Apps and Modules tab, complete these 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 StorEdge 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 to **nsrndmp_save**.

   This is the syntax for **nsrndmp_save**.

   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* **/savepoint**

   – To back up an NDMP client to a non-NDMP device, use the **-M** option.

   – To backup up to an NDMP device, use no **-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 as follows:

   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. For a list of the variables and rules for typing their values, see "Specifying Application Information" on page 644.

   For more information on the **nsrndmp_save** command, refer to the *NetWorker Command Reference Guide* or the UNIX man page.

6. In the Miscellaneous tab, select Yes for the NDMP attribute.

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 StorEdge 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 StorEdge EBS client configuration.

Specifying Application Information

In the Application Information attribute, list environment variables that the Sun StorEdge 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*. The default is ***n***.

  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. See "Mirapoint" on page 684.

Additional Considerations for HIST=Y

When HIST=n, NDMP backups with Sun StorEdge EBS 7.2 (and also 7.2.1) are always run at level FULL, even with a correct schedule and UPDATE=y (for more information, refer to the "UPDATE=y" bullet below).

A level FULL scheduled backups is performed if there is no previous fully browseable version of the same save set. Furthermore,

- If HIST=y, full and successive level backups perform properly.
- If HIST=y and at least one browseable save set exists, then future level backups of the same save set will backup at the specified level correctly.
- If HIST=n and no previous browseable 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 StorEdge EBS 7.2, 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 StorEdge EBS using **nsrndmp_save -T** ... **-l***1* ..., 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.



Important: The environment variables are case-sensitive. Check with the vendor for the latest information on supported values.

For application information variables for a specific vendor, see the particular vendor section in "Ancillary Product Information" on page 667. 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.

20

Additional Considerations for NDMP Operations

These sections provide information about various Sun StorEdge 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 Sun StorEdge EBS Console, Sun StorEdge EBS User program, or **nwbackup** program results in a full backup.

When the Interval attribute is set to less than 24 hours, the Sun StorEdge 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 higher, or set the Force Incremental attribute to No.

Group Parallelism Setting for NDMP Tape Server

By default, group Parallelism is set to zero, so the Sun StorEdge EBS server uses the value of the client Parallelism. Rather than depending on client Parallelism, set the group Parallelism attribute for NDMP group backups to equal the number of available NDMP drives. If the group Parallelism 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 group Parallelism attributes.

Performing NDMP Backups

This section describes backups by using configurations with the NDMP Tape Server and those with storage node support for NDMP.

Tape Server Support for NDMP Client Backup

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.

These sections describe how to perform an NDMP backup by using the **nsrndmp_save** command, Sun StorEdge EBS User, or the **nwbackup** program. For more information, refer to the *NetWorker Command Reference Guide* or the UNIX man pages.

Prerequisites

Perform a test backup after completion of these actions.

- 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 StorEdge 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.

  Note: To prevent a non-NDMP volume from being selected for an NDMP backup, create separate pools for NDMP and non-NDMP volumes.

- 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 StorEdge 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 only generated if the scheduled group is backed up in an ad hoc manner or through the schedule.

- In addition to the existing configuration support for NDMP clients to back up to an NDMP storage node, a Sun StorEdge EBS storage node can be added to the storage node list of the Client resource configuration to perform backups to non-NDMP devices
- Perform the test backup by running a scheduled group backup.

How to Perform a Local Backup by Using the Sun StorEdge EBS User Program

Use the Sun StorEdge EBS User program to perform a local backup that is not scheduled. A three-party backup *cannot* be performed with the Sun StorEdge EBS User program.

To perform a local backup:

1. In the Sun StorEdge EBS 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 Sun StorEdge EBS User program**, nwrecover**, or the **recover** command.

Note: Configure the client in the Administration window of the Sun StorEdge 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*. For more information on the **nsrndmp_save** command, refer to the *NetWorker Command Reference Guide* or the UNIX man page.

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 StorEdge EBS server and can be accessed through the Sun StorEdge 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 StorEdge EBS server. All environment settings specified for that group, such as clients, pools, and schedules, are used.

How to Perform a Three-Party Backup from the Command-prompt

Three-party backups involve two NDMP servers. For more information, see "Three-Party Backup with NDMP Tape Server" on page 623.

Note: For instructions on performing a remote backup, see "How to Perform a Remote Backup with the SnapImage / NDMP TapeServer Software" on page 689.

To perform a three-party backup:

1. Set up the servers as demonstrated in Figure 47 on page 624.
2. Install the Sun StorEdge EBS software on the server designated as the Sun StorEdge EBS server.

   Note: If the Sun StorEdge 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 StorEdge 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 StorEdge EBS server in the same manner as for any other NDMP backup.

   For instructions on configuring a Sun StorEdge EBS server, see "Configuring the Sun StorEdge EBS Server for NDMP Operations" on page 633.

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 StorEdge 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 StorEdge EBS server and any client without Sun StorEdge EBS software. An example of a DSA is an NDMP host that generates proprietary save data and sends that data to a Sun StorEdge EBS storage device to have a save set associated to it. Similarly, a DSA acts as an agent recover between a NetWorker server and any client without NetWorker 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.

To determine whether DAR is supported with a particular NDMP vendor, refer to the *EMC Legato Software Compatibility Guides.*

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 StorEdge EBS User, or **nwrecover**.



Important: 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 StorEdge EBS non-NDMP storage device (Sun StorEdge EBS storage node).

For more information on the **nsrndmp_save** command, refer to the *NetWorker Command Reference Guide* or the UNIX man page.

Once enabled, the savegroup starts the **nsrndmp_save** program on a specified host. The host could be any true Sun StorEdge EBS client (with or without Sun StorEdge EBS storage node software) or the Sun StorEdge EBS server host itself. The hostname is obtained from the option specified as part of the **nsrndmp_save** command (**-I** *hostname*) and provided as part of the Client resource.

Here is an example.

```
nsrndmp_save –c netapp760a –s hawkeye -I NWClient_host –M –T
dump /savepoint
```

These sections describe how to perform an NDMP backup by using the **nsrndmp_save** command and the **Sun StorEdge EBS User** program. For information on the **nsrndmp_save** command, refer to the *NetWorker Command Reference Guide* or the UNIX man page.

The **nsrndmp_save** program starts the **nsrdsa_save** program (which is part of the Sun StorEdge EBS client software package) to back up data to a Sun StorEdge 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 StorEdge 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 StorEdge 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 StorEdge 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 StorEdge EBS save sets and the volume data format is compatible with non-NDMP based Sun StorEdge 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 StorEdge EBS save instances on the same volume. Any DSA save stream can join a non-NDMP Sun StorEdge EBS save stream. Both can be sent to the same volume.

Performing a Local Backup from a Sun StorEdge EBS Client Computer

Use the Sun StorEdge EBS User program to perform a local backup that is not scheduled.

Note: A three-party backup cannot be performed with the Sun StorEdge EBS User program.

To perform a local backup:

1. In the Sun StorEdge EBS 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.

Note: UNIX users must use the **nwbackup** program.

Performing an NDMP Backup from the Command-prompt

Initiate NDMP backups from the command-prompt by using the **nsrndmp_save** command. With a backup performed from the command-prompt, files or directories cannot be browsed. To browse files or directories that have been backed up, use **nwrecover** (UNIX) or the **recover** command (UNIX).

Note: You must configure the client in the Administration window of the Sun StorEdge 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
-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 StorEdge 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 StorEdge 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.

For additional information on the **nsrndmp_save** command, refer to the *NetWorker Command Reference Guide* or the UNIX man page.

Performing an NDMP Three-Party Backup

Three-party backups involve two NDMP servers. For more information, see "Three-Party Backup with NDMP Tape Server" on page 623.

To perform a three-party backup:

1. Set up the servers as demonstrated in Figure 47 on page 624.
2. Install the Sun StorEdge EBS software on the server designated as the Sun StorEdge EBS server.

   Note: You do not need to install any additional software on the NDMP servers. If, when using a conventional tape or disk device, the Sun StorEdge EBS storage node is part of the configuration, then the bootstrap and file history can be saved conventionally. No additional setup and no explicit definition is needed for an NDMP storage device.
3. Configure the Sun StorEdge EBS server by using the Administration window of the Sun StorEdge EBS Console.
   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 StorEdge EBS server as you would for any other NDMP backup.

For instructions on configuring a Sun StorEdge EBS server, see "Configuring the Sun StorEdge EBS Server for NDMP Operations" on page 633.

Performing NDMP Recoveries

NDMP save set and directory-level recoveries are performed with Sun StorEdge 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 StorEdge EBS User (Windows), the **nsrndmp_recover** command is selected automatically.



Important: 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 StorEdge EBS storage device when **nsrdsa_recover** is started on the same host. The **nsrndmp_recover** program is spawned by these programs.

- **recover**
- **nwrecover**
- **winworkr** (**Sun StorEdge EBS User** graphical interface)
- Sun StorEdge EBS Console

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.



Important: 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 StorEdge 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 Destructive Save Set Recovery

These sections describe how to perform a destructive save set recovery by using **nwrecover** or from the command-prompt. With a destructive save set recovery, data is recovered to the original location and the files are overwritten. For instructions on performing a directed recovery, see "Performing a Directed Recovery" on page 661.

How to Perform a Recovery with the Sun StorEdge EBS User

To perform a destructive save set recovery from the Sun StorEdge EBS 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.
6. 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. For instructions on using this window, see "How to Perform a Directed Save Set Recovery" on page 662.

How to Perform a Save Set Recovery with the Sun StorEdge 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 StorEdge EBS server automatically selects the last save set on the list.
   - If you do not choose a save set instance to recover, the Sun StorEdge EBS server automatically selects a save set instance.
3. When all the files to recover are selected, click Recover.

How to Perform a Save Set Recovery from the Command-prompt

To perform a save set recovery from the command-prompt:

nsrndmp_recover -c *client* **-s** *server* **-S** *ssid/cloneid*

For example, to recover data to a particular drive (Windows) or raw device (UNIX) on a client named *venus*, with a server named *mars*, from the command-prompt, type this command:

```
nsrndmp_recover -c venus -s mars -S 2485094913
```

On a Microsoft Windows system, type this command:

```
nsrndmp_recover -s mars -c venus -r P:\ -S 2674606849
-m P:\
```

For more information on the **nsrndmp_recover** command, refer to the *NetWorker Command Reference Guide* or the UNIX man page.

Performing a Directory-Level Recovery

These sections describe how to perform a directory-level recovery from Sun StorEdge EBS User or from the command-prompt. For information about directed recoveries, see "Performing a Directed Recovery" on page 661.

How to Perform a Directory-Level Recovery from the Sun StorEdge EBS User Program

To perform a directory-level recovery from the Sun StorEdge EBS 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. For instructions on using this window, see "How to Perform a Directed Save Set Recovery" on page 662.

How to Perform a Directory-Level Recovery from the Command-prompt

To perform a directory-level recovery with the **nsrndmp_recover** command, use this syntax.

From a UNIX command-prompt, type this command:

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...**

For information on the **nsrndmp_recover** command, refer to the *NetWorker Command Reference Guide* or the UNIX man page.

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 StorEdge EBS Recover window, run the **nwrecover** program by typing **nwrecover** at the command-prompt.

   You can select a file that is listed in the Sun StorEdge EBS Recover window for recovery. To perform a recovery:

   - Browse the index for previous backups.
   - Select various versions of the file to recover
   - View your marked files
   - Relocate files upon recovery
   - Recover files from another authorized system.

     An authorized system is a client from which you have permission to recover data.

2. If necessary, choose a server by clicking Change Server on the speedbar. If applicable, chose a client by choosing Client from the Change menu.
3. Change the level of file information displayed, if needed.

   - 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. Highlight 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 from the Sun StorEdge EBS User Program

To perform a file-by-file recovery with Sun StorEdge EBS User:

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 mark 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 this command:

   recover -s *server*

2. When prompted, type the directory to browse, such as cd \.
3. Issue the appropriate commands to add the necessary files to the recover list.

   For a complete list of command options, refer to the *NetWorker Command Reference Guide*.

4. When all the required files have been added, type this command:

   recover

The files are recovered.

Note: On a Microsoft Windows a Windows system program named **recover.exe** is also located in *%SystemDrive% \ Winnt \ 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 StorEdge EBS_install_path*> \ *bin* is listed before *%SystemDrive% \ Winnt \ System32* in the PATH environment variable.

Performing a Directed Recovery

Sun StorEdge 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.



Important: 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 StorEdge 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 StorEdge 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 using the Sun StorEdge EBS User program or the **nsrndmp_recover** command.

- When using Sun StorEdge EBS User:
  a. From the Operation menu, select Save Set Recover.
  b. In the Source Client window, select the source client.
  c. In the Save Sets dialog box, select the backup version to recover.
  d. Click Recover Options.

     The NDMP Recover Options window appears.
  e. In the Destination Client text box, specify the client to receive the recovered data.
  f. 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*.
  g. 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.
  h. Click OK.
  i. In the Save Set window, click OK.
- When using **nsrndmp_recover**, type this command:

  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.

For more information on the **nsrndmp_recover** command, refer to the *NetWorker Command Reference Guide* or the UNIX man page.

How to Perform a Directed File-by-File Recovery

You can perform a directed file-by-file (Windows) recovery by using either the Sun StorEdge EBS User program or the **recover** command.

- When using the Sun StorEdge EBS User program:
  a. Click Recover to open the Source Client dialog box.
  b. Select the source client with the data to be recovered, and click OK.
  c. Select the source client with the data to recover, and click OK.
  d. In the left pane of the Recover window, select the save set that contains the files to recover.
  e. In the right pane, select the files to recover.
  f. Click Mark, then Start.
- When using the interactive **recover** program:
  a. From the command-prompt, type this command:

     recover -s *server*
  b. When prompted, type the directory to browse, such as cd \.
  c. Issue the appropriate commands to add the necessary files to the recover list.

     For a complete list of command options, refer to the *NetWorker Command Reference Guide* or the UNIX man page.
  d. 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 backslash (\) after the mount point. Here is an example, *target_hostname*::*mount_point*\.
  e. 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.

- When using the Sun StorEdge EBS recover program:
  a. Select Relocate from the Operations menu.
  b. For the Relocate attribute, use this syntax:

     target_hostname`::/`*mount_point*
  c. Complete the other attributes for a non-NDMP based recovery.
- When using the interactive **recover** program:
  a. From the command-prompt, type this command:

     `recover -s` *server*
  b. When prompted, type the directory to browse, such as cd /.
  c. Issue the appropriate commands to add the necessary files to the recover list.

     For a complete list of command options, refer to the *NetWorker Command Reference Guide*.
  d. 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/*.
  e. When all the required files have been added, type this command:

     `recover`

     The files are recovered to the specified location.

NDMP Cloning

NDMP cloning is set through the same Sun StorEdge EBS clone interface as non-NDMP based backups. For information on how to clone data, see "Chapter 9: Cloning Backups" on page 243.

Important: 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 way of **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 StorEdge 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 **mmlocate(8)**). The target node of a clone is found by the Clone Storage Nodes attribute of the Client resource in descending priority.

Cloning Functionality

This functionality applies to NDMP cloning.

- The Sun StorEdge EBS Console and the command-line interface support NDMP 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 StorEdge 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 in an ad hoc manner. For more information, see "Volume Cloning" on page 246 and "Recovering Cloned Data" on page 257.

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 StorEdge EBS Console through a scheduled group with automatic cloning. For more information on automatic cloning, see "Automatically Clone a Save Set" on page 250.

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 in an ad hoc manner, a minimum of two devices are required, one to read the existing data and one to write the cloned data. For more information, see "How to Manually Clone a Save Set" on page 248.

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 667
- "Filers That Act as Tape Servers Only" on page 686

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 668
- "EMC Celerra" on page 674
- "EMC Centera" on page 678
- "EMC IP4700" on page 678
- "Auspex" on page 680
- "Procom NetFORCE" on page 683
- "Mirapoint" on page 684
- "BlueArc" on page 685

Backup on NDMP servers is dependent on the data server provider type. Servers use the commands listed in Table 97 on page 667.

Table 97. NDMP Server Examples (Part 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.
NetApp	dump	Used for backing up filer data to tape. Also used to back up directories, quota trees, or entire volumes.

Table 97. NDMP Server Examples (Part 2 of 2)

Server	Command	Function
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 NetWorker software. Variables documented in NAS filer reference guides do not necessarily work for Sun StorEdge EBS backups. In these cases, the results may be unpredictable.

For example, the variable FILESYSTEM is supported by NetApp. However, with Sun StorEdge EBS 6.*x*, this setting was ignored and only the save set value was used. With Sun StorEdge 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

These 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 Legato 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 Save Sets

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.

Support for European Filenames

The Sun StorEdge EBS software supports the NDMP backup, recovery (save set, directory, and file-level), and interface display of European filenames by using a NetApp filer and supported Data ONTAP version.

The supported configurations are these ones.

- Nonlocalized Sun StorEdge EBS release 7.1 software for Solaris English 8 (or localized European version, if available) on NetApp filers running Data ONTAP 6.3.*x* and later.
- Nonlocalized Sun StorEdge EBS release 7.1 software for English Windows 2000 on NetApp filers running Data ONTAP 6.3.*x* and higher.

Contact Network Appliance Support to verify which languages each of the various Data ONTAP versions support. Not all Data ONTAP versions support the backup and recovery of all languages. NetApp is compliant with internationalization support. However, the ability of a particular Data ONTAP version to back up and recover the local language settings must be verified.

These sections contain important notes, requirements, and recommendations for supporting localized data with NDMP.

UTF-8 Versions Supported

The Sun StorEdge 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.

- If the NetApp filer's volume is configured for UTF-8 encoding / code set, the Sun StorEdge EBS server must also have the UTF-8 variable set to Yes in the Application Information attribute of the NDMP Client resource.
- If the filer's volume is not configured for UTF-8 encoding / code set, the Sun StorEdge EBS server must also have the UTF-8 variable set to No.

To support nonASCII characters when using NDMP, the UTF-8 settings on the filer and the Sun StorEdge EBS server must match.

To set the UTF-8 variable, perform one of these steps.

- If the data from the filer is *not* in UTF-8 format, type the following or leave the value blank.

  UTF8=n

  The default is ***n***.
- If the data from the filer is in UTF-8 format, type the following:

  UTF8=y

  The default is ***n***.

Single Save Sets Support Only One Code Set

A single save set supports data belonging 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 this command:

   `ndmpd off`

   The NetApp host confirms that the NDMP process has stopped.
3. To change the NDMP version, type this command:

   `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 this command:

   `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 634.

Important: The procedure that follows 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 2005
      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.

Note: 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.

Important: 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 98 on page 673 lists variables specific to NetApp that may be required. Type these variables in the Application Information attribute of the Client resource, as described in "How to Configure an NDMP Client Resource" on page 642.

Table 98. NetApp Application Information Variables (Part 1 of 2)

Variable	Required?	Values and Meaning
FILESYSTEM= *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. For additional information about DAR, see "Performing DAR Backups" on page 650.

Table 98. NetApp Application Information Variables (Part 2 of 2)

Variable	Required?	Values and Meaning
EXCLUDE=	No	This string value determines which files are excluded from backup. These rules apply. • The string is a filename. • 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. • 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=y / n	Yes for UTF-8 formatted data	Specifies whether the data on the filer is in UTF-8 format. To support nonASCII characters when using NDMP, the UTF-8 settings on the filer and the Sun StorEdge EBS server must match. If the data from the filer is not in UTF-8 format, set the value to **n**. If the data from the filer is in UTF-8 format, set the value to **y**. The default is **n**.

EMC Celerra

These sections provide information about limitations, autochanger configuration, and application information variables for the EMC Celerra host. For a complete list of supported versions, refer to the *EMC Legato Software Compatibility Guide.*

EMC Celerra Considerations with NDMP

These sections discuss NDMP issues that are specific to EMC Celerra filers.

EMC Celerra Backup of Root File System

Instead of specifying a save set of All for an EMC 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 EMC Celerra Data Movers

EMC 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, EMC 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.

How to Lower the NDMP Version on EMC Celerra

The NDMP version must be set individually on each data mover that requires the modification.

To set the NDMP version on EMC 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 EMC Celerra Control Station.
2. At the command-prompt, type this command:

   `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. Edit the *netd* file by typing this command.

   `vi netd`
4. At the end of the NDMP string in the *netd* file, add this command.

   `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 EMC 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

   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 an EMC 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 634.

Important: The procedure that follows is only an example. For the correct syntax and usage, refer to the EMC Celerra documentation.

To identify the device pathname, perform *one* of these actions.

- Re-create all devices and view all nondisks by typing this at the command-prompt.

  Celerra **server_devconfig** *movername* **-create -scsi -all**

  Celerra **server_devconfig** *movername* **-list -scsi -nondisks**

- Obtain the device pathname through the EMC Celerra Administrator program.

  To access the EMC Celerra Administrator program:

  a. In the browser window, type this URL.

     http://EMCcontrolhostIP:8000/top_level.htm.

  b. Type your username and password.

  c. Use the EMC 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 EMC 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 99 on page 678 lists Celerra-specific variables that may be required. Type these variables in the Application Information attribute of the Client resource, as described in "How to Configure an NDMP Client Resource" on page 642.

Table 99. EMC 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**. For additional information about DAR, see "Performing DAR Backups" on page 650.
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.

EMC Centera

For information about the EMC Centera™ Backup and Recovery Module Software, refer to the *Centera Backup and Recovery Module Release 2.01 Installation and Administrator's Guide* and the *Centera Backup and Recovery Module Release Supplement*. For a complete list of supported versions, refer to the *EMC Legato Software Compatibility Guide.*

EMC IP4700

These sections provide information about autochanger configuration and application information variables for the EMC IP4700 host. For a complete list of supported versions, refer to the *EMC Legato Software Compatibility Guide.*

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 634.

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 StorEdge 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 100 on page 679 lists IP4700-specific variables that may be required. Type these variables in the Application Information attribute of the Client resource, as described in "How to Configure an NDMP Client Resource" on page 642.

Table 100. EMC IP4700 Variables (Part 1 of 2)

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**.

Table 100. EMC IP4700 Variables (Part 2 of 2)

Variable	Required	Values and Meaning
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 only, 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. For a complete list of supported versions, refer to the *EMC Legato Software Compatibility Guide.*

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. For more information about snapshots, refer to the appropriate Auspex documentation.

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.

You can now back up large files.

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 higher.

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 Sun StorEdge EBS 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 634.

For information about adding a SCSI device, refer to the Auspex documentation.

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 101 on page 683 lists Auspex-specific variables that may be required. Type these variables in the Application Information attribute of the Client resource, as described in "How to Configure an NDMP Client Resource" on page 642.

Table 101. 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**. For additional information about DAR, see “"Performing DAR Backups" on page 650.
SNAPSHOT= y	No	Enables or disables snapshots. The default is **n**.

Procom NetFORCE

These sections provide information about autochanger configuration for the Procom host. For a complete list of supported versions, refer to the *EMC Legato Software Compatibility Guide.*

How to Find the Device Pathname for a Procom NetFORCE Host

To determine the device pathname, log in as root and type this 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 634.

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 *isp1tSSL[L]*, where *isp1* is the autochanger handle. The Fibre Channel device format is *ffx1tSSL[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 this command.

   ```
   /usr/sbin/nsrjb -HE
   ```

2. To read previously-labeled tapes back into the database, you must inventory the autochanger.

   ```
   /usr/sbin/nsrjb -IE
   ```

Mirapoint

These sections provide information about requirements and autochanger configuration for the Mirapoint host. For a complete list of supported versions, refer to the *EMC Legato Software Compatibility Guide*.

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/nrstn*, 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 634.

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 StorEdge 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 this one may appear.

  `NDMP` session—`Unknown environment variable` *name* `ignored.`

  You can ignore this message, as it does not impact the operation.

For complete information on configuring a BlueArc system for NDMP operations, refer to *BlueArc's* documentation.

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. Combined with the Sun StorEdge EBS software, the NetWorker SnapImage Module provides a complete, reliable, and efficient storage management solution.

For instructions on installing, configuring, and using the SnapImage software, refer to the SnapImage Module Installation and Administrator's Guides and Release Notes.

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. Note the following 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 NetWorker Administration window, the DinoStor TapeServer has one hostname with multiple device handles. Use the hostname to configure the autochanger with the NetWorker Administration window, as described in "Configuring Library Support for NDMP" on page 634.

To specify the device handles, perform these actions:

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 StorEdge EBS server, and the DinoStor TapeServer with a library attached.
2. Configure the Sun StorEdge 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 StorEdge EBS server as you would for any other NDMP backup.

For instructions on configuring the Client resource, see "How to Configure an NDMP Client Resource" on page 642.

20

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 StorEdge EBS Console Administration window, as described in "Configuring Library Support for NDMP" on page 634.

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 this 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.

Note: If the NDMP tape server is running version 2 or does not support NDMP_CONFIG interfaces, determine the autochanger handle by typing this 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, ensuring that you are logged in as root on the host, use *one* of these methods.

- Use the **ioscan** command by typing this command:

  `ioscan -fnC tape`

- Use the SAM utility.
  a. Start the utility.
  b. Select Peripheral Devices.
  c. Select Tape Drives.
  d. Highlight the appropriate tape drive and select Show Device Files from 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 StorEdge EBS software on the server designated as the Sun StorEdge EBS server.
2. Install the SnapImage Module / NDMP TapeServer software on the server designated as a tape server.
   - If you are backing up the data to the Sun StorEdge EBS server, install the SnapImage Module / NDMP TapeServer software on that Sun StorEdge EBS server.
   - 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.

   For instructions on installing the SnapImage Module / NDMP TapeServer software, refer to the NetWorker SnapImage Module Installation and Administrator's Guide and Release Notes.

   Note this restriction: 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 StorEdge 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 StorEdge EBS server as you would for any other NDMP backup.

      For instructions on configuring the Client resource, see "How to Configure an NDMP Client Resource" on page 642.

Chapter 21: SNMP Module

This chapter covers the following topics:

- "SNMP Traps" on page 691
- "Configuring StorEdge EBS SNMP Notifications" on page 692
- "Configuring SNMP Management Software" on page 695

The Sun StorEdge EBS Simple Network Management Protocol (SNMP) Module allows Sun StorEdge EBS servers to send notification messages to SNMP management agents.

SNMP-enabled network management software must be configured to accept traps from the Sun StorEdge EBS server. For detailed information about SNMP management operations refer to your network management documentation.

SNMP Traps

The StorEdge EBS SNMP Module uses *traps* to communicate Sun StorEdge EBS event notifications to SNMP management stations. A trap is an unsolicited notification sent from the SNMP agent (such as the Sun StorEdge EBS server) to the SNMP event manager.

The types of traps that the Sun StorEdge EBS server sends are determined when the StorEdge EBS SNMP notification is configured within the Sun StorEdge EBS server. Typical traps include warnings, critical errors, and other messages from the Sun StorEdge EBS server. For instructions on configuring the StorEdge EBS SNMP notification, see "Configuring StorEdge EBS SNMP Notifications" on page 692.

Configuring StorEdge EBS SNMP Notifications

Sun StorEdge EBS software provides notifications to a variety of resources about Sun StorEdge EBS server events. The StorEdge EBS 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 optional **nsrtrap** command-line switches, such as the SNMP community and the trap type.

In order to configure the StorEdge EBS SNMP notification, you must first enable the StorEdge EBS SNMP module. For information about enabling the StorEdge EBS SNMP module, see the appropriate *Sun StorEdge EBS Installation Guide*.

Command-Line Options for nsrtrap

The StorEdge EBS SNMP Module uses the **nsrtrap** program to communicate SNMP traps from the Sun StorEdge EBS server to the SNMP management software. Table 102 on page 692 lists the command-line options that can be included in the Action attribute when the SNMP notification is configured.

Table 102. Command-Line Options for nsrtrap (Part 1 of 2)

Option	Description
-c *community*	Specifies the SNMP community that is authorized to receive traps from the Sun StorEdge 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 StorEdge 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 StorEdge EBS 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 StorEdge 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 StorEdge EBS notification.

Table 102. Command-Line Options for nsrtrap (Part 2 of 2)

Option	Description
-s *specific_type*	A generic setting that can be used to identify the type of trap the Sun StorEdge 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 StorEdge 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 StorEdge 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 StorEdge 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.

How to Modify the Preconfigured StorEdge EBS SNMP Notification

The Sun StorEdge 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. For more information about command-line options, see "Command-Line Options for nsrtrap" on page 692.

5. Click OK.

   The events and priorities associated with the preconfigured SNMP notification cannot be modified. To set different events and priorities for the SNMP notification, see "How to Create StorEdge EBS SNMP Notifications" on page 694.

How to Create StorEdge EBS SNMP Notifications

To create additional StorEdge EBS 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, enter the following:

   ```
   /usr/sbin/nsr/nsrstrap 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. For more information about command-line options, see "Command-Line Options for nsrtrap" on page 692.

8. Click OK.

Configuring SNMP Management Software

In order for the SNMP management software to accept traps sent by Sun StorEdge 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 StorEdge 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 StorEdge 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 StorEdge EBS servers, an icon of each Sun StorEdge EBS server appears on the network management console. The following are examples of how the software can then be configured:

- To indicate that a trap was received (for example, the Sun StorEdge 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 StorEdge 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. For more information about setting the **-s** *specific-type* option, see "Command-Line Options for nsrtrap" on page 692.

Chapter 22: Support for Microsoft Removable Storage

This chapter covers the following topics:

- "Microsoft Removable Storage" on page 697
- "Types of Support for Removable Storage" on page 699
- "Removable Storage Media Pools" on page 702
- "Usage Scenarios" on page 704
- "Configuring the Sun StorEdge EBS Server for Removable Storage" on page 706
- "Library Operations with Removable Storage" on page 710
- "Device Cleaning" on page 720
- "Considerations for Uninstalling the Sun StorEdge EBS Software" on page 720

Note: The procedures in this chapter assume that you have administrator privileges on the Sun StorEdge EBS server or storage node host computer for which removable storage is being configured.

Microsoft Removable Storage

Microsoft Removable Storage is a service that is installed as a standard component of Windows 2000 and Windows Server 2003.

Removable Storage acts as an agent to the Sun StorEdge EBS software (and other data management applications) to perform the actual mounting, labeling, and tracking of media. Removable Storage controls the data storage hardware's doors, drives, and robotics, and performs uniform drive cleaning operations. Removable Storage does not provide volume or file management. Those tasks are performed by data management applications.

22

A primary advantage in using Removable Storage is that it allows multiple data management applications to share access to the same storage hardware.

Removable Storage Libraries and Removable Storage Systems

A Removable Storage system is one or more libraries under the control of the Removable Storage service on a single computer.

The essential components of a Removable Storage library are data storage media (tape or disk volumes) and one or more drives used to read and write data on the volumes. Although Microsoft Removable Storage supports both autochanger libraries and stand-alone drives, Sun StorEdge EBS software supports use of Removable Storage for autochanger libraries only. Stand-alone drives are controlled directly, without Removable Storage.

Note: In this chapter, the term *library* refers to a Removable Storage library. The term *autochanger* refers to a variety of robotic backup devices, including autoloaders, carousels, datawheels, jukeboxes, and near-line storage.

The Removable Storage Database

The Removable Storage database (also called the Windows NT Media Services database) tracks configuration information about the objects in a Removable Storage system.

The Removable Storage service updates its database whenever a configuration or status change occurs within a Removable Storage system. For example, if a volume is mounted in a Removable Storage system, the Removable Storage service commands the library hardware to mount the volume, then updates its database to reflect the change in the volume's location and status.

Database Backup and Recovery

The Removable Storage database files, stored in *%SystemRoot% \ system32 \ NtmsData*, are backed-up and recovered as a part of the Sun StorEdge EBS SYSTEM DB or VSS SYSTEM SERVICES save set. For details about backup and recovery of SYSTEM and VSS SYSTEM save sets, see "Chapter 13: Recovering Data" on page 367.

The Removable Storage service must be running during a Removable Storage database recovery. Before beginning a SYSTEM DB or VSS SYSTEM SERVICES save set recovery, verify that the Removable Storage service is running.

After the recovery, Windows requires a system reboot. Windows does not allow more than one recovery of the Removable Storage database during the same session. An additional recovery of this database requires a server reboot.

Types of Support for Removable Storage

The Sun StorEdge EBS software interacts with Removable Storage systems through an application programming interface (API). Support for this API is installed as a standard component during Sun StorEdge EBS server or storage node setup.

Use of Removable Storage is optional. You can disable Removable Storage, for example, so that Sun StorEdge EBS software can have exclusive use of a library. If this is the case, disregard this chapter and follow the procedures in "Chapter 11: Device Operations" on page 297.

Licensing

Sun StorEdge EBS licensing for devices enabled for Removable Storage is available on a per-slot basis. Purchase a license for the number of logical slots required for the Sun StorEdge EBS volumes. For information about logical slots, see "Logical Slot Numbers versus Physical Slot Locations" on page 711.

Operating System Support

Table 103 on page 699 shows Sun StorEdge EBS Removable Storage support for each Windows operating system that Sun StorEdge EBS servers or storage nodes support.

Table 103. Removable Storage Support by Operating System

Operating System	Sun StorEdge EBS Removable Storage Support
Windows Server 2003	Full Removable Storage support
Windows 2000	Full Removable Storage support
Windows NT 4.0	Sun StorEdge EBS Administrator access to Windows Server 2003 or Windows 2000 Sun StorEdge EBS servers only. (Removable Storage is not available on Windows NT 4.0.)

22

Note: Interaction between Removable Storage and Sun StorEdge EBS servers or clients on UNIX is not supported.

Hardware Support

Table 104 on page 700 lists the storage library hardware Sun StorEdge EBS software supports for use with Removable Storage.

Table 104. Library Hardware Supported for Removable Storage

Format	Library Hardware
4 mm	HP SureStore Galactica, fully configured
	Sony TSL-S11000
8 mm	Exabyte 430M
	Exabyte 8900 Mammoth / 2
8 mm AIT-1	Qualstar TLS series (4210, 4210A, 4212 FCO, 4220, 4420 FCO)
8 mm AIT-2	ADIC Scalar 100 with large port
	HP StorageWorks SSL2020
DLT	ADIC DS9000 series (DS9200D, DS9300D, DS9400D, DS9700D, DS9800D)
	ADIC Scalar 218
	Dell PowerVault 110T series (DLT1, DLT4000, DLT7000)
	ODI Neo Series 2000 or 4000
SuperDLT	Dell PowerVault 110T-SDLT
	ODI Neo Series 2000 or 4000
LTO	Dell PowerVault 110T-LTO

Limitations

Release 7.*x* of the Sun StorEdge EBS software does not support the following features for use with Removable Storage systems:

- Interaction with Sun StorEdge EBS servers or clients that are running on UNIX
- DDS
- NDMP
- Sun StorEdge EBS device autodetection

  Use **jbconfig** instead. For more information, see "Configuring the Sun StorEdge EBS Server for Removable Storage" on page 706.
- Sun StorEdge EBS autoclean

  Use Removable Storage cleaning instead. For more information about cleaning devices, see "Device Cleaning" on page 720.
- The following Sun StorEdge EBS commands are not supported with Removable Storage systems:
  - Any **sji*** command, for example, **sjimm** and **sjirelem**.

    For a complete list of Sun StorEdge EBS **sji** commands, refer to the *NetWorker Command Reference Guide*.
  - The **mt** command
  - The **inquire** command
  - The **nsrjb -d** and **nsrjb -w** commands, for Cartridge Access Port (CAP) deposit and withdraw

Note: If your library has a CAP, mail slot, mailbox, or import/export element, the Sun StorEdge EBS **nsrjb -d** and **nsrjb -w** commands cannot be used.

22

Removable Storage Media Pools

The Removable Storage service organizes media into hierarchical groupings called *media pools*. Each volume in a Removable Storage system must belong to a single, type-specific pool, which defines the properties used to manage that media type. There are two general categories of Removable Storage media pools: system media pools and application media pools.

Note: Do not confuse Removable Storage media pools with Sun StorEdge EBS pools. Microsoft Removable Storage uses its own set of pools for organizing media. In this chapter, the term *pool* or *media pool* refers to a Removable Storage media pool, unless otherwise specified. Continue to use Sun StorEdge EBS pools to sort and manage save sets in a Removable Storage system, but at a lower level For more information about Sun StorEdge EBS pools, see "Chapter 7: Sorting Backup Data" on page 203.

A Removable Storage media pool can hold only one type of media. For example, if the Sun StorEdge EBS server has a 4-mm DDS library, only pools defined exclusively for 4-mm DDS media can contain that library's volumes.

As a data management application, Sun StorEdge EBS software interacts with Removable Storage to access media and, when necessary, commands Removable Storage to move media from one pool to another.

System Media Pools

System media pools are the containers Removable Storage uses for unallocated media. At the root of the media pool hierarchy, Removable Storage creates the following types of system media pools and manages the media in them to provide the amount of data storage each application needs:

- *Free* pools — Contain media that is available to any application. Removable Storage labels each volume in a free pool with a *free media label*, which overwrites any existing label and destroys any data on the volume.
- *Import* pools — Contain volumes that are not currently allocated, but have a label that Removable Storage is able to recognize by referring to the appropriate media label library (MLL). For more information about the Sun StorEdge EBS MLL, see "Media Label Library" on page 712.
- *Unrecognized* pools — Contains new (blank) media, or media with a label Removable Storage does not recognize. Media cannot be allocated directly from an unrecognized pool. It must first be moved to a free or import pool.

Each system media pool contains one pool for each media type that is present in the Removable Storage system. For example, if the Sun StorEdge EBS server has a 4-mm DDS library and a DLT library, the free, import, and unrecognized pools each contain type-specific pools for 4-mm DDS media and DLT media.

Application Media Pools

Application media pools contain volumes that are allocated for the exclusive use of a particular data management application. Multiple applications can create and maintain their respective pools within a Removable Storage system. The application can then share access to the library hardware while protecting their own data volumes.

Sun StorEdge EBS Media Pools

The Sun StorEdge EBS server creates the following application media pools:

- The Sun StorEdge EBS pool is created in the media pool *root*. It is the top-level container for all Sun StorEdge EBS media pools.
- In the Sun StorEdge EBS pool, a type-specific pool is created for each media type that is allocated to the Sun StorEdge EBS software. These pools contain all of the volumes of a particular media type, even if the volumes reside in different libraries.

For more information about how these pools are used, see "Sun StorEdge EBS Media Allocation and Deallocation" on page 712.

Viewing Media Pools

Windows provides a Microsoft Management Console (MMC) snap-in for viewing and administering Removable Storage systems.

To view the media pool hierarchy and the volumes each pool contains:

1. On the Windows desktop, right-click My Computer and select Manage.
2. Expand *Storage* \ *Removable Storage* \ *Media Pools*.



Important: Do not use the Removable Storage MMC snap-in to manage Sun StorEdge EBS volumes except when instructed to do so in this chapter.

22

Usage Scenarios

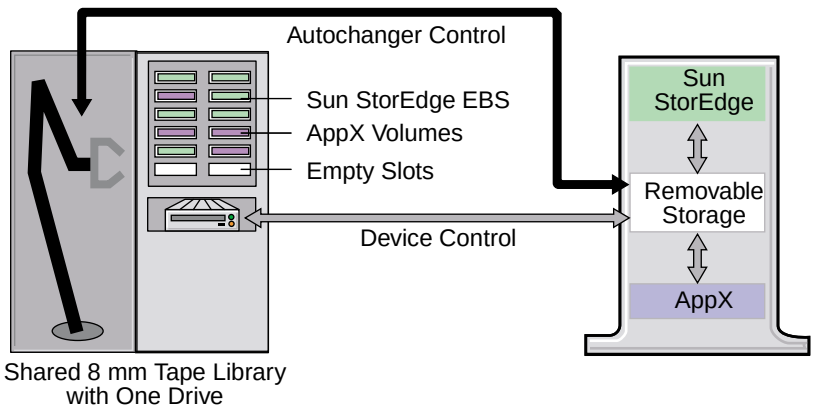
The following examples illustrate a few of the many configuration options available for using Removable Storage with Sun StorEdge EBS software.

AppX represents any data management application that supports Removable Storage (for example, Windows NT Backup or Backup Exec).

Example: One Shared Library with One Drive

In Figure 48 on page 704, the Sun StorEdge EBS server and AppX are installed on a computer that has one 8 mm tape library with one drive. Removable Storage allocates media to each application from the shared library.

Figure 48. Usage Scenario: One Shared Library



In this example, the Sun StorEdge EBS application media pool hierarchy might look like this:

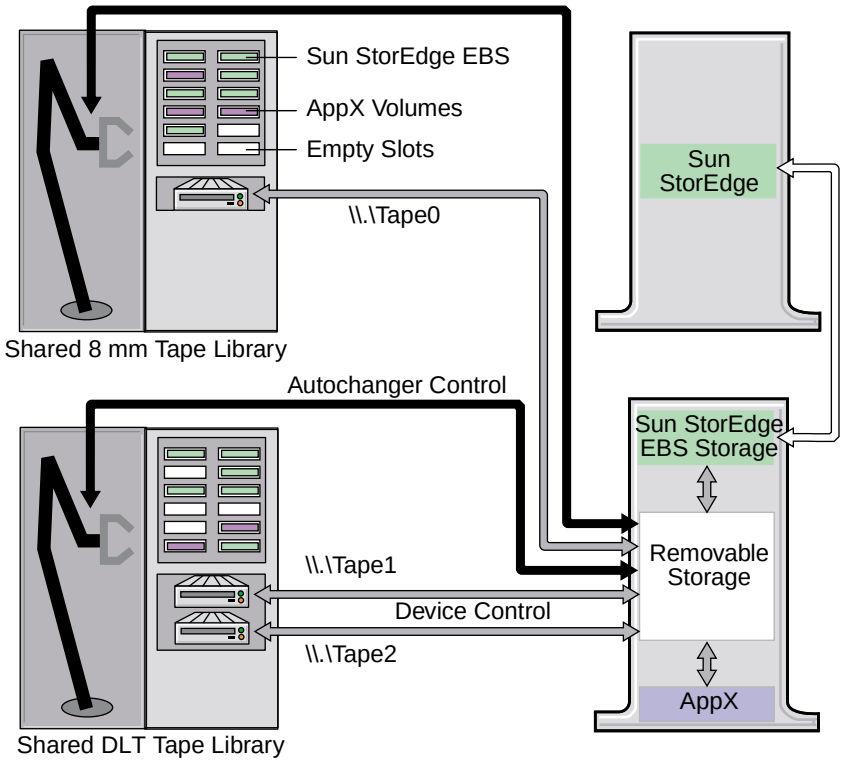
```
Sun StorEdge EBS
    8 mm
    nw_volume.001   nw_volume.004
    nw_volume.002   nw_volume.005
    nw_volume.003   nw_volume.006
```

The Sun StorEdge EBS server and AppX compete for access to the drive. Each must wait for the other to finish before Removable Storage can mount its media.

Example: Two Shared Libraries of Different Media Types

In Figure 49 on page 705, a storage node and AppX are installed on a computer that has an 8 mm library with one drive and a DLT library with two drives.

Figure 49. Usage Scenario: Two Shared Libraries



The Sun StorEdge EBS storage node and AppX can each allocate media from both libraries, and they compete for access to all drives.

In this example, the contents of the Sun StorEdge EBS application media pools might look like this:

```
Sun StorEdge EBS
    8 mm
    AccountingFull.001        AccountingFull.004
    AccountingFull.002        AccountingFull.005
    AccountingFull.003        AccountingFull.006
```

```
DLT
AccountingFull_c.001      AccountingFull_c.004
AccountingFull_c.002      AccountingFull_c.005
AccountingFull_c.003      AccountingFull_c.006
```

The Sun StorEdge EBS pools might be configured like this:

```
Sun StorEdge EBS pool 1:
   Name:     Accounting Full
   Type:     Backup
   Group:    Accounting
   Level:    Full
   Devices:  \\.\Tape0

Sun StorEdge EBS pool 2:
   Name:     Accounting Full Clone
   Type:     Backup Clone
   Group:    Accounting
   Level:    Full
   Devices:  \\.\Tape1
             \\.\Tape2
```

Sun StorEdge EBS *pool 1* directs all save streams to the device associated with the 8 mm library, while Sun StorEdge EBS *pool 2* directs the clone save streams to the two devices associated with the DLT library.

For information about working with Sun StorEdge EBS pools, see "Chapter 7: Sorting Backup Data" on page 203.

Configuring the Sun StorEdge EBS Server for Removable Storage

Configuring a Sun StorEdge EBS server or storage node to use a Removable Storage system, is a three step process:

- "Step 1: Enable or Disable Removable Storage" on page 706
- "Step 2: Monitor the Removable Storage Work Queue" on page 707
- "Step 3: Create Removable Storage Resources" on page 708

Step 1: Enable or Disable Removable Storage

Disable Removable Storage on each library that is connected to a Sun StorEdge EBS server or storage node to create a SCSI jukebox resource. However, you must re-enable Removable Storage feature before you can configure the

software to use it as a Removable Storage jukebox resource. If the Sun StorEdge EBS server or storage node has a library that should not be configured for Removable Storage, ensure that Removable Storage is disabled on that library.



Important: Removable Storage can be enabled or disabled at three levels: library, drive, and volume. To use Removable Storage, ensure that library, drive, and volume are all enabled. By default, it is enabled at all three levels.

How to Enable or Disable Removable Storage

To enable or disable Removable Storage:

1. On the Windows desktop, right-click My Computer and select Manage.
2. In the left pane of the Computer Management window, expand *Storage \ Removable Storage \ Physical Locations*.
3. Right-click the library name and select Properties.
4. Perform one of the following in the library's Properties dialog box (under the General tab):
   - If Sun StorEdge EBS software is to use Removable Storage, select Enable library.
   - If Sun StorEdge EBS software is *not* to use Removable Storage, clear Enable Library.
5. Repeat steps 3 and 4 for each library connected to the Sun StorEdge EBS server or storage node host computer.

Step 2: Monitor the Removable Storage Work Queue

Removable Storage can only perform one task at a time on a library system. While the storage library hardware is busy carrying out a task, other requests for Removable Storage operations are sent to the work queue. Some operations, such as attempting to run the Sun StorEdge EBS **jbconfig** command, can cause the following error message to appear:

```
Could not start Removable Storage session
```

How to Monitor the Removable Storage Work Queue

To monitor the Removable Storage work queue:

1. On the Windows desktop, right-click My Computer and select Manage.
2. In the Computer Management window, expand *Storage \ Removable Storage*, and then select the Work Queue icon.
3. Queued tasks are listed in the Computer Management window. For more information about monitoring and managing the work queue, refer to the Microsoft Removable Storage documentation.

Step 3: Create Removable Storage Resources

Sun StorEdge EBS software communicates with a Removable Storage system through a resource called a Microsoft Removable Storage jukebox. To create this resource and one or more associated Device resources, run **jbconfig** from the command-prompt. You cannot use the Sun StorEdge EBS autodetection feature to create these resources.

Create one Microsoft Removable Storage jukebox resource for each Removable Storage-enabled library connected to the Sun StorEdge EBS server or storage node. If you have multiple libraries of the same media type, all volumes of that type are managed in the same application media pool.

Preliminary Requirements

Before you run **jbconfig**, perform the following:

1. If you have more than one library connected to the Sun StorEdge EBS server or storage node host computer, determine the ID of each SCSI library by running **inquire** at the command-prompt.

Note: The **inquire** command cannot be run with Removable Storage enabled. For instructions on how to disable Removable Storage, see "Step 1: Enable or Disable Removable Storage" on page 706.

The following example shows two HP libraries, each with one tape drive:

```
scsidev@2.5.0:HP  C1557A  U812|Tape, \\.\Tape0
scsidev@2.5.1:HP  C1557A  U812|Autochanger (Jukebox), \\.\Changer0
scsidev@2.6.0:HP  C1557A  U709|Tape, \\.\Tape1
scsidev@2.6.1:HP  C1557A  U709|Autochanger (Jukebox), \\.\Changer1
```

2. Ensure that the name to be used is not already in use. For information about checking library names, "Chapter 3: Configuring Libraries and Devices" on page 101.

Creating a Microsoft Removable Storage Jukebox Resource

If there are no libraries with Removable Storage enabled on the system, **jbconfig** displays the following error message and then quits:

```
There are no enabled Removable Storage libraries to be
configured.
```

Enable any libraries you intend to configure for Removable Storage. For more information, see "Step 1: Enable or Disable Removable Storage" on page 706.

How to Create a Removable Storage Jukebox on a Sun StorEdge EBS Server

To create a Microsoft Removable Storage jukebox on a Sun StorEdge EBS server:

1. Check the Removable Storage work queue to verify that the library you want to configure is not busy. For information on monitoring the work queue, see "Step 2: Monitor the Removable Storage Work Queue" on page 707.
2. Start the Windows command-prompt.
3. Enter the following command:

   jbconfig

4. Respond to the **jbconfig** prompts as appropriate for your system. Examples of nondefault responses here are bold. Press Enter to accept the default response for any prompt.

   ```
   1) Configure an AlphaStor Library.
   2) Configure an Autodetected SCSI Jukebox.
   3) Configure an Autodetected NDMP SCSI Jukebox.
   4) Configure an SJI Jukebox.
   5) Configure an STL Silo.
   6) Configure a Microsoft Removable Storage Jukebox.
   What kind of Jukebox are you configuring? [1] 6
   ```

How to Create a Removable Storage Jukebox on a Storage Node

To create a Microsoft Removable Storage jukebox on a Sun StorEdge EBS storage node:

1. Check the Removable Storage work queue to verify that the library you want to configure is not busy. For information on monitoring the work queue, see "Step 2: Monitor the Removable Storage Work Queue" on page 707.
2. Start the Windows command-prompt.
3. Enter the following command:

   `jbconfig -s server`

4. Respond to the **jbconfig** prompts as appropriate. Press Enter to accept a default answer. Examples of nondefault responses here are bold.

```
On a storage node, the hostname is a prefix to the jukebox name.
Enter the hostname to use as a prefix? [squid.legato.com]
using 'squid.legato.com' as the hostname prefix
    1) Configure an AlphaStor Library.
    2) Configure an Autodetected SCSI Jukebox.
    3) Configure an Autodetected NDMP SCSI Jukebox.
    4) Configure an SJI Jukebox.
    5) Configure an STL Silo.
    6) Configure a Microsoft Removable Storage Jukebox.
What kind of Jukebox are you configuring? [1] 6
```

Library Operations with Removable Storage



Important: Do not use the Removable Storage MMC snap-in, the Removable Storage Manager (RSM) command-line interface, or the Sun StorEdge EBS **nsrjb** command at the command-prompt to allocate or deallocate Sun StorEdge EBS volumes. Doing so causes the Sun StorEdge EBS media database to become out of sync with the Removable Storage database.

Library Operations



Important: To allow storage device sharing between Sun StorEdge EBS software and other data management applications, do not manually mount volumes in a Removable Storage system. The Sun StorEdge EBS software automatically mounts and dismounts volumes as necessary. Volumes are dismounted after 10 minutes of device idle time.

To access Library Operations:

1. In the left pane of the Libraries window, select the server.
2. In the right pane, select the device.
3. From the Libraries menu, select Administer Library.

   The Library window appears.
4. From the middle pane of the Library window, select the device.
5. From the Devices menu, select Library Operations.
6. Choose the required operation from the list.

Logical Slot Numbers versus Physical Slot Locations

For a library with Removable Storage disabled, the slot numbers in the Libraries window represent the physical location of each volume. For a library with Removable Storage enabled, the numbers shown are *logical slot numbers*, and do not indicate the volume's physical location within the library.

Note: Logical slot number assignments can change as you add or remove volumes. For example, if Sun StorEdge EBS software has two volumes allocated and you add a third, the volume that was assigned to slot 2 might get pushed down to slot 3, and the new volume might be assigned to slot 2.

How to Determine a Volume's Physical Location

To determine the physical location of a volume:

1. Use the Sun StorEdge EBS Management Console to determine the logical slot numbers of the volume:
   a. In the left pane of the Libraries window, select the server.
   b. In the right pane, select the device.
   c. From the file menu, select Libraries>Administer Library. The Library window appears.
   d. From the middle pane of the Library window, select the device.

      The logical slot numbers are listed in the far right pane of the window.
2. Use the Windows Device Manager program to determine the physical slot numbers of a volume. For example:
   a. On the Windows desktop, right-click My Computer and select Manage>Device Manager. The Computer Management window appears.
   b. In the Computer Management window, expand *Storage \ Removable Storage \ Physical Locations \ library_name \ Media*. The Location column displays the volume's physical slot location.

Sun StorEdge EBS Media Allocation and Deallocation

The Sun StorEdge EBS software interacts with Removable Storage to allocate and deallocate media according to predefined policies.

Media Label Library

A Media Label Library (MLL) is a software component that interprets the format of a particular application's volume labels and provides that information to the Removable Storage service.

The Sun StorEdge EBS MLL, contained in a single file named *mll_lgto.dll*, is installed in *%SystemRoot% \ System32* and registered automatically during Sun StorEdge EBS server or storage node setup. Removable Storage uses the Sun StorEdge EBS MLL to interpret Sun StorEdge EBS volume labels and determine the Sun StorEdge EBS media allocation and deallocation policies.

Media Allocation

When you select Add in Library Operations, the Sun StorEdge EBS software commands Removable Storage to search the media pools in a predefined order and allocate the first appropriate volume it finds.

For example, if the library uses 4-mm DDS tapes, the Sun StorEdge EBS software searches for media pools in the following order:

1. *Removable Storage* \ *Media Pools* \ *Legato* Sun StorEdge EBS \ *4mm DDS* is searched for Sun StorEdge EBS volumes in that library that were previously allocated, but are now deallocated.
2. *Removable Storage* \ *Media Pools* \ *Import* \ *4mm DDS* is searched for Sun StorEdge EBS volumes in that library that have been imported, but have not yet been allocated.
3. *Removable Storage* \ *Media Pools* \ *Free* \ *4mm DDS* is searched for a new or recycled volume in that library.

If an available volume is not found in the library, an informational message appears.

Media Deallocation

When you select Remove in Library Operations, the Sun StorEdge EBS software does the following:

- Deletes the selected volume from its media database.
- Commands Removable Storage to change the volume's status from Allocated to Available.

For more information about removing volumes, see "Removing Volumes" on page 718.

Note: Due to a Microsoft limitation, data management applications cannot move volumes from an application media pool to a system media pool. When you remove a volume, the status is changed to Available but the volume is kept in the Sun StorEdge EBS media pool. To recycle the volume, use the Removable Storage MMC snap-in to move it to a free media pool, where it will be overwritten with a free media label.

Importing Sun StorEdge EBS Volumes

When you insert a Sun StorEdge EBS volume into a Removable Storage system, the next time Removable Storage inventories the library it calls the Sun StorEdge EBS MLL to read the label, and then places the volume in the appropriate pool.

Note: If a Sun StorEdge EBS release earlier than 7.0 is installed on the Sun StorEdge EBS server at the time of the inventory, Removable Storage cannot recognize Sun StorEdge EBS labels, and places Sun StorEdge EBS volumes in an unrecognized pool. The volumes cannot then be moved to an import pool and cannot be allocated to a Sun StorEdge EBS server until you update to Sun StorEdge EBS release 7.*x* or higher and perform a full Removable Storage inventory of the library.

You may need to import volumes in the following situations:

- After an update to Sun StorEdge EBS release 7.*x* or higher
- When a library is reconfigured for Removable Storage control

How to Import a Sun StorEdge EBS Volume

The following procedure assumes that the Sun StorEdge EBS media database is intact. If the media database is not intact, you must recover it before you attempt to import Sun StorEdge EBS volumes. For instructions on how to recover the media database, refer to the *NetWorker Disaster Recovery Guide*.

To import a Sun StorEdge EBS volume:

1. If the Sun StorEdge EBS volume is not already in the library, insert it into an available slot.
2. On the Windows desktop, right-click My Computer and select Manage.
3. Expand *Storage* \ *Removable Storage* \ *Media Pools* \ *Import* and check the appropriate type-specific pool to verify that Removable Storage correctly identified the Sun StorEdge EBS volume.

   Note: If the Sun StorEdge EBS volume does not appear in the appropriate import pool, refer to Microsoft documentation for information about performing a Removable Storage inventory of a library's contents.
4. Open Library Operations through the Console. For more information, see "Library Operations" on page 711.

5. From the devices menu, select Devices>Library Operations>Add. Sun StorEdge EBS software checks the media pools according to its allocation policy and allocates a volume, if possible. For information on media allocation, see "Sun StorEdge EBS Media Allocation and Deallocation" on page 712.

   When a volume is first allocated, it appears in the Slot list with a logical slot number assigned, but with "unlabeled" for the volume name.

   Note: If you now check the media pools in the Removable Storage MMC snap-in, you see that Removable Storage moved the volume from the import pool to the appropriate Sun StorEdge EBS pool, and changed its state from Available to Allocated.

6. Click Devices>Library Operations>Inventory.
7. Specify the slot number or range of slot numbers to inventory, and click OK.

   During the inventory, the volume information is registered in the media database. When the inventory is complete, the volume name appears in the volume list.

Adding New Media

When you insert blank volume into a Removable Storage system, the next time Removable Storage inventories the library it places the new volume in an unrecognized pool. Before the volume can be allocated, a free media label must be written on it.

To add new media:

1. Insert the new volume into the library.
2. On the Windows desktop, right-click My Computer and select Manage.
3. Expand *Storage \ Removable Storage \ Media Pools \ Unrecognized* and check for the new media in the appropriate pool.

   Note: If the new volume does not appear in a type-specific pool under the unrecognized pool, refer to the Microsoft documentation for information about performing a Removable Storage inventory of a library's contents.

4. Click and drag each new volume to the appropriate free pool.
5. When prompted, click OK to have Removable Storage write a free media label on the volume.

Note: To monitor the progress of this and other Removable Storage operations check the task list in *Storage \ Removable Storage \ Work Queue.*

6. From the Console, navigate to Library Operations. For more information, see "Library Operations" on page 711.
   a. From the middle pane of the Libraries window, select the device. and then click Devices>Library Operations>Add.
   b. The Sun StorEdge EBS software checks the media pools according to its allocation policy and allocates a volume, if possible.

      Each added volume appears in the Slot list with a logical slot number assigned, but with "unlabeled" for the volume name.
   c. From the file menu, select Devices>Library Operations>Label. A Sun StorEdge EBS label is written on the new volume and added to the media database.

Note: For Removable Storage, enter a +1 in the Barcodes attribute when the List option is selected. Using a +1 will add one volume within the Removable Storage Service for use by NetWorker.

For more information about volume labeling, see "Chapter 11: Device Operations" on page 297.

Using Auto Media Management with Removable Storage

Use the following procedure to configure a Removable Storage jukebox to use Sun StorEdge EBS Auto Media Management. With Auto Media Management enabled, the Sun StorEdge EBS software automatically mounts and labels allocated volumes when needed.

How to Configure Auto Media Management

To configure Auto Media Management:

1. Insert one or more blank volumes into the library.
2. On the Windows desktop, right-click My Computer and select Manage.
3. Expand *Storage \ Removable Storage \ Media Pools \ Unrecognized* and check for the new media in the appropriate pool.

   Note: If the new media does not appear in a type-specific pool under the unrecognized pool, refer to the Microsoft documentation for information about performing a Removable Storage inventory of a library's contents.
4. Click and drag each new volume to the appropriate free pool.

5. When prompted, click OK to have Removable Storage write a free media label on the volume.
6. From the Console, navigate to Library Operations. For more information, see "Library Operations" on page 711.
   a. From the left pane of the Libraries window, right-click the device and select Properties.
   b. From the Properties window, select the Configuration tab and enable the Auto media management attribute.
7. From the Console, navigate to Library Operations. For more information, see "Library Operations" on page 711.
   a. From the middle pane of the Libraries window, select the device. and then click Devices>Library Operations>Add.
   b. The Sun StorEdge EBS software checks the media pools according to its allocation policy and allocates a volume, if possible.

      Each volume added appears in the Slot list with a logical slot number assigned, but with "unlabeled" for the volume name.
   c. From the file menu, select Devices>Library Operations>Label. A Sun StorEdge EBS label is written on the new volume and added to the media database.

      For more information about volume labeling, see "Chapter 11: Device Operations" on page 297.
8. From the Console, navigate to Library Operations. For more information, see "Library Operations" on page 711.
   a. From the middle pane of the Libraries window, select the device. and then click Devices>Library Operations>Add.
   b. The Sun StorEdge EBS software checks the media pools according to its allocation policy and allocates a volume, if possible.

      Each volume added appears in the Slot list with a logical slot number assigned, but with "unlabeled" for the volume name.
   c. From the file menu, select Devices>Library Operations>Label. A Sun StorEdge EBS label is written on the new volume and added to the media database.

      For more information about volume labeling, see "Chapter 11: Device Operations" on page 297.

Removing Volumes

When removing a volume from a Removable Storage system, the Sun StorEdge EBS server does the following:

- Deletes a volume from the media database.
- Directs Removable Storage to change the volume's status.

Note: To remove a volume, barcodes must be entered. Although barcodes can be viewed in the Console Devices window in the Library details pane, they can not be copied and pasted. To copy the barcode values, use the **nsrjb** command to display the barcodes in the barcode column.

Remove volumes for the following reasons:

- To recycle a volume to a free pool and make it available to other applications.
- To move a volume to a different library.
- To temporarily deallocate a volume, and then reallocate it later.

How to Recycle a Sun StorEdge EBS Volume



Important: To recycle a Sun StorEdge EBS volume in a Removable Storage library, use only the following procedure. Do not use the Sun StorEdge EBS **erase** command or any other tool to erase the volume's contents. Doing so will get the Sun StorEdge EBS software out of sync with Removable Storage.

To recycle a Sun StorEdge EBS volume and make it available to other applications:

1. In the left pane of the Libraries window, select the server.
2. In the right pane, select the device.
3. From the file menu, select Libraries>Administer Library.

   The Library window appears.
4. From the right pane of the Libraries window, right-click on the volume to be recycled.
5. From the drop-down list, select Recycle.
6. From the Recycle dialog box, select the recycling policy that you want to apply to the volume:

- Auto
- Manual

7. Click OK.
8. On the Windows desktop, right-click My Computer and select Manage.
9. Expand *Storage \ Removable Storage \ Media Pools \ Legato Sun StorEdge EBS* and check for the deallocated volumes. Recyclable volumes are indicated by an *Available* status.
10. To recycle volumes, click-and-drag the volume icon to the appropriate free pool. When prompted for confirmation, click OK.

Performing a Reset

When you select the Reset option in the Library window, the Sun StorEdge EBS software dismounts any known mounted volume, then performs a consistency check between the Sun StorEdge EBS media database and the Removable Storage database. Any volume that appears in the Sun StorEdge EBS media database but not in the Removable Storage database is deleted from the media database. Any volume that appears in a Sun StorEdge EBS pool in the Removable Storage database, but not in the media database is registered in the media database.

Perform a Sun StorEdge EBS reset in the following situations:

- If a library is unable to complete an operation, or appears to be unresponsive.
- If the following Sun StorEdge EBS error message appears when attempting to mount an allocated, labeled volume:

  `Cannot mount volume`
- If you used the Removable Storage MMC snap-in to manage Sun StorEdge EBS volumes.

Note: If a volume was mounted in a device without Removable Storage or the Sun StorEdge EBS software being aware of it (for example, using the front panel controls on the library hardware), you must perform a hardware reset or cycle power on the library to eject the volume.

How to Perform a Reset by Using Sun StorEdge EBS Management Console

To reset the volume from the Console:

1. Open the Library Operations dialog box. For more information, see "Library Operations" on page 711.
2. From the file menu, select Devices>Library Operations>Reset.

How to Perform a Reset at the Command-prompt

1. Log in with administrator privileges to the Sun StorEdge EBS server host.
2. At the Windows command-prompt, enter:

   `nsrjb -vHE`

Device Cleaning

The Sun StorEdge EBS autocleaning feature cannot be used to configure automatic cleaning of devices in a Removable Storage system. To configure device cleaning, refer to the Microsoft documentation.

Considerations for Uninstalling the Sun StorEdge EBS Software

After the Sun StorEdge EBS server or storage node software is uninstalled, the following Removable Storage components remain intact:

- The Sun StorEdge EBS MLL, *%SystemRoot% \ System32 \ mll_lgto.dll*
- Registry keys that were created for *mll_lgto.dll* during Sun StorEdge EBS installation
- Sun StorEdge EBS application media pools
- Sun StorEdge EBS volumes, which are deallocated but still contained in a Sun StorEdge EBS media pool with a status of Available

These components allow Removable Storage to continue to recognize Sun StorEdge EBS volume labels, and keep Sun StorEdge EBS backup data secure. To recycle Sun StorEdge EBS volumes to a free pool after uninstalling the Sun StorEdge EBS software, see "Recycle a Sun StorEdge EBS Volume" on page 718.

For more information about uninstalling Sun StorEdge EBS software, refer to the *Sun StorEdge EBS Installation Guide*.

Chapter 23: Support for Microsoft SharePoint Portal Server

This chapter covers the following topics:

- "Microsoft SPS" on page 721
- "The SHAREPOINT Save Set" on page 723
- "SPS Backup" on page 724
- "SPS Recovery" on page 726
- "Performing an SPS Disaster Recovery" on page 729

Microsoft SPS

Microsoft SPS is a web portal development solution that allows you to find, share, and publish information by using Microsoft Web Parts technology.

SPS uses the Microsoft Web Storage System (similar to Microsoft Exchange Server) as a data store. SPS features include content indexing, document management, content searching, and a collaborative application platform.

Sun StorEdge EBS Software Support for SPS

Sun StorEdge EBS software provides backup, recovery, and full disaster recovery support for Microsoft SPS.

Note: Sun StorEdge EBS software currently supports only the Portal Server edition of Microsoft SharePoint, not the SharePoint Team Services edition (the SharePoint edition that is included with Microsoft Office XP).

System Requirements

Sun StorEdge EBS software support for SPS requires the following minimum system configuration:

- Microsoft Windows 2000 Server or Windows 2000 Advanced Server with Service Pack 2 (Intel x86)
- Microsoft SharePoint Portal Server 2001 (with or without Service Pack 1)
- Sun StorEdge EBS release 7.*x* on the Sun StorEdge EBS server and the SPS server (Sun StorEdge EBS client)

Note: Microsoft Windows Server 2003 does not support Microsoft SharePoint Portal Server 2001.

For installation procedures and configuration information, refer to the appropriate Microsoft documentation and the *Sun StorEdge EBS Installation Guide*. For current compatibility information, refer to the *EMC Legato Software Compatibility Guide*.

Sun StorEdge EBS Licensing

A Sun StorEdge EBS SharePoint Portal Server license is required for each SPS server (Sun StorEdge EBS client) to perform backups of the SHAREPOINT save set. Sun StorEdge EBS SharePoint Portal Server licensing is not required to perform SHAREPOINT recoveries.

For information about enabler codes for Sun StorEdge EBS software licensing, refer to the *Sun StorEdge EBS Installation Guide*.

Required Services

The following services must be running while Sun StorEdge EBS SPS backup and recovery are performed:

- Internet Information Service
- IIS Admin
- Microsoft Search
- SharePoint Portal Server
- Microsoft Exchange Information Store
- Sun StorEdge EBS Backup and Recovery Server (Sun StorEdge EBS server)
- Sun StorEdge EBS Remote Exec Service (Sun StorEdge EBS server and client)

To verify that these services are running, refer to the Microsoft Windows documentation.

The SHAREPOINT Save Set

The Sun StorEdge EBS SPS backup and recovery processes are similar to the processes for backup and recovery of the Windows SYSTEM save sets. During a Sun StorEdge EBS SPS backup, a single save set called *SHAREPOINT* is saved. The SHAREPOINT save set can be restored only in its entirety.

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.
  - 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. However, 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.

SPS Backup

You can perform a Sun StorEdge EBS SPS backup in any of the following ways:

- As part of a regularly scheduled Group backup.
- As a manual backup using the Sun StorEdge EBS User program.
- From the Windows command-prompt.

Note: Sun StorEdge EBS software currently supports only level full backups for SPS. Also, you cannot back up individual SPS workspaces or documents. You must back up the entire SPS server.

During backup, SPS suspends crawls and may briefly pause content indexing and publishing operations. Although performance may be slow, most user functionality remains available during the backup, including:

- All dashboard sites
- Document management
- Search
- Subscription
- Discussion

The index does not include any newly published documents. However, new documents are not available for searching until after the backup is complete.

Any attempt to create or delete an SPS workspace is suspended until the backup is complete. If a workspace creation or deletion is already underway when you start the backup, the backup fails and the workspace creation or deletion succeeds.

If the SPS service is not running when you attempt to perform an SPS backup, the backup fails, and the following error messages appear:

```
save: Could not stat SHAREPOINT:: The data area passed to
a system call is too small.

save: save of connecting directories failed.
```



Important: If backup of an SPS server is terminated before successful completion, the server is left in an unstable state and you must reboot. Otherwise, subsequent backup or recovery operations will fail.

Scheduled SPS Backup

You can configure scheduled SPS backups by using the Sun StorEdge EBS Administrator program. For details about configuring scheduled backups, see "Chapter 2: Backing Up Data" on page 53.

The following Sun StorEdge EBS settings are required for scheduled SPS backups.

Resource	SPS Requirement
Client	Enter **SHAREPOINT:** for the Save Set attribute. **Note**: The SHAREPOINT save set is included in save set All.
Schedule	Set the scheduled SPS backups for level full only. No other backup level is supported.

How to Perform an Ad Hoc SPS Backup

To perform an ad hoc SPS backup:

1. Log in with administrator privileges to the SPS server.
2. Verify that the required services are running (see "Required Services" on page 722).
3. In the Sun StorEdge EBS User program, click the Backup toolbar button.
4. In the Backup Configuration window, select the SHAREPOINT save set, then click the Mark toolbar button.
5. Click the Start toolbar button. The time required for the backup depends on the amount of data, network traffic, server load, and tape positioning. When the backup is complete, a message similar to the following appears:

```
Backup SHAREPOINT: 298 MB 00:05:11 4 files
Backup completion time: 1-04-06 3:27p
```

How to Perform an SPS Backup from the Command-prompt

To perform a Sun StorEdge EBS SPS backup from the command-prompt:

1. Log in with administrator privileges to the SPS server.
2. Verify that the required services are running (see "Required Services" on page 722).

3. At the command-prompt, change to the <*Sun StorEdge EBS_install_path*> \ *bin* directory.
4. Enter the following command:

   `save -s` *servername* `"SHAREPOINT:"`

   where *servername* is the name of the Sun StorEdge EBS server performing the backup.

How to Verify an SPS Backup

To verify that a SHAREPOINT save set was successfully saved:

1. In the Administration interface, click Media.
2. In the expanded left pane, select Volumes.
3. In the right pane, double-click the volume name.
4. Check the save set size. A SHAREPOINT save set of 1 KB or less indicates that the backup failed.

SPS Recovery

You can recover an SPS server to the host from which it was backed up, or to a different host. During recovery, SPS becomes unavailable to users, and all existing content is overwritten. If users attempt to access SPS during the recovery, they receive an "unavailable" message.



Important: Sun StorEdge EBS software uses the SHAREPOINT save set to back up and recover the Microsoft Web Storage System database, the central data store for SPS. This database must be backed up and recovered in its entirety. Due to Microsoft SPS design, if recovery of a SHAREPOINT save set is terminated before the database is fully restored, essential services (Microsoft Exchange Information Store and SharePoint Portal Server) will fail to start and the SPS server will not be operational.

To return the server to an operational state, uninstall SPS, and then reinstall SPS. For more information, refer to the Microsoft SPS documentation.

The SPS reinstallation requirement applies to incomplete recovery of the SHAREPOINT save set only. Incomplete recovery of other parts of the server do not require SPS reinstallation.

Notes:

- Sun StorEdge EBS software does not support directed recovery of the SHAREPOINT save set. To recover the SHAREPOINT save set to a server other than the one from which it was saved, see "Recovery to a Different Server" on page 728.
- After an SPS recovery, you might be prompted to reboot the computer. Rebooting is not necessary.

How to Perform an SPS Recovery with the Sun StorEdge EBS User Program

To perform an SPS recovery from the Sun StorEdge EBS User program:

1. Log in with administrator privileges to the SPS server.
2. Verify that the required services are running (see "Required Services" on page 722).
3. Start the Sun StorEdge EBS User program.
4. Click the Recover toolbar button.
5. In the Source Client dialog box, click OK to select the local client.
6. In the Destination Client dialog box, click OK to select the local client.
7. In the Recover window, select the SHAREPOINT save set, then click the Mark toolbar button.
8. Click the Start toolbar button to start the recovery. The time required for the recovery depends on the amount of data, network traffic, server load, and tape positioning. When the recovery is complete, a message similar to the following appears:

```
Received 1 matching file(s) from NSRserver servername
Recovery completion time: 1-04-06 4:27p
```

How to Perform an SPS Recovery from the Command-prompt

To perform a Sun StorEdge EBS SPS recovery from the command-prompt:

1. Log in with administrator privileges to the SPS server.
2. Verify that the required services are running (see "Required Services" on page 722).
3. At the command-prompt, change to the <Sun StorEdge EBS_*install_path*> \ *bin* directory.
4. Enter the following command:

   `recover -s` *servername* `-N "SHAREPOINT:"`

   where *servername* is the name of the Sun StorEdge EBS server that performed the SPS backup

Recovery to a Different Server

To recover a SHAREPOINT save set to a host other than the one from which it was backed up:

1. Determine the save set ID (*ssid*) of the SHAREPOINT save set to be recovered. To browse a list of valid ssid's:
   a. In the Administration interface, click Media.
   b. In the expanded left pane, select Volumes.
   c. In the right pane, double-click the volume name.
   d. Note the SSID column for the ssid of the SHAREPOINT save set for use later in this procedure.
2. Log in with administrator privileges to the new SPS server.
3. Ensure that the new SPS server host meets the system requirements, with all the necessary software installed and configured (see "System Requirements" on page 722).
4. Verify that the required services are running (see "Required Services" on page 722).
5. At the command-prompt, change to the <Sun StorEdge EBS_*install_path*> \ *bin* directory.
6. Enter the following command:

   `recover -s` *servername* `-iY -S` *ssid*

   where:

- *servername* is the name of the Sun StorEdge EBS server that performed the SPS backup
- *ssid* is the save set ID of the SHAREPOINT save set obtained at the beginning of this procedure.

Performing an SPS Disaster Recovery

If you experience complete failure of an SPS server host, perform a disaster recovery as follows:

1. Recover the SPS server host. For instructions, refer to the *NetWorker Disaster Recovery Guide*.
2. Perform a new SPS installation on the recovered host computer. For details, refer to Microsoft documentation. (The installation of a new SPS reinstalls the Web Storage System, which is not saved during the SPS backup and could not be restored during the disaster recovery.)

   Note: During the SPS installation, skip the Workspace Creation step. Any workspace you create would be completely overwritten when you recover the SHAREPOINT save set.
3. Recover the SHAREPOINT save set. For instructions, see "SPS Recovery" on page 726.

23

Chapter 24: Support for Microsoft Automated System Recovery

This chapter covers the following topics:

- "Microsoft ASR" on page 731
- "Sun StorEdge EBS Support for ASR" on page 732
- "ASR Backup" on page 734
- "Creating an ASR Disk" on page 735
- "Using the ASR Disk to Recover a Sun StorEdge EBS Client" on page 739
- "Limitations and Special Considerations" on page 743

Microsoft ASR

Microsoft Automated System Recovery (ASR) is a Windows XP and Windows Server 2003 feature that enables backup and recovery applications, such as the Sun StorEdge EBS software, 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.

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 StorEdge EBS Support for ASR

Sun StorEdge EBS software supports ASR on Sun StorEdge EBS clients that run Windows XP Professional or Windows Server 2003. ASR is not supported for Sun StorEdge EBS servers or storage nodes.

ASR operates only when the client is backed up to a Sun StorEdge EBS Server that is running Microsoft Windows 2000 or Windows Server 2003.

ASR is *not* supported for Sun StorEdge EBS clients configured with a shared disk. This applies to stand-alone clients, as well as clients operating in an MSCS cluster environment. For disaster recovery information about these configurations, refer to the *NetWorker Disaster Recovery Guide.*

The Sun StorEdge EBS ASR Save Set

The Sun StorEdge 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 StorEdge EBS SYSTEM save sets, ASR is backed up automatically at level full.

Usage Overview

Because ASR is the recommended Sun StorEdge EBS disaster recovery method for Windows XP Professional and Windows Server 2003 Sun StorEdge 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 StorEdge EBS Administrator program.
- In an ad hoc backup by using the Sun StorEdge EBS User program.

You can create an ASR disk by using either of these methods:

- By restoring the ASR save set on the Sun StorEdge EBS client computer.
- By performing a directed recovery of the ASR save set to a different Sun StorEdge EBS client computer.

In the event of a disaster on the Sun StorEdge EBS client, use the ASR disk to perform an ASR recovery. This must be done locally, on the failed Sun StorEdge EBS client computer.

Note: If a current ASR disk or ASR save set is not available for the failed Sun StorEdge EBS client computer, you must use the legacy Sun StorEdge EBS disaster recovery method. For more information, refer to the *NetWorker Disaster Recovery Guide*.

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. For more information, see "Post ASR Disk Creation Task" on page 738.

ASR Backup



Important: Include ASR in a scheduled group backup for each Windows XP Professional and Windows Server 2003 Sun StorEdge EBS client. When you use the ASR disk to recover a Sun StorEdge EBS client computer, any data and configuration changes made since the last ASR backup will be lost.

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:

1. In the Save Set attribute for the Client resource, enter All.

   Alternatively, select at least 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:

For more information about specifying save sets for a scheduled backup, see "Chapter 2: Backing Up Data" on page 53.

How to Include the ASR or VSS ASR DISK Save Set in a Manual Backup

Backup of data by selecting My Computer is the preferred method of performing ad hoc 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 client's ad hoc backup:

1. In the Sun StorEdge EBS User program, click the Backup button and select My Computer to save all data.

   Alternatively, select at least 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:

For more information about ad hoc backups, see "Chapter 2: Backing Up Data" on page 53.

Creating an ASR Disk

Before an ASR recovery, create an ASR disk for the Sun StorEdge EBS client computer that will be recovered. An ASR disk is created using the Sun StorEdge EBS 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 StorEdge 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. For more information about backing up an ASR DISK save set, see "ASR Backup" on page 734.
- Multiple floppy diskettes will be required to create an ASR disk.

How to Create an ASR Disk Locally

If the Sun StorEdge EBS client computer for the new ASR disk is not functional, perform a directed recovery of the ASR save set. For instructions, see "How to Create an ASR Disk by Using Directed Recovery" on page 737.

To create an ASR disk locally:

1. Log in with administrator privileges to the Sun StorEdge EBS client computer for which you want to create the ASR disk.
2. In the Sun StorEdge EBS User program, click the Recover button.
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 the Start button. 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 the following options:
   - Yes — During recovery of the Sun StorEdge 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 an ad hoc 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 StorEdge 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 an ad hoc 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 "Post ASR Disk Creation Task" on page 738.

How to Create an ASR Disk by Using Directed Recovery

If the Sun StorEdge 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 StorEdge EBS administration, and in cases where you cannot access the *%SystemRoot% \ Repair \ nsr* directory on the Sun StorEdge 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 StorEdge EBS client computer.
2. In the Sun StorEdge EBS User program, click the Recover button.
3. In the Source Client dialog box, select the Sun StorEdge 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 the Start button. 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 the following options:

- Yes — During recovery of the Sun StorEdge 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 an ad hoc 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 StorEdge 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 an ad hoc 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 "Post ASR Disk Creation Task" on page 738.

Post 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 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 the following convention to name client 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 StorEdge 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 *NetWorker Disaster Recovery Guide*.

During an ASR recovery, the following 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 StorEdge EBS backup save sets to recover.
- The Sun StorEdge 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 the following:

- A current ASR disk for the computer being recovered (see "Creating an ASR Disk" on page 735).

  If a current ASR disk or ASR save set is not available for the failed computer, use the legacy Sun StorEdge EBS disaster recovery method. For procedures, refer to the *NetWorker Disaster Recovery Guide*.
- 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. For more information about this restriction, see "OEM Recovery CDs Are Not Supported" on page 744.

- The latest Sun StorEdge EBS backup for the computer you are recovering.

  Note: If you routinely move Sun StorEdge EBS backup media to an offsite location for safekeeping, ensure that all necessary volumes are available before beginning the recovery. To list the media associated with the files you want to recover, run **mminfo -mv** from the command-prompt. For more information about the **mminfo** command, refer to the *NetWorker Command Reference Guide* or UNIX man pages.

How to Perform an ASR Recovery

To perform an ASR recovery on a Sun StorEdge 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.

   Notes:

   - 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 StorEdge 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 the following legacy save sets:

SYSTEM STATE:\
SYSTEM DB:\
SYSTEM FILES:\

If VSS is licensed and enabled, the following VSS save sets are included:

VSS SYSTEM BOOT:\
VSS SYSTEM FILESET:\

Notes:

- VSS USER DATA, VSS OTHER, and VSS SYSTEM SERVICES do not appear because they are not required to boot from ASR mode.
- For more information about the "pause during recovery" option, see "Creating an ASR Disk" on page 735.

By default, the displayed save sets represent the most recent backup. You can view and select previous backups by entering 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. For more information about **nsr_getdate**, refer to the *NetWorker Command Reference Guide* or UNIX man pages.



Important: The VSS components and certain components cannot be correctly restored during ASR recovery. To ensure that all necessary components are properly recovered, see "Components That Require Special Handling After an ASR Recovery" on page 741.

Components That Require Special Handling After an ASR Recovery

Due to limitations in Microsoft ASR functionality, the following 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 StorEdge EBS client being recovered uses any of these components, perform the following procedure after an ASR recovery:

1. Log in with administrator privileges to the target computer.
2. Start the Sun StorEdge EBS User program.
3. Click the Recover button.
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. In the Recover window, select the SYSTEM STATE save set:
   a. Check for the presence of the COM+ Registration Database component.
   b. 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 the following components:

      Disk Quota Database

      Windows Management Instrumentation Database

   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 the Start button.

How to Verify the Sun StorEdge EBS Client Recovery

Note: 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.

To verify the Sun StorEdge EBS client recovery:

1. Reboot the Sun StorEdge EBS client computer and verify that the Sun StorEdge EBS Remote Exec and Sun StorEdge EBS Power Monitor services have started.

   Note: You can disable the Sun StorEdge EBS Power Monitor service if it is not needed. For more information, see "Appendix D: Additional Features of the Microsoft Windows Server" on page 837.

2. Use the Windows Event Viewer to examine the event logs for errors. In particular, check for the following:
   - Service startup errors related to the Windows system state
   - Errors regarding the recovery of Windows system-protected files

   For information about how the Sun StorEdge EBS software handles the Windows system state and system-protected files, see "Appendix D: Additional Features of the Microsoft Windows Server" on page 837.
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.

Limitations and Special Considerations

The following limitations and special considerations apply to ASR 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. For more information about this option, see "Creating an ASR Disk" on page 735.
- When the ASR recovery operation pauses, clear the FAT16 partition to exclude it from the recovery. For more information about performing ASR recoveries, see "Using the ASR Disk to Recover a Sun StorEdge EBS Client" on page 739.

After you have completed the ASR recovery and rebooted, recover the FAT16 partition in a separate, non-ASR Sun StorEdge EBS recovery operation. For information about recovering data, see "Chapter 13: Recovering Data" on page 367.

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.



Important: To perform an ASR recovery, you *must* have an official Microsoft Windows installation CD for the version of Windows you are recovering. If you have only an OEM recovery CD that includes the Windows installation files, refer to the *NetWorker Disaster Recovery Guide*. Also, refer to OEM documentation as necessary.

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.

Chapter 25: Volume Shadow Copy Service

This chapter covers the following topics:

- "Overview of VSS" on page 745
- "VSS and the Backup Process" on page 746
- "Controlling VSS from NetWorker Software" on page 748

Overview of VSS

Volume Shadow Copy Service (VSS) is a Microsoft technology that acts as a coordinator among all the components that create, archive, modify, back up, and restore data, including:

- The operating system
- Storage hardware
- Applications
- Utility or backup programs, such as NetWorker software

VSS allows for the creation of a point-in-time *snapshot*, or temporary copy, of a volume. Instead of backing up data directly from the physical file system, data is backed up from the snapshot. In addition, VSS allows for a single, point-in-time capture of the system state.

In NetWorker software releases 7.1 and later, if a VSS client license is present, NetWorker software uses VSS technology to create snapshot backups of volumes and exact copies of files, including all open files. For example, 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 StorEdge EBS PowerSnap Module software. For more information, refer to the PowerSnap Module documentation.

VSS and the Backup Process

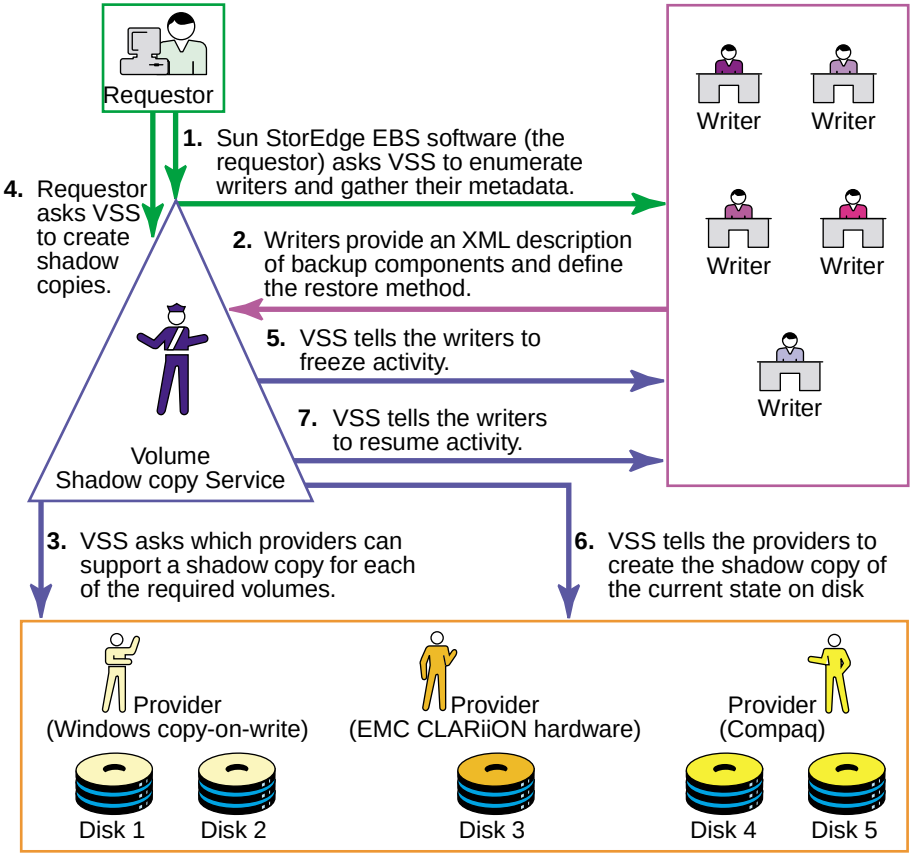
In VSS terms, NetWorker software is a *requestor*—an application that needs data from other applications or services. When a requestor needs data from an application or service, the following process occurs:

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. 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. VSS then calls a *provider* to capture the requested data. 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. See Figure 50 on page 747 for a graphical representation of the VSS backup process.

Figure 50. VSS Backup Process



NetWorker software backs up data from the point-in-time snapshot that is created during this process. Any subsequent data access is performed on the snapshot, *not* the 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, NetWorker 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 the following conditions are true:

- The writer associated with the application is present on the system and active.
- The writer appears on the list of supported writers in the NSRLA database.
- A user has not disabled the writer.

If these conditions are all true for a particular writer, NetWorker software defaults to backing up data by using VSS technology. If any of the conditions are false for a particular writer, the data served by that writer is excluded from the backup operation.

List of Supported Writers

During a VSS backup operation, NetWorker software validates each writer against a list of supported writers. As part of a software release, or between releases, there may be updates to the list of supported writers. For a list of the currently supported writers, refer to the *EMC Legato Software Compatibility Guide*.

Controlling VSS from NetWorker Software

By default, NetWorker software release 7.1 and later uses VSS technology to back up a client that is also running release 7.1 or later, and is licensed to use VSS. For VSS SYSTEM save sets, this means NetWorker software uses VSS for most save sets and writers (for details, see "VSS SYSTEM Save Sets" on page 807). For the file system, this means the software attempts to take a

snapshot of each drive, but if it fails, then it saves the file system by using the *legacy method* (that is, no snapshot is taken). During a given backup for an individual client, either the VSS method or the legacy method is used, but not both.

There may be times when you need finer control over how NetWorker software uses VSS. For example, if you need to disable a writer or completely disable VSS. You can control VSS from the Administration window, the Sun StorEdge EBS User program, or the command-prompt.

For information about:

- Controlling VSS from the Administration window, see "How to Control VSS from the Administration Window" on page 749.
- Controlling VSS from the Sun StorEdge EBS User program, see "How to Control VSS from the Sun StorEdge EBS Client" on page 751.
- Controlling VSS from the command-prompt, see "How to Control VSS from the Command-prompt" on page 752.

Since VSS is a new technology, run a test backup and recovery operation to see how these operations work with VSS.

Note: To control VSS when backing up a client running NetWorker software release 7.1 or later, where the Sun StorEdge EBS *server* is running a prior release of NetWorker software, refer to the documentation that accompanied the VSS client license.

How to Control 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, enter the appropriate command according to Table 105 on page 753, then click OK. Separate multiple commands with a semicolon (;). If the Save Operations attribute is left blank, NetWorker software backs up data by using VSS if a VSS client license is present.

Notes:

- The Save Operations attribute does not support Sun StorEdge EBS Module save sets. If a Sun StorEdge EBS Module save set name is entered in the window, the backup will fail.
- If you enter a VSS command in the Save Operations attribute of the Administration window, the command will be executed 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 will be executed only if the backup is started by using the Sun StorEdge EBS User program on the client.
- When entering 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 Sun StorEdge EBS User program and browse the corresponding save set. Keep in mind that the list of writers appearing under each save set is dynamically determined at runtime. A writer may not appear under its corresponding save set if:
  - NetWorker software does not support it.
  - Its associated application, service, or database is not currently running on the system.
  - It has already been disabled by using the Sun StorEdge EBS User program.



Important: Only use the Save Operations attribute for clients running NetWorker software release 7.1 or later. If anything is entered in this attribute for a client running a prior NetWorker software release, the backup will fail.

How to Control VSS from the Sun StorEdge EBS Client

If you open the Backup window to browse data before opening the Local Save Operations dialog box, you will not be able to modify Local Save Operations. The Local Save Operations dialog box becomes read-only once the Backup window is opened. To modify Local Save Operations after opening the Backup window, exit the Sun StorEdge EBS User program and restart it.

If you are going to modify Local Save Operations, run only one instance of the Sun StorEdge EBS User program at a time.

To control VSS from the Sun StorEdge EBS client:

1. Start the Sun StorEdge EBS User program.
2. From the Options menu, select Save Operations.
3. In the Local Save Operations dialog box, enter the appropriate command according to Table 105 on page 753, then click OK. Separate multiple commands with a semicolon (;). If the Local Save Operations dialog box is left empty, NetWorker software backs up data by using VSS if a VSS client license is present.
    - Local Save Operations does not support Sun StorEdge EBS Module save sets. If a Sun StorEdge EBS Module save set name is entered in the window, the backup will fail.
    - If you enter a VSS command in the Save Operations attribute of the Administration window, the command will be executed 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 will be executed only if the backup is started using the Sun StorEdge EBS User program on the client.
    - When entering 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 appearing under each save set is dynamically determined at runtime. A writer may not appear under its corresponding save set if:
        – NetWorker software does not support it.
        – Its associated application, service, or database is not currently running on the system.
        – It has already been disabled using the Sun StorEdge EBS User program.

How to Control VSS from the Command-prompt

VSS can be controlled from the command-prompt on a Sun StorEdge EBS client or the Console server by using the **-o** option and the Save Operations commands in Table 105 on page 753, 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*, enter the following:

```
save -s jupiter -o "vss:*=off" "C:\myfile"
```

Although the server name is not required in the preceding command example, include it to ensure that the **save** command finds the correct server. Separate multiple Save Operation commands with a semicolon (;).

For more information about the **save**, **savefs**, and **nsrarchive** commands and command options, refer to the *NetWorker Command Reference Guide* or the UNIX man pages.

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.

VSS Commands

Table 105 on page 753 lists the commands and syntax used to control VSS. These commands may be entered the following locations:

- Save Operations attribute of the Administration window on the Console server.
- Local Save Operations dialog box of the Sun StorEdge EBS User program on the Sun StorEdge EBS client.
- Command-prompt on the Sun StorEdge EBS client or the Console server.

Table 105. VSS Commands (Part 1 of 2)

Task	Entry	Result
Enable VSS.		Leaving the attribute empty will result in NetWorker 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 StorEdge EBS Module exists for that application and is installed and configured on the system. **Note:** When a writer is disabled, NetWorker software still processes the writer so its files are skipped during the file system backup.

Table 105. VSS Commands (Part 2 of 2)

Task	Entry	Result
Instruct NetWorker 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 Exchange Server Databases

This release of NetWorker software does not support VSS backups of 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, NetWorker 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, NetWorker software excludes the entire writer from the backup operation.

If you used a previous release of NetWorker 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.

Chapter 26: Troubleshooting

This chapter covers the following topics:

- "Before Contacting Technical Support" on page 755
- "Sun StorEdge EBS Functionality Issues" on page 758
- "Devices and Autochangers" on page 771
- "Sun StorEdge EBS Locale and Code Set Support" on page 776
- "Resource Database Notes" on page 777
- "Network and Server Communication Errors" on page 779
- "Sun StorEdge EBS Archiving and Retrieval" on page 790
- "Storage Nodes" on page 791
- "Console Error Messages and Corrective Actions" on page 793
- "Console Log Files" on page 798
- "Console Troubleshooting Notes and Tips" on page 799

Before Contacting Technical Support

If the solutions in this chapter do not solve the problem, go to *http://softwaresupport.emc.com* for technical assistance. Provide the following information:

- The software version of the Sun StorEdge EBS component.
- The operating system version.

  For example, for Solaris, use the following command:

  `uname -a`

  For AIX, use the following command:

  `oslevel`

- The hardware configuration.
- Information about devices and other SCSI IDs.

  For example, AIX, DYNIX / ptx, IRIX, Linux, and Solaris, use the following command:

  `/usr/sbin/inquire`

  For HP-UX, use the following command

  `/etc/ioscan`
- 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, use the following command:

  `pkginfo -x SUNWebsn`

  For AIX, use the following command:

  `lslpp -l | grep Sun`
- Be able to supply the following information:
  - How to reproduce the problem.
  - The exact error messages.
  - Number of times you have seen the problem.
  - Whether the Sun StorEdge EBS command was successful before you made any changes and, if so, the changes you made.

Determining the Version of Sun StorEdge EBS Software Running on a Client

To determine the version of the Sun StorEdge EBS software running on a client, use either the Sun StorEdge EBS client interface or the **nsradmin** command.

How to Determine the Software Version by Using the Client Interface

To determine the software version by using the client interface:

1. Open the Sun StorEdge EBS client program:
   - On Windows, open the Sun StorEdge EBS User program.
   - On UNIX, open one of the Sun StorEdge EBS client programs such as **nwbackup**.

   For information about opening these programs, see "Sun StorEdge EBS Client Interface" on page 38.

2. From the Help menu, select About Sun StorEdge EBS User.

   The Sun StorEdge EBS version number appears in the About dialog box.

3. Click OK to close the dialog box.

How to Determine the Client Software Version by Using nsradmin

To determine the client version by using **nsradmin**:

1. At the command-prompt, enter:

   nsradmin

2. At the **nsradmin** command-prompt, enter:

   nsradmin> **show name; NetWorker version**

   nsradmin> **print type: NSR client**

This displays the version of Sun StorEdge EBS software running on each client.

Creating Reports for Troubleshooting

The **nsrsup** program gathers system information, including Sun StorEdge EBS software specific details and generates a HTML report that can be used for troubleshooting. For information about the **nsrsup** command, refer to the *NetWorker Command Reference Guide* or the UNIX man page.

Note: The *nsrsup.exe* interface replaced the previous DOS batch file version: *nsrsup.bat* in the Sun StorEdge EBS 7.0 release.

To run the **nsrsup** program in interactive mode:

1. Navigate to the <*Sun StorEdge EBS_install_path*> \ *bin* directory.
2. Enter the following at the command-prompt to run the **nsrsup** program and generate an HTML format report:

   nsrsup

   or

   nsrsup --html

3. Follow the prompts to finish generating the report.

Note: It could take some time for the **nsrsup** command to complete.

The reports generated by the **nsrsup** program are written to the client's file system. For example:

- *%SystemDrive% \ Temp \ nsrsuplog.html*
- *[location specified by %TEMP% environment variable] \ nsrsuplog.html*
- */tmp/nsrsup.pid.out.html*
- */nsr/tmp/nsrsup.pid.out.html*

Note: Subsequent HTML reports are written to *nsrsuplog2.html*.

Displaying Diagnostic Mode Attributes

Sun StorEdge 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. View the properties of any resource to see diagnostic attributes.

Sun StorEdge EBS Functionality Issues

This section describes workarounds for Sun StorEdge EBS issues.

Backup and Recovery

This section covers backup and recovery operations.

Multiple nsrexecd Processes

You might notice multiple instances of the **nsrexecd** process running on the Sun StorEdge EBS server during a backup. This is normal behavior and happens because one **nsrexecd** process is started for each active save stream.

Checking the Sun StorEdge EBS Services

If you have trouble starting Sun StorEdge 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 the following commands:

  ps -ef | grep nsr

  ps -ax | grep nsr

You should receive a response similar to the following:

```
12217 ?        S  0:09 /usr/sbin/nsr/nsrexecd -s jupiter
12221 ?        S  2:23 /usr/sbin/nsr/nsrd
12230 ?        S  0:00 /usr/sbin/nsr/nsrmmdbd
12231 ?        S  0:01 /usr/sbin/nsr/nsrindexd
12232 ?        S  0:00 /usr/sbin/nsr/nsrmmd -n 1
12234 ?        S  0:00 /usr/sbin/nsr/nsrmmd -n 2
12410 pts/8    S  0:00 grep nsr
```

If daemons are not present, start the Sun StorEdge EBS daemons.

Unsynchronized Client Clock Errors

This situation affects computers with Sun StorEdge EBS client software release 5.0 or later. If the setting for the system clock on a Sun StorEdge EBS client differs from that of the Sun StorEdge 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, these messages appear, but the schedule is still honored. To avoid receiving these messages, 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 AM and backups are scheduled to occur at 2:01 AM. At 2:00 AM, the clock is moved forward to 3:00 AM. 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.

Clone ID Timestamp Does Not Reflect the Time the Clone was Created

To guarantee that cloned savesets created on different storage nodes do not have the same timestamp, the Sun StorEdge EBS software assigns a timestamp to cloned savesets 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 the following:

- Stop the conflicting group from running by clicking Stop in the Group Control window. For more information, see "How to Stop a Group" on page 478.
- Shut down and restart the Sun StorEdge EBS services. For more information, see, "Manually Stopping and Starting Sun StorEdge EBS Services" on page 49.

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. For information on save set recovery, see "Recovering by Save Set Selection" on page 379.

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.

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 StorEdge 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, use the **nsr_shutdown** command to shut down the Sun StorEdge EBS server.

Restart the Sun StorEdge EBS server manually.

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, the following 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 StorEdge 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, the following 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 the following 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 StorEdge EBS binaries.

Illegal Characters in Configurations

When naming label templates, directives, groups, policies, and schedules, the following characters are not allowed:

```
/ \ * [ ] ( ) $ ! ^ ’ " ? ; ‘ ~ < > & | { }
```

When naming label templates, directives, groups, policies, and schedules, the following characters are not allowed:

```
/ \ * [ ] ( ) $ ! ^ " ’ ; ? ~ & | { }
```

Hostname Aliases

Savegroups fail when aliases are improperly defined. Under certain conditions (usually related to improperly configured DNS servers or hosts files), the Sun StorEdge 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 StorEdge EBS aliases field.

If you encounter any of the following situations, a client alias problem might be the cause:

- The following 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 */nsr/index* (the directory that contains the indexes) there are two directories for the same client using 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.



Important: Do not include aliases that are shared by other hosts in the Aliases attribute.

Directory Pathname Restrictions

A file manager (but not Windows Explorer) restriction causes errors when a pathname contains too many characters. To avoid these errors, use pathnames with fewer than 128 characters.

Backup of Recursive Subdirectories

The backup fails when attempting to back up 420 or more recursive subdirectories. The maximum number of recursive subdirectories that can be backed up is 419.

Failed Recover Operation Directly after New Installation

If you attempt to start the **nwrecover** program immediately after installing Sun StorEdge EBS software for the first time, the following error message appears:

```
nwrecover: Program not found.
```

To avoid the problem, perform a Sun StorEdge EBS backup on the client.

Recovering Files from an Interrupted Backup

If you terminate a backup by stopping the Sun StorEdge EBS services, you cannot recover the files because the media database is not updated when the services stop. Consequently, the Sun StorEdge 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, the following 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, this message might appear because:

- The clocks on the client and server are not synchronized.
- The **savegrp** session begins before midnight and ends after midnight.

Cannot Perform 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 StorEdge 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 an ad hoc backup of the file. For more information, see "Ad Hoc Backups" on page 66.

Renamed Clients Cannot Recover Old Backups

The Sun StorEdge 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 using the old client name, perform a directed recovery by directing data saved under the old client name to the new client. For information about performing directed recoveries, see "Directing Recoveries to Another Client" on page 387.

Client File Index Errors

The following issues are related to client file indexes.

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 might 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 StorEdge 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 may 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 Online Indexes for a Sun StorEdge EBS Server" on page 395.

No Notification of Client File Index Size Growth

The Sun StorEdge 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. For information on how to manage the Sun StorEdge EBS client file indexes, see "Reducing Client File Index Size" on page 560. For more information, refer to the **nsrls**, **nsrck**, and **nsrim** man pages, or to the *NetWorker Command Reference Guide* or the UNIX man pages.

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, the following 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 may occur:

- The recovery might stop immediately.
- The files that still need to recovered are listed.
- Messages similar to the following 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 scrolling filenames and checksum message indicate that a recovery was not stopped cleanly. A fix is planned for a future release.

RPC Error

If Sun StorEdge EBS has trouble backing up a directory path, a message similar to the following 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, the following 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, Windows 2000, and Windows NT 4.0 have a set desktop heap size limitation that may produce the following error message when exceeded:

```
The application failed to initialize properly
```

For information about this problem and how to correct it, refer to *Microsoft* Knowledge Base article 142676 on the Microsoft website.

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 StorEdge EBS software.

- Avoid using the All save set. Instead, specify each drive letter and SYSTEM or VSS SYSTEM save set separately. For more information about the All save set, see "Scheduling Predefined Save Sets for Backup" on page 64.

Disk Label Errors

If a nonoptical device was configured as an optical device, the following 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 StorEdge 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 StorEdge EBS server.

The information to recover server indexes is stored in the media database on the Sun StorEdge EBS server.



Important: The **savegrp** no longer performs a server index backup for servers not in an active group.

Copy Violation

If Sun StorEdge EBS software is installed on multiple servers and the same Sun StorEdge EBS enabler code was used for them all, messages similar to the following 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 StorEdge EBS software from the extra servers.
3. Restart the Sun StorEdge EBS services on the server where the backups are to go.

Xview Errors

If the following error message appears when you attempt to start the Sun StorEdge EBS Console, it means that the client is not authorized to display Sun StorEdge 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 Sun StorEdge EBS Console:

1. From the client machine, invoke the **xhost** command:

   xhost *server_name*

2. Remotely log in to the Sun StorEdge EBS server and issue the **setenv** command:

   setenv DISPLAY *client_name*:**0.0**

For command shells other than **csh** enter:

DISPLAY=*client_name***:0.0**

export DISPLAY

Converting Sparse Files to Fully Allocated Files

The Sun StorEdge 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 the following 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.

   For information about setting the inactivity timeout attribute for the group, see "How to Edit a Group" on page 138.

The mminfo -N Command Is Case-Sensitive Regarding Save Set Names

When querying the media database 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 StorEdge EBS server stores the save set name in uppercase in the media database. For example, if the save set name representing the drive partition was entered in the Client resource in lowercase (*c:*), you must query 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 in the directory were changed, the renamed directory is not included in subsequent incremental backups.

To avoid this problem, rather than renaming the original directory, copy the directory and then rename the copy. This causes the Sun StorEdge EBS software to treat the renamed copy of the directory as new. After naming a directory, either manually back up the directory or run a level full backup.

If you attempt to recover files that reside in the renamed directory, the renamed directory will not be visible when browsing backups. To recover the files, change the browse time to a time immediately before the directory was renamed, and then recover files under the old directory name. Relocate the recovered files to the renamed directory, or copy them after the recovery.

Devices and Autochangers

This section explains how to resolve problems with devices and autochangers.



Important: 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. For information about displaying hidden attributes, see "Displaying Diagnostic Mode Attributes" on page 758. For information about these attributes, refer to the field help or the *NetWorker Command Reference Guide* or the UNIX man pages.



Important: Do not change time-related attributes unless advised to do so by a Technical Support representative.

Maintenance Commands

Sun StorEdge EBS device driver software provides maintenance commands, such as **lusbinfo** and **lusdebug**, for diagnosing problems on tape devices and autochangers. For information about these commands, refer to the *NetWorker Command Reference Guide* or the UNIX man pages.

Autodetected SCSI Jukebox Option Causes Server to Hang

If an autodetected SCSI jukebox is installed using **jbconfig**, and the server hangs:

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 the following message appears:

   ```
   Jukebox has been added successfully.
   ```

Autochanger Inventory Problems

The following 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 Sun StorEdge EBS 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 StorEdge EBS server.
3. Reset the autochanger by entering the following command:

   nsrjb -Hv

4. Perform an inventory by entering the following command:

   nsrjb -Iv

   After the inventory operation completes, the Sun StorEdge EBS server can use the autochanger.

   For complete information on the use of the **nsrjb** command, refer to the **nsrjb** man page or see "Chapter 3: Configuring Libraries and Devices" on page 101.

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 StorEdge EBS server, the following 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 StorEdge 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 StorEdge 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 StorEdge EBS filemarks consume space on the tape.

  The Sun StorEdge 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.

For detailed information, refer to *EMC Legato Technical Bulletin 176: NetWorker and Tape Capacity Discrepancies (UNIX | NetWare),* available at www.sun.com.

Server Cannot Access Autochanger Control Port

The control port controls the autochanger loading mechanism. To verify whether the control port is properly connected, refer to the autochanger hardware installation manual. If you cannot determine whether the control port is working, contact the autochanger vendor for assistance.

Nonrewinding Device Requirement

Use a nonrewinding device for Sun StorEdge EBS backups. The Sun StorEdge 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, run the following command to mount only the _AF_readonly device:

`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 StorEdge EBS services.
2. Shut down and restart the autochanger with the TZ89 drives.
3. When the autochanger is back online, restart Sun StorEdge EBS services.

   This will reset Sun StorEdge EBS so that it stops trying to unload the drive.
4. Use the following settings for the sleep time attributes:
   - Eject Sleep: 18 secs
   - Unload Sleep: 40 secs
   - Load Sleep: 40 secs

   For more information, see "Additional Attributes in the Autochanger Resource" on page 772.
5. Attempt to unload the drive again. If the drive fails to unload, repeat this procedure and increase the sleep times.

Message Displayed When CDI Enabled on NDMP or Disk File Type Device

If the CDI feature is enabled while using an NDMP tape device or file type device, a message similar to the following appears in the Sun StorEdge 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 StorEdge EBS server or storage node that has multiple Network interface cards (NIC). To prevent this situation, add the domain name of each additional NIC to the Aliases attribute in the Client resource that is set up for the Sun StorEdge EBS server or storage node. For information about editing a Client resource, see "How to Edit a Client" on page 568.

Sun StorEdge EBS Locale and Code Set Support

Sun StorEdge 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 StorEdge EBS software might function normally in such an environment, but might not function normally if the locale is changed (the previously existing indexes can become invalid).

Locale Support

Sun StorEdge EBS software supports *only* the following mixed-locale environments:

- English server and a supported localized client
- English client with a supported localized server

Sun StorEdge EBS functionality varies depending on whether the software is localized:

- A localized version of Sun StorEdge EBS software on a localized operating system running a supported locale displays localized filenames correctly.
- An English version of Sun StorEdge EBS software, running on either a localized or English version of an operating system, might not display non-English filenames correctly.

Code Set Support

Sun StorEdge EBS software for Microsoft Windows supports the Windows standard code set, which is 16-bit Unicode. The Japanese versions of Windows and Sun StorEdge EBS software use Shift-JIS character encoding.

Resource Database Notes

With Sun StorEdge EBS release 6.2 and later, resource information resides in the following directory structure:

<Sun StorEdge EBS_install_path> \ res \ nsrdb \ 00
.
.
<Sun StorEdge EBS_install_path> \ res \ nsrdb \ 09

Within this directory structure, 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 StorEdge EBS client (**nsrexecd**) continues to use the *<Sun StorEdge EBS_install_path> \ res \ nsrla.res* file.

Viewing Resources

You can view all Sun StorEdge EBS resources (and administer a Sun StorEdge EBS release 6.2 or later server) through the Administration window.



Important: Although you can view the contents of the new Sun StorEdge 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 StorEdge 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 StorEdge EBS resource database files can be caused by a power outage, operating system crash, or manual editing of the database. If the Sun StorEdge EBS server is unable to read the resource files upon startup, messages similar to the following are written to the *daemon.log* file:

```
nsrd: WARNING: NSR configuration database detected
invalid resource /.../00019803aa14713c89456b41

nsrd: Invalid resource saved at
/.../00019803aa14713c89456b41
```

The Sun StorEdge 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 *http://softwaresupport.emc.com* 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 StorEdge EBS environment.

You can also find communication troubleshooting information in *Sun Technical Bulletin 299: IP Naming in Heterogeneous Environments (UNIX | NT | NetWare | Windows 95)*, available at *www.sun.com*.

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, the following message appears:

client_name`: ` *server_name* `cannot request command execution`

After installation, if you want the client to accept backup requests from other Sun StorEdge EBS servers, you must add the name of the additional Sun StorEdge EBS servers to the *servers* file.

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 StorEdge EBS client should contain the following names for a Sun StorEdge EBS server named *mars* in the *jupiter.com* domain:

```
mars
mars.jupiter.com
```

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.

The preferred method of editing the *servers* file is to run the Sun StorEdge EBS Setup program in maintenance mode and edit the Allowed Servers list. For detailed instructions, refer to the installation guide for your operating system.

Unapproved Server Error during Client Setup

If you add a Windows client to a UNIX Sun StorEdge EBS server, and the UNIX server hostname is not included in the Windows client's *servers* file, you might receive the following 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 adding 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 StorEdge EBS server. If aliases are not recognized, the server may be disabled with the following 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*.

The following 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.

UNIX Communication Issues

This section covers communication issues on UNIX networks.

How to Troubleshoot IP Errors

If any of the following 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 StorEdge EBS services failing to start
- Sun StorEdge 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 Sun Technical Support. This enables you to email or fax the steps and error message to Sun.
2. Set up a host table for each Sun StorEdge EBS client and Sun StorEdge EBS server. For more information, see "How to Set Up a Host Table" on page 782.
3. Disable other name servers to simplify testing. For more information, see "How to Disable the DNS Server" on page 783.
4. Use **ping** to establish basic connectivity. For more information, see "How to Use ping to Verify Network Connections" on page 783.
5. Use **rpcinfo** to verify that sessions can be established and that portmapping is correct. See "How to Use rpcinfo to Verify That Sessions Can Be Established" on page 784 for more information.

How to Set Up a Host Table

Troubleshoot IP problems using only the host table. Troubleshooting using only the host table does not mean you cannot use a name service, for example, DNS, with Sun StorEdge EBS software. Run tests 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 StorEdge EBS client, list the client and the Sun StorEdge 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 StorEdge EBS server, list the Sun StorEdge 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 "How to Use ping to Verify Network Connections" on page 783 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.

How to Disable 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 StorEdge EBS clients and the Sun StorEdge 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 StorEdge 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.

How to Set 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
   ```

How to Use ping to Verify Network Connections

After you create the host table, test it using **ping**.

On the Sun StorEdge EBS client, **ping** the following:

- 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

The following example shows how to **ping** the client short name and client long name from a Sun StorEdge EBS client called *mars* in the *jupiter* domain:

ping *mars*

ping *mars.jupiter.com*

On the Sun StorEdge EBS server, **ping** the following:

- 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

How to Use 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, Sun 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 StorEdge 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 entering **rpcinfo** at the command-prompt. For more information, refer to the UNIX man page. 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
```

Replacement for Functional Use of portmapper and rpcbind

The **nsrexecd** process combines the functionality of **portmapper** and **rpcbind**. The **nsrexecd** process starts a **nsrexecd** subprocess that handles the **portmapper** functions. As a result, you see more than one instance of the **nsrexecd** process.

Binding to Server Errors

Sun StorEdge 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 themselves with the registration service provided by the portmapper.

If the Sun StorEdge EBS services are not running and a Sun StorEdge EBS service is requested, the following messages appear in the savegroup completion mail:

```
Server not available
RPC error, no remote program registered
```

These messages indicate that the Sun StorEdge EBS services **nsrd**, **nsrexecd**, **nsrindexd**, **nsrmmd**, and **nsrmmdbd** might not be running.

Saving Remote File Systems

The following 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 106 on page 786 lists the location for the boot-time file. For information about **nsrexecd**, refer to the **nsrexecd** man page.

Table 106. Boot-Time File Locations

Operating System	Boot-Time File
AIX	*/etc/rc.nsr*
HPUX	*/etc/rc*
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*

Windows Communication Issues

This section covers communication issues on Windows networks.

Certified Protocols

This release of Sun StorEdge 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 StorEdge EBS software at this time.

Resolving Connection Problems

This section addresses DNS and local host file-related connection problems and solutions.

During installation, the following 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. The following sections note 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*

  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 the following order:

```
Server A Host Table
127.0.0.1            localhost     loopback
xxx.xxx.xxx.xxx      server_a      server_a.domain.com
xxx.xxx.xxx.xxx      client_a      client_a.domain.com

Client A Host Table
127.0.0.1            localhost     loopback
xxx.xxx.xxx.xxx      client_a      client_a.domain.com
xxx.xxx.xxx.xxx      server_a      server_a.domain.com
```

- Ensure that the hostname is set correctly.

  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 entering the following 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 StorEdge 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 StorEdge EBS Software

If you experience Sun StorEdge EBS software problems, ensure that:

- You are using the **-s server** argument when starting a Sun StorEdge EBS utility from the command-prompt.
- There is a space in the path for the **-s server** argument or the path to the Sun StorEdge EBS program.
- The Sun StorEdge EBS server to which you are connecting is running.

Tips

The following tips can help prevent problems using the Sun StorEdge 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 StorEdge EBS Server Address

When the TCP/IP address changes on the Sun StorEdge EBS server, the Sun StorEdge 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 Sun Customer Service. If you do not reregister the software within 14 days, the Sun StorEdge EBS software stops working.

If you are using DHCP, use a static IP address for the Sun StorEdge EBS server.

New.Net and Sun StorEdge 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 StorEdge EBS software. This causes many Sun StorEdge EBS programs, including **save.exe**, to fail on exit. This is a New.Net problem that Sun StorEdge 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.



Important: Do *not* manually delete the *newdotnet.dll* file. Doing so renders the system unusable.

Troubleshooting Port Configurations

This section provides guidelines for troubleshooting port configurations.

- Cannot bind socket to service port in configured range on system *hostname*

  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*

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 StorEdge EBS portmapper. This conflict occurs because Maestro and Exceed software install a service named **HCL Inetd**. You cannot successfully install the Sun StorEdge 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 StorEdge EBS software (the Maestro and Exceed software continue to run).

If you installed both the Sun StorEdge EBS software and the Hummingbird software and you have difficulties with portmappers, contact Hummingbird Technical Support.

Sun StorEdge 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 StorEdge 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 StorEdge EBS administrator privileges for *root@client_system* in the Administrator attribute in the Server resource. However, be aware that Sun StorEdge 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, which appears as a slash (/) in the Archives list in the **nwretrieve** program.

To retrieve archives separately, archive them separately.

Cloned Archives Do Not Appear in Sun StorEdge EBS Retrieve Program

When you search for an annotation in the **nwretrieve** program, the Archives attribute does not display archive clones.

To locate the clones, start the query without specifying a Search Annotation attribute. If that query returns too many archives, use **mminfo** to locate the archive clone with the same save set ID (ssid) as the archive to be retrieved.

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.

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 the following 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, the following 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 StorEdge EBS services.
2. Remove the environment setting for NSR_MMDCONTROL.
3. Restart the Sun StorEdge 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. For more information, see "Timeout Attribute for Storage Node Operations" on page 105.

Console Error Messages and Corrective Actions

Table 107 on page 793 provides a list of Console error messages or symptoms and corrective actions to take.

Table 107. Error Messages or Symptoms (Part 1 of 5)

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 StorEdge EBS software temporary enabler code has already expired.	Log out, then stop and restart the Console server.

Table 107. Error Messages or Symptoms (Part 2 of 5)

Error Message or Symptom	Possible Cause	Corrective Action
Application window is unresponsive.	Insufficient disk space on the file system where the Console database is installed.	• Ensure that the Console server is running. (See "Console Troubleshooting Notes and Tips" on page 799.) If it is not, close all application windows and check *gstd.log* for errors. • Back up and move the Console database, if necessary. • On a Windows system, run InstallShield with the Repair option to move the database to a different drive.
	Application ran out of memory.	Close all instances of the application and restart it.
	Another dialog box is open in the Console window or Administration window.	Close any open dialog boxes or error messages.
Connection refused: no further information. or Problem contacting server *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 *gstd.log* for errors.
	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. See "Recovering the Console Database" on page 399.

Table 107. Error Messages or Symptoms (Part 3 of 5)

Error Message or Symptom	Possible Cause	Corrective Action
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 *install.log* to determine the HTTP port used by the Console server. See "The Install Log" on page 798.
	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 107. Error Messages or Symptoms (Part 4 of 5)

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 Sun.	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. • If License Manager is installed, but not running, start it. For details, refer to the *License Manager Installation and Administrator's Guide.*
	Sun StorEdge EBS client was stopped, but the License Manager was not stopped, and then the Sun StorEdge EBS client was restarted. Although both services are now running, Sun StorEdge 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 StorEdge EBS software. 2. Stop License Manager if it is running. 3. Restart License Manager. 4. Restart the Sun StorEdge EBS software.
License allocation failed.	Temporary license for Sun StorEdge EBS software is expired.	Enter enabler codes and register the product.

Table 107. Error Messages or Symptoms (Part 5 of 5)

Error Message or Symptom	Possible Cause	Corrective Action
License managed event indicates that license is expiring / expired even though it has been authorized.	License has been authorized within the last 24 hours.	None needed. To remove the managed event from the display, dismiss the event or it will be deleted within 24 hours.
Logging of debug messages has stopped. NOTICE: *alloc /opt:* file system full.	Disk space on the */opt* file system is nearly depleted.	Allocate more disk space.
Event disappears from the Events window.	Another user dismissed it, or the problem that was causing the event no longer exists.	None
Dialog box: "Java Web Start – Download Error" with the message, "Unable to launch Sun StorEdge EBS Management 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.)	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.)
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 the following 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.

Table 108 on page 798 describes the location of the *install.log* file for each Console package.

Table 108. Installation Log File Locations

Package	Solaris Default Location	Windows Default Location
GST Server	*/opt/LGTObgst/logs*	*C:\Program Files\Legato\Management\GST\logs*
Console	*/opt/LGTOcon/logs*	
Administration	*/opt/LGTOnwa/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* is checked. If *gstd.log* has reached its maximum size, the Console server starts a new *gstd.log* file. The *gstd.log* file is placed in the following default locations:

- On Solaris: */opt/LGTObgst/logs*
- On Windows: *%SystemDrive%\Program Files\Legato\Management\GST\logs*

Table 109 on page 799 lists the variables that control the *gstd.log file*.

Table 109. Environment Variables for the GSTD Log

Variable Name	Description
GST_MAXLOGSIZE	Sets the maximum size of *gstd.log* 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 *gstd.log*. Can also be set from the System Options dialog box.

For more information, see "How to Set a System Option" on page 504.

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, check two things:

- *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 StorEdge 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.
- 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 AM) or have **ntp** run as a service and synchronize continuously.

How to Check on a Windows System

On a Windows computer:

- Using the Services window:
  a. From the Start menu, select Control Panel > Administrative Tools > Services.
  b. Verify that **EMC LEGATO GST Service** is running.

How to Check on a Solaris System

- To check whether the **gst** server process is running, enter the following:
  `/usr/bin/ps -ef | grep gstd`
  If the **gst** process is running, a result similar to the following appears:
  `root61401 0  12:54:10 ?0:03 /opt/SUNWebsg/bin/gstd`
- To check whether the database process is running, enter the following:
  `/usr/bin/ps -ef | grep dbsrv`
  If the database server is running, a result similar to the following appears:
  `root61401 0  12:54:10 ?0:03`
  `/opt/SUNWebsg/sybasa/bin/dbsrv9`
- To check whether the web server process is running, enter the following:
  `/usr/bin/ps -ef | grep gsttclsh`
  If the web process is running, a result similar to the following appears:
  `root61401 0  12:54:10 ?0:03 /opt/SUNWebsg/bin/gsttclsh`

Enabling JavaScript

If JavaScript becomes disabled, the Console server will not launch. Check the web browser's settings and re-enable 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.

Appendix A: SYSTEM and VSS SYSTEM Save Sets

This appendix covers the following topics:

- "SYSTEM Save Sets" on page 801
- "VSS SYSTEM Save Sets" on page 807

SYSTEM Save Sets

The SYSTEM save sets discussed in this section work with the following operating systems:

- Windows 2000
- Windows XP Professional
- Windows Server 2003 (with no VSS client license, or with VSS disabled)

Components of the SYSTEM STATE Save Set

As part of the SYSTEM STATE save set, Sun StorEdge EBS software backs up all Windows system state components except the SFP component. For information about how Sun StorEdge EBS software treats the SFP component, see "Components of the SYSTEM FILES Save Set" on page 803.

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 110 on page 802 lists these basic components and references to their backup and recovery procedures.

Table 110. 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)	See "Internet Information Server" on page 845.
Registry	See "Windows Registry" on page 845.
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 the following 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 111 on page 802.

Table 111. SYSTEM STATE Save Set Optional Components (Part 1 of 2)

SYSTEM STATE Optional Component	Special Backup and Restore Considerations
Active Directory (AD)	See "Active Directory" on page 837.
Certificate Server	None

Table 111. SYSTEM STATE Save Set Optional Components (Part 2 of 2)

SYSTEM STATE Optional Component	Special Backup and Restore Considerations
Cluster Server	See "Chapter 19: Cluster Support" on page 587.
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, SFP 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 StorEdge EBS SYSTEM STATE save set, backing up or recovering the system state would be extremely resource-intensive. Therefore, NetWorker software backs up the SFP component in its own save set called SYSTEM FILES.

System File Protection

For a full backup operation, specification of 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.



Important: 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 Sun StorEdge EBS User program, if only the SYSTEM FILES or SYSTEM STATE save set is marked for recovery, a dialog box displays a warning that both of these save sets should be recovered together. No such warning is provided for recoveries performed at the command-prompt.

Not all SFP files necessarily exist on a computer at any given time. Additional SFP files are sometimes installed automatically when a new Windows system component is installed. When this occurs, the new files have a creation date that corresponds to the system component installation date. But the files have a modification date that corresponds to the creation date of the Windows distribution. (The modification date of a new SFP file is usually the same modification date of the already existing SFP files.)

When a backup of save set *All* or the SYSTEM FILES save set detects new SFP files, it checks for the more recent of the file creation and file modification dates. If the more recent date is after the *as of time*, NetWorker software backs up *all* of the system-protected files.

Components of the SYSTEM DB Save Set

The Sun StorEdge EBS SYSTEM DB save set is used to back up the Windows Server 2003, Windows XP Professional, or Windows 2000 system databases that are installed and started. Table 112 on page 805 lists the system databases that are present by default and references to their backup and recovery procedures.

Table 112. SYSTEM DB Save Set Basic Components

SYSTEM DB Basic Component	Special Backup and Restore Considerations
Content Index Server (CIS)	See "Backing Up the Windows Content Index Server" on page 83 and "Restoring a Windows Content Index Server" on page 412.
Disk Quota Database	See "Permissions Requirements for Disk Quota Database Backup" on page 86.
Removable Storage Database	See "Microsoft Removable Storage" on page 697.
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. For information about TSL backup and recovery by using NetWorker software, see "Terminal Server Licensing" on page 847.

Windows Databases Not Included in the SYSTEM DB Save Set

Table 113 on page 806 lists the additional Windows Server 2003, Windows XP Professional, and Windows 2000 databases that NetWorker software supports. Each of these databases may be installed optionally. These databases are backed up as part of the file system, and *not* as part of any SYSTEM save set.

Table 113. Windows Databases Not in SYSTEM DB Save Set

Database	Backup and Restore Procedures
Distributed File System (DFS)	See "Appendix C: Backing Up and Restoring a Microsoft DFS" on page 831.
Encrypting File System (EFS)	See "Encrypting File System" on page 844.
Event logs	See "Event Logs" on page 844.
Sparse files	See "Sparse Files" on page 846.

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.

For more information about NetWorker software support for SPS, see "Chapter 23: Support for Microsoft SharePoint Portal Server" on page 721.

Components of the ASR Save Set

The 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:

- 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.

For more information about NetWorker software support for ASR, see "Chapter 24: Support for Microsoft Automated System Recovery" on page 731.

VSS SYSTEM Save Sets

The VSS SYSTEM save sets discussed in this section work with Windows Server 2003, where VSS is licensed and enabled.

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 NetWorker

software runs. For information about how NetWorker software treats the SFP component, see "Components of the VSS SYSTEM FILESET Save Set" on page 808.

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 StorEdge EBS VSS SYSTEM BOOT save set, backing up or recovering the system state would be extremely resource-intensive. To relieve the strain on resources, NetWorker software backs up the SFP component as part of the VSS SYSTEM FILESET save set. For more information about the SFP component, see "Components of the SYSTEM FILES Save Set" on page 803.

System File Protection

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.



Important: 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. For more information, refer to the *NetWorker Disaster Recovery Guide*.

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 StorEdge EBS backup of save set *All* or the VSS SYSTEM FILESET save set detects new SFP files, NetWorker 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, NetWorker 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. This component of the save set is recovered by using the legacy method of recovery. All other components of this save set are dynamically generated and may differ each time NetWorker software runs.

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 NetWorker 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 NetWorker 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.

A

Appendix B: Firewall Support

This appendix covers the following topics:

- "Firewall Support" on page 811
- "Configuring Firewall Support" on page 816
- "Troubleshooting Port Configuration Problems" on page 823
- "Connections between the Console Server and Clients" on page 824
- "Connections between the Console Server and Sun StorEdge EBS Hosts" on page 826
- "Additional Firewall Support Considerations" on page 828

Firewall Support

This appendix assumes that the user has an understanding of firewall and *Transmission Control Protocol/Internet Protocol (TCP/IP)* protocol basics. This section addresses firewall support specifically in regard to Sun StorEdge EBS software.

Firewall support enables the backup of Sun StorEdge EBS clients and storage nodes that are separated from the Sun StorEdge EBS server by a packet-filtering firewall. A packet-filtering firewall looks at each IP packet entering or leaving the network. The firewall accepts or rejects packets based on user-defined rules.

To enable firewall support, calculate the number of *service ports* (also known as listening ports) used by each Sun StorEdge EBS component to listen for the other Sun StorEdge EBS components' requests. Use the service port calculations to configure the firewall to provide sufficient open service ports for Sun StorEdge EBS software communication. See "Determining the Minimum Number of Service Ports Required" on page 818 for guidelines.

Note: For information about configuring Microsoft Windows XP SP2 Firewall for Sun StorEdge EBS clients, refer to *Sun Technical Bulletin No. 399*. For information about using Windows Server 2003 SP1 Security Configuration wizard with Sun StorEdge EBS clients, refer to *Sun Technical Bulletin No. 410*.

Sun StorEdge EBS Software Configurations with Service Ports

Two common Sun StorEdge EBS software configurations exist:

- Storage node on the same side of the firewall as the Sun StorEdge EBS client (Figure 51 on page 813).
- Storage node on the same side of the firewall as the Sun StorEdge EBS server (Figure 52 on page 815).

Each Sun StorEdge EBS service has one service port, except for **nsrexecd**, which has two fixed service ports: 7937 and 7938.

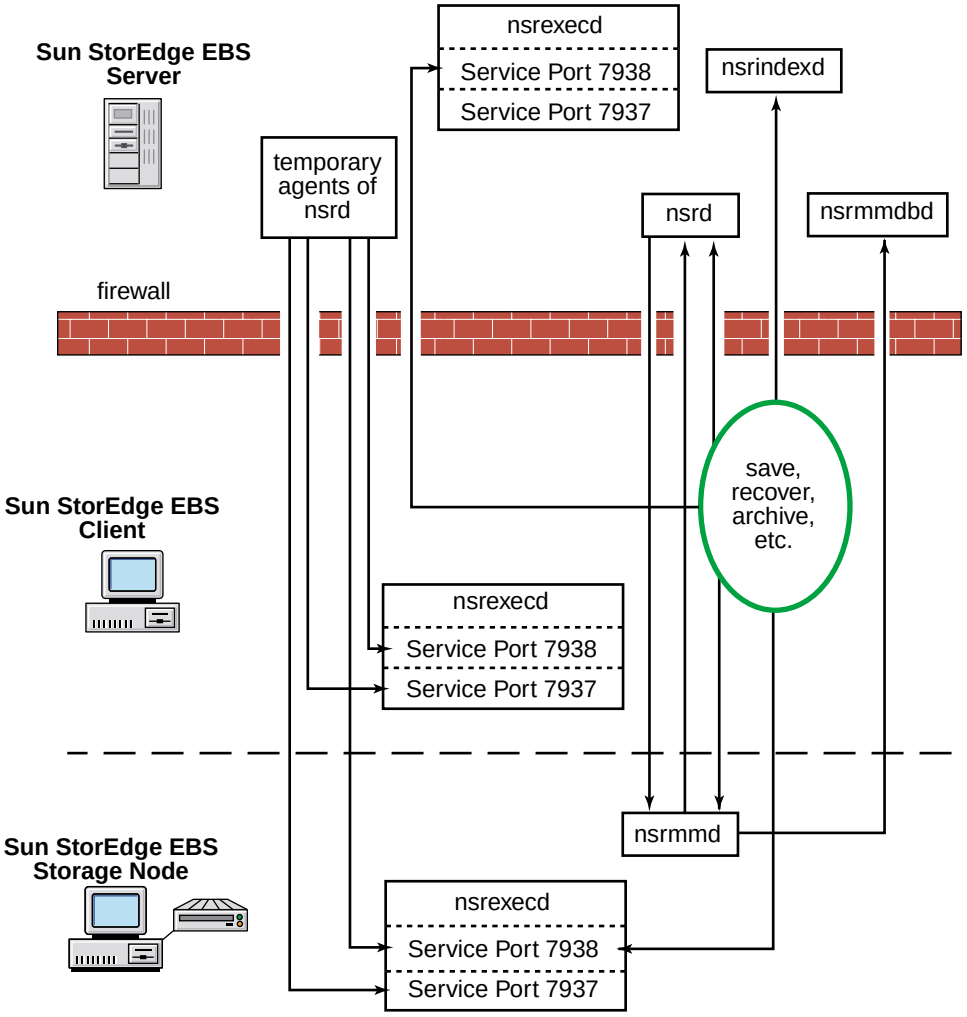
- Port 7937 is the service port used for starting commands.
- Port 7938 is the service port used for the Sun portmapper.

The port numbers used by the services, except for **nsrexecd**, are assigned from the service port range that is set in the Sun StorEdge EBS software. One service port is used for each **nsrmmd** process.

Configuration 1: Storage Node Adjacent to Sun StorEdge EBS Client

Figure 51 on page 813 shows the storage node located next to the Sun StorEdge EBS client. Some users prefer this configuration because it enables a direct connection from the Sun StorEdge EBS client to the storage node without going through the firewall. As a result, large volumes of data can be transmitted more easily.

Figure 51. Storage Node Adjacent to the Sun StorEdge EBS Client



The services that facilitate communication among the three computers are shown with arrows to indicate the direction in which the communication is initiated. Once initiated, TCP/IP communication moves in both directions.

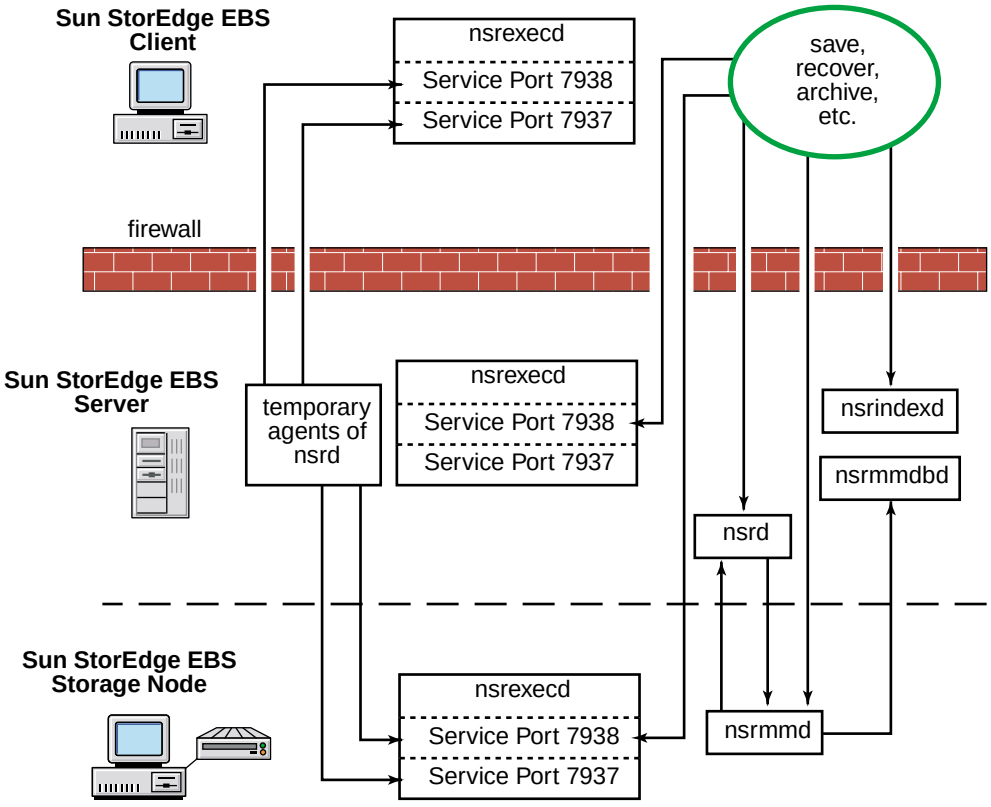
As depicted in Figure 51 on page 813, the service ports must be available. This allows the following:

- The Sun StorEdge EBS client can communicate with the **nsrexecd** (through port 7938), **nsrd**, and **nsrindexd** services on the Sun StorEdge EBS server. These communications are required for archive requests and by such commands as **save** and **recover**.
- The **nsrd** service or agents of **nsrd** on the Sun StorEdge EBS server can communicate with the following services:
  - **nsrexecd** on the Sun StorEdge EBS client (through ports 7937 and 7938).
  - **nsrexecd** on the storage node (through ports 7937 and 7938).
  - **nsrmmd** on the storage node.
- The Sun StorEdge EBS client can communicate with the **nsrexecd** (on port 7938) and **nsrmmd** services on the storage node.

Configuration 2: Storage Node Adjacent to Sun StorEdge EBS Server

Some users choose to put the storage node next to the Sun StorEdge EBS server behind the firewall, to make the storage of backed-up data more secure. This configuration is shown in Figure 52 on page 815.

Figure 52. Storage Node Adjacent to the Sun StorEdge EBS Server



The services that facilitate communication among the three computers are shown with arrows to indicate the direction in which the communication is initiated. Once initiated, TCP/IP communication moves in both directions.

As shown in Figure 52 on page 815, the service ports must be available in order to allow the following:

- The Sun StorEdge EBS client can communicate with the **nsrexecd** (through port 7938), **nsrd**, and **nsrindexd** services on the Sun StorEdge EBS server. These communications are required for commands, such as **save** and **recover**.

- The **nsrd** service, communicate with the following services:
  - **nsrexecd** on the Sun StorEdge EBS client (through ports 7937 and 7938).
  - **nsrexecd** on the storage node (through ports 7937 and 7938).
  - **nsrmmd** on the storage node.
- The Sun StorEdge EBS client can communicate with the **nsrexecd** (through port 7938) and **nsrmmd** services on the storage node.

Note: A third type of Sun StorEdge EBS configuration that does not include a storage node also exists. When configured for that function, the Sun StorEdge EBS server performs the tasks of a storage node. In that case, the **nsrmmd** services run on the Sun StorEdge EBS server, and the other details shown in Figure 52 on page 815 stay the same.

Configuring Firewall Support

To configure the service ports for the Sun StorEdge EBS server, the Sun StorEdge EBS client, and the storage node, do *one* of the following:

- Use the default settings.
- Restrict the number of ports used by the Sun StorEdge EBS software.

Note: If the configured service ports range is not large enough, the associated services and processes cannot communicate through the firewall. For more information, see "Troubleshooting Port Configuration Problems" on page 823.

Configuring Firewall Support with Default Settings

Configure the ports on the firewall using the default settings as follows:

1. Set the firewall rules to allow inbound and outbound TCP packets to and from the service ports.
2. Set the range to 7937-9936, the default range used by Sun StorEdge EBS services.
3. Ensure that Dynamic Network Address Translation (NAT) is either disabled for each Sun StorEdge EBS host or that scheduled backups are disabled for each host that connects through a dynamic NAT firewall. For more information, see "Dynamic Network Address Translation (NAT) Firewall Support" on page 829.

Tips for Configuring Port Ranges

Review the following information before configuring Sun StorEdge EBS port ranges:

- The resource (*nsrla.res*) file stores port ranges that are permitted as a **nsr** port range resource. Do not modify the resource file directly. Instead, define the ranges by using the Sun StorEdge EBS Console as described in "Configuring the Sun StorEdge EBS Port Settings" on page 820.

  The **nsrexecd** service manages the **nsr** port range resources. It must be the first Sun StorEdge EBS service to start, as it does during system initialization.

  Note: If the Sun StorEdge EBS software is started manually, ensure that the **nsrexecd** service is the first one started. Otherwise, ports might be assigned randomly.

- After changing the service port ranges, restart the Sun StorEdge EBS software and make any corresponding modifications to the firewall rules.

  Note: If the port ranges have been restricted, wait briefly before restarting the Sun StorEdge EBS software in order for the ports to become available again.

- Do not change the TIME_WAIT intervals to reduce the port range demand. A change in the TIME_WAIT interval has minimal effect on the port ranges. If the interval is set too low, it could cause the corruption of new processes.
- Static Network Address Translation (NAT) of the Sun StorEdge EBS server IP address may be used. Its use applies only for as long as Sun StorEdge EBS clients can resolve the Sun StorEdge EBS server hostname by using both forward and reverse lookups.
- Refer to *Sun Technical Bulletin No. 388* for considerations regarding the use of stateful inspection firewalls.

Restricting the Port Range

To restrict the port range (rather than use the default settings):

1. Determine the minimum port range required by each Sun StorEdge EBS component. See "Sun StorEdge EBS Server Requirements" on page 818, "Sun StorEdge EBS Client Requirements" on page 819, and "Storage Node Requirements" on page 820, respectively.
2. Configure the Sun StorEdge EBS server port settings. See "Configuring the Sun StorEdge EBS Port Settings" on page 820.

3. Set the firewall rules to allow inbound and outbound TCP packets to and from the service ports configured in the Sun StorEdge EBS software in step 2.
4. Ensure that Dynamic Network Address Translation (NAT) is either disabled for each Sun StorEdge EBS host or that scheduled backups are disabled for each host that connects through a dynamic NAT firewall. For more information, see "Dynamic Network Address Translation (NAT) Firewall Support" on page 829.

Determining the Minimum Number of Service Ports Required

This section describes the minimum service ports required for the Sun StorEdge EBS server, the Sun StorEdge EBS client, and the storage node.

To help calculate the minimum service ports required for the Sun StorEdge EBS software, examine how the ports are used on a computer. Use the **netstat -a** or **rpcinfo -p** command to obtain this information.

Sun StorEdge EBS Server Requirements

To determine the minimum number of service ports required by a Sun StorEdge EBS server, use the following calculation:

9 + 2 x [number of devices] + [number of jukeboxes] + 3 [if the Console server is running on the same host, 0 if not]

Table 114 on page 818 lists the required number of Sun StorEdge EBS server ports:

Table 114. Sun StorEdge EBS Server Port Requirements

Service	Number of Ports
nsrd	1
nsrmmdbd	1
nsrindexd	1
nsrmmgd	1 (only used by the Sun StorEdge EBS server)
nsrjobd	1

Table 114. Sun StorEdge EBS Server Port Requirements

Service	Number of Ports
nsrcpd	1
gstd	3 (one each for RPC, HTTP requests, and reporting) see "Connections between the Console Server and Clients" on page 824

In addition to the port requirements listed, ports for the Sun StorEdge EBS client and devices are needed if they are running on the same host.



Example: Setting the Server's Service Port Range

If the configuration includes four tape devices, 13 service ports must be allocated to the Sun StorEdge EBS server. In this case, set the service port range for the Sun StorEdge EBS server to be 7937-7949, or use the ranges of 7937-7938 (required ports) and 8001-8011 (different port numbers could be used here).

Sun StorEdge EBS Client Requirements

Table 115 on page 819 lists the required number of Sun StorEdge EBS client ports.

Table 115. Client Port Requirements

Service/Feature	Number of Ports
nsrexecd	2 (one each for portmapper and remote execution)
Client Configuration Wizard	1 per module client configuration wizard running
Local Host resources	1 for administering local host resources through the Administration window.
NDMP	1 transient port used when an NDMP filer is backed up

The Sun StorEdge EBS client service, **nsrexecd,** and the Sun portmapper bind locally to service ports 7937 and 7938, respectively.

Storage Node Requirements

To determine the minimum number of service ports required by a Sun StorEdge EBS storage node, use the following calculation:

2 + 2 x [number of devices] + [number of jukeboxes]

This includes the following services and ports:

- 2 is the number of ports required for the **nsrexecd** service (must be ports 7937 and 7938).
- **nsrmmd**: 2 per device (an extra port may be needed when spanning tapes)
- **nsrlcpd**: 1 per media library

In addition to the ports listed, ports for the Sun StorEdge EBS client are needed.



Example: Setting the Storage Node's Service Port Range

If the configuration includes 4 tape devices, 10 service ports must be allocated to the Sun StorEdge EBS storage node. In this case, set the service port range for the Sun StorEdge EBS storage node to be 7937-7946, or use the ranges of 7937-7938 (required ports) and 8001-8008 (different port numbers could be used here).

Configuring the Sun StorEdge EBS Port Settings

Configure the ports for the server, client, and storage nodes by changing the assigned default settings by using the Sun StorEdge EBS Console or **nsrports** program.

Note: The range of ports assigned does not have to be contiguous. To set the ports to noncontiguous ranges, enter spaces between the ranges.

- To use the Sun StorEdge EBS Console to configure the service port ranges:
  1. Set up the appropriate privileges for the NSR system port ranges resource. For more information, see "Access Privileges for Authentication Configuration" on page 572.
  2. Open the Sun StorEdge EBS Administration window and click Configuration.
  3. Right-click a host from the list of Local Hosts and select Configure Port Ranges.
  4. In the General Tab, enter a range of service ports in the Service Ports field. Service ports are those used by Sun StorEdge EBS to provide services.

     To enter noncontiguous port ranges, include a space between the two ranges: for example, 7937-7938 8001-8010.

     Note: The Connection port range of 0-0 should not be changed.
  5. If needed, click the System tab and define a list of administrators who can edit resources for this host.
  6. Click OK to save the changes.
- To use the **nsrports** program to configure the service port ranges, enter the following:

  `nsrports -s` *server* **`[-S|-C]`** *range*

  For example, to define noncontiguous ranges of service ports:

  `nsrports -S 7937-7938 8001-8010`

  Table 116 on page 821 describes the **nsrports** command options.

Table 116. Options for the nsrports Command

Option	Description
-s *server*	Specifies the system to contact.
-S	Sets the system's service ports range to the specified range.
-C	Sets the system's connection ports range to the specified range. By default Sun StorEdge EBS defines a range of 0-0 for connection ports.

B

TCP Packet Considerations

Set firewall rules to allow inbound and outbound TCP packets to and from the following service ports:

- On the Sun StorEdge EBS server — open ports between the Sun StorEdge EBS clients and the Sun StorEdge EBS server.
- On the Sun StorEdge EBS storage node — open ports between the Sun StorEdge EBS clients and the Sun StorEdge EBS storage node.
- On port 7937 — open the port between the Sun StorEdge EBS server and both of the following:
  - Sun StorEdge EBS clients
  - Sun StorEdge EBS storage node on the client side of the firewall
- On port 7938 — open the port for inbound and outbound requests.

Sample rpcinfo Output

The **rpcinfo -p** command can be used to determine port allocations. The sample **rpcinfo** output in Table 117 on page 822 shows port allocations for five devices and the ten services required by the Sun StorEdge EBS server.

- Two ports, 7937 and 7938, are reserved for the **nsrexecd** service. Note that port 7938 is not listed in the table of **rpcinfo** output, because it acts as the portmapper, which contains the information in the table.
- Five ports, 7942 to 7946, are allocated for the five **nsrmmd** services.
- Three ports, 7939 to 7941, are shared among nine other RPC services.

Table 117. Sample of rpcinfo Output (Part 1 of 2)

Program	Version Number	Protocol	Port	Program or Service[a]
390113	1	tcp	7937	**nsrexec**
390103	2	tcp	7939	**nsrd**
390109	2	udp	7939	**nsrstat**
390109	2	tcp	7939	**nsrstat**
390110	1	tcp	7939	**nsrjb**
390107	5, 6	tcp	7940	**nsrmmdbd**

Table 117. Sample of rpcinfo Output (Part 2 of 2)

Program	Version Number	Protocol	Port	Program or Service[a]
390105	5, 6	tcp	7941	**nsrindexd**
390104	105	tcp	7946	**nsrmmd**
390104	205	tcp	7942	**nsrmmd**
390104	305	tcp	7943	**nsrmmd**
390104	405	tcp	7944	**nsrmmd**
390104	505	tcp	7945	**nsrmmd**

a. Programs that register themselves to the portmapper do not have to use their program name as the service name. For example, **nsrexecd** registers its service as the **nsrexec** service rather than the **nsrexecd** service.

Troubleshooting Port Configuration Problems

This section provides guidelines for troubleshooting common error messages that might occur after the ports have been configured.

- Cannot bind socket to service port in configured range on system *hostname.*

  The configured service port range is not large enough for the Sun StorEdge EBS server services to listen for client requests.

  To resolve this problem, do the following:

  – Increase the service port range configured for the specified hostname.

  – Set the firewall rules to allow inbound and outbound TCP packets within the port range specified in the previous choice.

- Remote system error (severity 4, number 12).

  Some ports in the service port range configured to the Sun StorEdge EBS software are not allowed in the firewall rules. (Enough ports are configured in the Sun StorEdge EBS software, but they are not all open in the firewall.)

To resolve this problem, do the following:

– Use the **nsrports** command to check the port range configured in the Sun StorEdge EBS software.

– Determine whether the firewall rules allow inbound and outbound TCP packets between the Sun StorEdge EBS client and server.

B

Connections between the Console Server and Clients

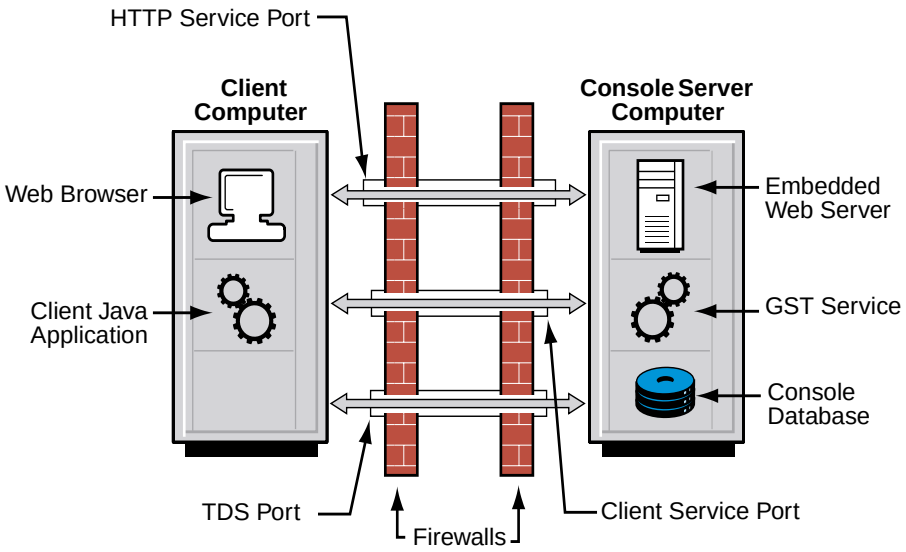
Three ports are necessary for connections between the Console server (**gstd**) and Console clients. Table 118 on page 824 describes the three ports:

Table 118. Console Client/Server Communications Ports

Port	Purpose	Setting on Console Server
HTTP Service (Web Server)	Opens the Welcome screen, downloads the JAR files, and delivers online help pages.	Default 9000 Can be changed during installation of Console server.
Client Service	Logs in to the Console server. Carries status and commands between the Console clients and the GST service on the Console server.	Default 9001 Can be changed during installation of Console server.
TDS (Database Communications)	Carries data from the Console database to the Console clients for displaying reports.	Always 2638 Cannot be changed.

Figure 53 on page 825 illustrates the open ports required for connections between a client and a Console server through a firewall. The firewalls protecting the Console server and the client must be configured to allow communication over these three ports.

Figure 53. Open Ports Required for Communications through a Firewall



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.

Connections between the Console Server and Sun StorEdge EBS Hosts

This section describes considerations for connecting the Console server with Sun StorEdge EBS hosts.

Connections for a Sun StorEdge EBS Server

Because the Console server computer must also be a Sun StorEdge EBS client, it must back up to a Sun StorEdge EBS server. The Sun StorEdge EBS server should be behind the same firewall as the Console server. The ports that need to be opened are described under.

Another possible configuration of the Console server is to install the Sun StorEdge EBS server or storage node software in addition to the Sun StorEdge EBS client software. In either case, there may be additional port requirements to connect to computers outside the firewall.

The Console server connects to Sun StorEdge EBS servers for a number of reasons:

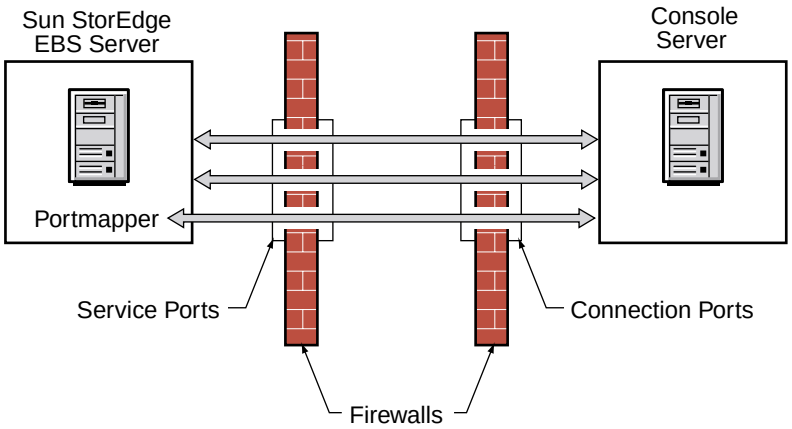
- Monitor Sun StorEdge EBS servers for managed events, and collect data for reports.
- Gather status information for the Administration interface.
- Query Sun StorEdge EBS servers for library changes.
- Send the commands executed from the Administration window to the Sun StorEdge EBS servers.

It is necessary to open ports on the firewall protecting the servers, as follows:

- If a firewall is protecting a Console server, open a set of ports on that server's firewall.
- If a firewall is protecting a Sun StorEdge EBS server, open a set of ports on that server's firewall.
- If separate firewalls protect both servers, open ports on the firewalls protecting each respective server.

Figure 54 on page 827 illustrates a connection to a server through a firewall.

Figure 54. Connecting to a Sun StorEdge EBS Server through a Firewall



How to Configure a Sun StorEdge EBS Server and Its Firewall

Use this procedure to enable a Sun StorEdge EBS server to make connections to a Console server or to Console clients when they are separated by a firewall.

Note: If the Sun StorEdge EBS server is already configured to back up Sun StorEdge EBS clients outside the firewall, and if sufficient ports are available, the following procedure is not necessary:

To configure a Sun StorEdge EBS server to use service ports:

1. Determine the range of service ports needed for the Sun StorEdge EBS server.

   Open a range of ports equal to *ten* plus the number of devices (such as tape, file, or optical) attached to the Sun StorEdge EBS server. The number of devices should equal the number of instances of **nsrmmd** on the computer.

   For example, a Sun StorEdge EBS server connected to a library with four tape devices and one file device requires 15 ports (10 + 4 tape devices + 1 file device).

2. Configure the Sun StorEdge EBS client that is the Sun StorEdge EBS server to use the range of service ports determined in step 1. (Every Sun StorEdge EBS server has itself as a client.)

Note: At this point, communication between the Sun StorEdge EBS server and the Console server is still blocked by a firewall. Use the local **nsradmin** program to configure the Sun StorEdge EBS client.

Choose any range of ports between 1024 and 32767, as long as it does not interfere with any other ports in use. The range 7939–9936 is recommended. Allow for the ports used by **nsrexecd**: 7937–7938. For detailed instructions on how to configure port settings, see "Configuring the Sun StorEdge EBS Port Settings" on page 820.

3. Use the firewall software to accept inbound and outbound TCP/UDP packets from the ranges of ports selected in step 2.
4. Restart the Sun StorEdge EBS server.

Connections for a Sun StorEdge EBS Client

If the Console server is separated from any Sun StorEdge EBS client by a firewall, ensure that ports 7937 and 7938 are opened in the firewall. These ports are required for administering Sun StorEdge EBS client local host resources through the Administration window.

Additional Firewall Support Considerations

This section contains additional considerations for firewall support.

UDP Connections

During startup, the Console server attempts to contact or ping the hosts that it is monitoring for managed events and data collection. The Console server first attempts to use User Datagram Protocol (UDP) to make the connection, then switches to TCP for all connections after that.

The Console server uses UDP to send a remote procedure call (RPC) and asks the host to perform a null procedure. The null procedure causes the host to return a successful result to the Console server.

When the Console server makes a UDP connection, it first attempts to connect by using port 7938, the port reserved for the Sun portmapper. If this connection fails, the Console server tries again by using port 111.



Important: Many network administrators prevent all UDP connections across a firewall because it is possible to make the connection *appear* to come from an authorized user.

The Console server initially pings by using UDP. The response is much faster than with TCP in cases where some of the managed hosts are unavailable.

The UDP connections can be enabled by setting the RPC Ping via the UDP when Connecting to the Sun StorEdge EBS system. Having this option enabled results in a large number of messages. To prevent UDP connections disable this option.

However, preventing UDP connections introduces the risk of a slowdown in the Console software. The Console software must wait for the TCP connection to time out before determining whether the target host is up and running. As a result, a large number of unavailable hosts (due to power loss, bad network connections, or hostname changes, for example) can cause the Console software to start slowly.

Dynamic Network Address Translation (NAT) Firewall Support

Scheduled backups and directed recoveries are not supported for Sun StorEdge EBS clients that connect to the Sun StorEdge EBS server through a dynamic Network Address Translation (NAT) firewall. However, they *are* supported for Sun StorEdge EBS clients that connect through a static NAT firewall.

Disable dynamic NAT firewalls for each Sun StorEdge EBS host, or disable scheduled backups for Sun StorEdge EBS hosts that connect through a dynamic NAT firewall.

To disable dynamic NAT firewall support, refer to the router documentation.

To disable scheduled backups:

1. Open the Client resource for editing. For more information, see "How to Edit a Client" on page 568.
2. Unselect Scheduled Backup.
3. Click OK.

B

Appendix C: Backing Up and Restoring a Microsoft DFS

This appendix covers the following topics:

- "Overview of a Microsoft Distributed File System" on page 831
- "DFS Backups and Restores" on page 832

Overview of a Microsoft Distributed File System

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 StorEdge EBS software does not treat DFS junctions the same as files or directories for backup and recovery. However, for simplicity, DFS junctions appear just like files and directories in the Sun StorEdge EBS 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

A complete Sun StorEdge EBS backup of a DFS requires backups of the following:

- DFS topology information
- DFS junctions
- DFS destination directories (shared directories connected to DFS links)



Important: The Sun StorEdge 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 a part of AD, a component of the SYSTEM STATE or VSS SYSTEM BOOT save set on domain controllers.

Registry-based DFS topology information is backed up as a part of the Windows registry, a component of the DFS host server's SYSTEM STATE or VSS SYSTEM BOOT save set.

For more information about the SYSTEM or VSS SYSTEM save sets, see "Appendix A: SYSTEM and VSS SYSTEM Save Sets" on page 801

Note: For Sun StorEdge EBS software on Windows NT 4.0, use the Microsoft DFS Administrator to save the DFS topology information before initiating a backup. The Windows DFS Administrator does not provide this capability.

How to Configure a Scheduled DFS Backup



Important: To avoid inconsistencies among the various save sets, configure a scheduled backup that includes the DFS topology information, junctions, and destination directories.

Note: For more information about configuring a scheduled backup, see "Chapter 2: Backing Up Data" on page 53.

To configure a scheduled backup for a DFS:

1. From the Administration window, include the following clients in the Sun StorEdge 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 StorEdge EBS group named DFS, then make each of the preceding clients a member of the DFS group.
2. Enter the following save sets in the Save Set attribute 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 the following 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:
```

How to Restore a DFS



Important: The SYSTEM STATE save set, or the VSS SYSTEM BOOT save set and the AD can be restored as a single unit only. For more information, see "Appendix A: SYSTEM and VSS SYSTEM Save Sets" on page 801.

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.

   For recovery procedures, see "Chapter 13: Recovering Data" on page 367.
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.

C

Appendix D: Additional Features of the Microsoft Windows Server

This appendix covers the following topics:

- "Active Directory" on page 837
- "Encrypting File System" on page 844
- "Event Logs" on page 844
- "Internet Information Server" on page 845
- "Windows Registry" on page 845
- "Sparse Files" on page 846
- "Terminal Server Licensing" on page 847
- "Windows Change Journal" on page 848
- "Advanced Configuration and Power Interface" on page 852

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 logon, authentication, directory searches, and access to other shared resources.

Backing Up Active Directory

The Sun StorEdge 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. For more information, refer to the Microsoft Windows Server Resource Kit documentation on AD.

How to Prepare 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, edit the *%SystemRoot% \ system32 \ drivers \ etc \ hosts* file and *%SystemRoot% \ system32 \ drivers \ etc \ lmhosts* file to include the name and IP address of the Sun StorEdge EBS server (the backup server for the computer).
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. For more information, refer to the *NetWorker Disaster Recovery Guide*.)



Important: 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. For information about modifying the *boot.ini* file for multiple-boot configurations, refer to the Microsoft Windows documentation.

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.

How to Perform a Nonauthoritative Restore

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. For more information, see "How to Prepare to Recover the System State on a Domain Controller" on page 838.
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:
     a. SYSTEM FILES or VSS SYSTEM FILESET
     b. SYSTEM STATE or VSS SYSTEM BOOT

     The command used to recover each save set should look similar to the following:

     Sun StorEdge EBS_install_path`\bin\recover.exe -iY` [`-s` *Sun StorEdge EBS_server_name*] `-N "`*saveset_name*`"`

     Note: To avoid use of the Windows version of **recover.exe**, either include <*Sun StorEdge EBS_install_path*> \ *bin* \ *recover.exe* at the command-prompt, or ensure that <*Sun StorEdge EBS_install_path*> \ *bin* is listed before *%SystemRoot%* \ *System32* in the $PATH environment variable.
   - To initiate the operation from the Sun StorEdge EBS 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 "How to Perform a Nonauthoritative Restore" on page 839.
2. Use the **xcopy** command to copy *sysvol\domain* to another location.

   This preserves group policy, and file and directory permissions.
3. Run the Windows **ntdsutil** command-line utility.

   The **ntdsutil** prompt appears. The **ntdsutil** utility is a command interface similar to the Sun StorEdge EBS **recover** interface. For help in the **ntdsutil** utility, enter **?** at the **ntdsutil** prompt.
4. At the **ntdsutil** prompt, enter:

   `NTDSUTIL: authoritative restore`
5. To perform an authoritative restore:
   - Of the entire database, enter:

     `NTDSUTIL: restore database`
   - Of a subtree or individual object, enter:

     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`

   For more information, refer to the Microsoft Windows Server Resource Kit documentation on AD.
6. Exit the **ntdsutil** utility by entering **quit** at each successive **ntdsutil** prompt until the command-prompt appears.
7. Copy the entire restored *SYSVOL\domain* directory and all of its subdirectories to a new location.
8. Restart the domain controller in normal mode.
9. Log in to the domain controller.
10. Wait for the SYSVOL share to be published. This can take several minutes while the recovered domain controller synchronizes with its replication partners.

11. After the SYSVOL share has been published, perform one of the following:
    - For an authoritative restore of the entire database, copy the entire *SYSVOL\domain* directory tree from the new location (see step 7) to the existing *SYSVOL\domain* directory. When the copy operation prompts for confirmation, select "Yes to All."
    - For an authoritative restore of only a part of AD that includes Policy objects, copy the policy folders from the new location (see step 7) to the existing policy folders. Policy objects can be identified by their global universal ID (GUID). Policy folders are located in *SYSVOL\domain\Policies*.

Checking Authoritative Restore Results

To ensure that authoritative restore is working properly, perform the procedures in this section with Windows domain administrator privileges.

How to 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 StorEdge EBS client software on both nodes. If either node is to function as a Sun StorEdge 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 841.
8. Ensure that the authoritative restore worked correctly, as described in "How to Check the Authoritative Restore Results" on page 843.

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, the following 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 the following when backing up or recovering files or folders that are encrypted with EFS:

- Sun StorEdge EBS software will not encrypt or compress a file already encrypted by Windows.
- 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 StorEdge 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. For instructions on encrypting the re-created folder, refer to Microsoft Windows documentation.

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 StorEdge 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.

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 StorEdge 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. For information about creating a registry key to specify an alternate metabase location, refer to the Microsoft documentation. Relocation of the IIS metabase is not supported in IIS version 6.0.



Important: If you perform a Sun StorEdge 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. For information, refer to the Microsoft NNTP documentation.

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 Sun StorEdge EBS User program, the registry is a component

of the SYSTEM STATE or VSS SYSTEM BOOT save set. The Sun StorEdge 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 StorEdge EBS Releases

The Windows NT registry save set is no longer supported in the Sun StorEdge EBS software. In Windows Server 2003, Windows XP Professional, and Windows 2000, the registry is part of the system state. The Sun StorEdge 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.



Important: Attempting to restore a Windows NT 4.0 registry save set to a computer 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*

For more information, refer to Microsoft Knowledge Base article 233427 on the Microsoft website.

Sparse Files

The NTFS sparse files feature enables a program to create huge files without actually committing disk space for every byte. The Sun StorEdge 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, 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, edit the Windows 2000 registry to reflect the TSL changes that were not backed up.

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.

How to Recover 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 entering the number that Microsoft provides.
3. From the Sun StorEdge EBS 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 licenses.

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 StorEdge 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 StorEdge 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*.

Sun StorEdge 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 StorEdge EBS Change Journal Manager.

The Sun StorEdge EBS Change Journal Manager is installed during the Sun StorEdge EBS software setup and can be run from Start>Programs>Sun StorEdge EBS. The Sun StorEdge EBS Change Journal Manager allows you to:

- Enable or disable the Change Journal for each NTFS 5.0 volume.
- Enable or disable the Sun StorEdge EBS software's use of each volume's Change Journal.
- Set parameters that control the size of the Change Journal log file.

How Sun StorEdge EBS Software Uses the Change Journal

When configured to use the Change Journal, the Sun StorEdge 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. For more information about level and incremental backups, see "Backup Levels" on page 154.



Important: The Sun StorEdge 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 StorEdge 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 StorEdge EBS **save**

command's **-t** argument. For information about the Sun StorEdge EBS **save** command, refer to the *NetWorker Command Reference Guide* or the UNIX man page.

Sun StorEdge EBS Save Criteria When Using the Change Journal

A Sun StorEdge 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 StorEdge 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 119 on page 850 shows the Sun StorEdge EBS save triggers when using the Change Journal.

Table 119. Sun StorEdge EBS Save Triggers (Part 1 of 2)

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.

Table 119. Sun StorEdge EBS Save Triggers (Part 2 of 2)

Change	Trigger
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 StorEdge EBS Software to Use the Change Journal

The Sun StorEdge EBS Change Journal Manager enables one to view or edit the Change Journal configuration of each volume in the Sun StorEdge EBS server or client host computer.

How to Configure Sun StorEdge EBS Software to Use the Change Journal

To view or edit the Change Journal configuration:

1. On the Sun StorEdge EBS server or client host computer, select Start>Programs>Sun StorEdge EBS>Sun StorEdge EBS Change Journal Manager.
2. In the Sun StorEdge 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.

   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 StorEdge EBS software, ensure that the Sun StorEdge EBS Uses Change Journal checkbox is selected. To disable the Change Journal, ensure that this checkbox is clear.

   If you select All NTFS Volumes, and some (but not all) of the local NTFS 5.0 volumes already have this option enabled, the checkbox appears shaded. To change the setting of this option for all NTFS 5.0 volumes, click the checkbox until it is selected or clear, but not shaded.

4. To enable the Change Journal for use with Microsoft, select the Enable Change Journal For Selected Volumes checkbox. To disable the Change Journal, ensure that this checkbox is clear.

   If you select All NTFS Volumes, and some (but not all) of the local NTFS 5.0 volumes already have this option enabled, the checkbox appears shaded. To change the setting of this option for all NTFS 5.0 volumes, click this box until it is selected or clear, but not dimmed.

   Note: When you enable the Change Journal, there may be a delay of several minutes before logging begins.

5. Edit the following values to control the size of the Change Journal log file. If the Change Journal is already enabled for the selected volume or volumes, disable it before changing either of these values.

   - % Of Volume For Log File — The maximum amount of the volume's storage space that can be used for the Change Journal log file. The allowable range is 0.01% to 2.0% of the volume's capacity.
   - % Of Log For Allocation Delta — The amount by which the Change Journal log file can expand if additional space is needed. This is also the amount that will be purged from the beginning of the log when the file has reached its maximum size. The allowable range is 12% to 25% of the % Of Volume For Log File.

6. Click Apply to save any configuration changes.
7. Click OK to exit Sun StorEdge EBS Change Journal Manager.

Advanced Configuration and Power Interface

Sun StorEdge EBS software supports the Windows Server 2003, Windows XP Professional, and Windows 2000 Microsoft Advanced Configuration and Power Interface (ACPI), which is also called OnNow.

Sun StorEdge EBS Support for ACPI

Support for ACPI is provided by the *Sun StorEdge EBS Power Monitor* service. The executable for this service (<*Sun StorEdge EBS_install_path*> \ *bin* \ *nsrpm.exe*) is installed and configured for automatic startup during Sun StorEdge EBS setup.

For scheduled backups of Windows Server 2003, Windows XP Professional, and Windows 2000 Sun StorEdge EBS clients, each client must be able to respond when a Sun StorEdge EBS server contacts it for backup. Therefore, by

default, the Sun StorEdge EBS Power Monitor service does not allow a Sun StorEdge EBS server or client host operating system and its network interface to enter the ACPI standby mode.

However, the Sun StorEdge 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 StorEdge EBS software does not prevent the host's power-management policies from forcing the system to power down. The Sun StorEdge 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 StorEdge 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 StorEdge EBS client or server host in standby or hibernation mode *while* a Sun StorEdge EBS backup or restore operation (involving that host) is in progress.

  Doing so will yield unpredictable results.
- If a Sun StorEdge EBS server is powered down while Sun StorEdge 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 StorEdge EBS's tape processes will have incorrect positioning information.

D

Appendix E: UNIX and Linux Platform-Specific Notes

This appendix covers the following topics:

- "Solaris" on page 855
- "Linux" on page 856
- "HP-UX" on page 857
- "AIX" on page 867
- "IRIX" on page 868

Solaris

This section provides information specific to Sun StorEdge EBS software that runs on the Solaris platform.

Sun StorEdge EBS Executables Not Found in Expected Location for Solaris Client

On Solaris, Sun StorEdge EBS executables are installed by default in */usr/sbin/nsr*. If you start a group backup on a Sun StorEdge 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 StorEdge EBS executables in */usr/sbin/nsr*. This is because the **savefs** command is not in the search path.

To solve this issue, set the hidden attribute Executable Path for the client.

How to Set the Executable Path

To set the Executable Path:

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 StorEdge 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 StorEdge 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.

Linux

This section provides information about Sun StorEdge EBS software that runs on the Linux platform.

Backup Consideration 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. For more information on configuring the Linux SCSI subsystem, refer to *The Linux Documentation Project* website. 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 StorEdge 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 StorEdge EBS server 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
```

For more information on creating devices, refer to the OS vendor documentation.

Linux Journaled File System Support

Backup and recovery operations are supported on the following Linux journaled file systems:

- ext3
- reiserfs
- jfs
- xfs



Important: 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. For information about directives, see "Chapter 6: Directives" on page 185.

HP-UX

This section provides information specific to Sun StorEdge 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:

- Log in as root on the Sun StorEdge EBS server or storage node and enter the **ioscan -f** command.

  Use a SCSI address within the range of 0 to 6. The primary hard disk is usually on SCSI address 6.



Important: 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.

E

Installing the SCSI Pass-Through Driver

There are two procedures for installing a SCSI pass-through (spt) driver, depending on what hardware is installed:

- For a SCSI pass-through driver on an HP 9000 Series 800 containing an HP-PB interface card, see "How to Install an HP-PB or NIO Pass-Through Driver" on page 858.
- For a SCSI pass-through driver on an HP 9000 Series 700 or 800 containing an HSC or PCI interface card, see "How to Install a GSC, HSC, or PCI Pass-Through Driver" on page 861.

How to Install an HP-PB or NIO Pass-Through Driver

For more information about the installation of an HP-PB or NIO Pass-Through Driver, refer to the HP-UX man page for **scsi_pt**.

To install the HP-PB or NIO pass-through driver:

1. Enter the following command to get more information about the peripherals attached to the system:

   ioscan -f

   A screen output similar to the following appears:

```
Class        I  H/W PathDriverS/W StateH/W TypeDescription
ext_bus      7  4/8c720CLAIMEDINTERFACEGSC add-on
                           FAST/WIDE SCSI
target       4  4/8.1tgtCLAIMEDDEVICE
tape         6  4/8.1.0stapeCLAIMEDDEVICESTK 4890
ext_bus      8  8/4/4scsi3CLAIMEDINTERFACEHP 28696A-wide
                           SCSE ID=7
target       5  8/4/4.0targetCLAIMEDDEVICE
tape         7  8/4/4.0.0tape2CLAIMEDDEVICEQuantum DLT4000
ext_bus      12 10/4/4scsi3CLAIMEDINTERFACEHP 28696A-Wide
                           SCSI ID=7
target       6  10/4/4.6targetCLAIMEDDEVICE
unknown      -1 10/4/4.6.0UNCLAIMEDUNKNOWNHP C6280-7000
```

   The screen output shows:

   - An HSC fast / wide differential SCSI at instance 7, which is connected to an STK 4890 tape drive at SCSI address 1.
   - An HP-PB fast / wide differential SCSI adapter at instance 8, which is connected to a Quantum DLT4000 tape driver at SCSI address 0.
   - An HP-PB fast / wide differential SCSI adapter at instance 12, which is connected to the robotics controller for an HP C6280-7000 at address 6.

   The robotics controller is listed as UNCLAIMED / UNKNOWN, because the HP-PB pass-through has not been configured to communicate with it.

2. To complete the configuration, go to the build area:

   cd /stand/build

3. Create a system from the existing kernel:

   /usr/lbin/sysadm/system_prep -s *system*

   Check whether the SCSI pass-through driver is already configured into the current kernel. The driver is represented by the name *spt*:

   grep spt system

4. If no *spt* line prints, add a driver statement to the SCSI drivers area of the system file, for example, */stand/build/system*:

   spt

5. Append a driver line to the system file, as follows:

   driver *path* **spt**

   where *path* is the complete hardware path of the device (from the **ioscan** command).

   The previous **ioscan** display with the library robotics connected to the hardware path *10/4/4.6.0* would read:

   ```
   driver 10/4/4.6.0
   ```

   A driver statement is required because the system cannot autoconfigure the SCSI pass-through driver. One statement is required for the robotics controller on each library.

   Note: Adding a line such as **driver** *path* **spt** may prevent the system from recognizing tape drives added to that SCSI bus. You may need to add lines for **driver** *path* **spt** for each tape drive that the system was previously able to recognize.

6. Save the old system file:

   mv /stand/system /stand/system.prev

7. Move the new system to the current one:

   mv ./system /stand/system

8. Build a new kernel:

   /usr/sbin/mk_kernel

9. Verify that the pass-through driver has been built into the new kernel:

   what ./vmunix_test | grep scsi_pt

   A line similar to the following appears:

   ```
   scsi_pt.c Revision: 10.20 $Date: 06/07/18 14:38:27 $
   ```

10. Save the old kernel:

    mv /stand/vmunix /stand/vmunix.prev

11. Move the new kernel to be the current one:

    mv /vmunix_test /stand/vmunix

12. Boot the system from the new kernel:

    /usr/sbin/shutdown -r 0

13. Verify the peripheral configuration. After the system reboots, enter the following command:

    ioscan -f -C spt

    Output from the **ioscan** should resemble the following:

```
Class   I    H/W PathDriverS/W StateH/W TypeDescription
spt     0    10/4/4.6.0SPTCLAIMEDdeviceHP C6280-7000
```

14. Proceed to "How to Create a Device File" on page 862.

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 [Return].
3. Select Drivers and press [Return].
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 Create a Device File" on page 862.

   Note: 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 [Return].
6. From the Actions menu, select Create a New Kernel and press [Return].
7. When prompted with “Are you sure?” indicate Yes, and press [Return].
8. The Creating Kernel message appears, followed by the Move Kernel Message. Select OK and press [Return].

   The system reboots.
9. Proceed to "How to Create a Device File" on page 862.

Installing the Autochanger Hardware

To install the autochanger hardware, follow the instructions in the vendor's hardware installation manuals. Use the SCSI addresses that were selected in the section, "Selecting SCSI Addresses for the Autochanger" on page 858, to set any required switches on the autochanger for the devices and robotics controller.

How to Shut Down the System

Before installing the autochanger device hardware, shut down the system.

To shut down the server:

1. Halt the server by typing the following command:

   sync

   shutdown -h *time*

   where *time* is the grace period, in seconds, for users to log off before the system shuts down. The default grace period is 60 seconds.

2. Turn off the power to the server.

How to Start the System

To start the system:

1. Turn the main power switch of the autochanger to the On position.
2. Turn on the power to the server.

   The server reboots.

How to Create a Device File

To create 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.

4. If the device entry is not defined, create a device file. Use the **mknod** command to create a device file to access the robotics controller. For example:

   `/user/sbin/mknod /dev/sji` *devfilename* `c` *majornum minornum*

   - *devfilename* is the user-defined name of the device file.
   - *majornum* is the major number from the **lsdev** command.
   - *minornum* is the minor number. It must be entered in the following format:

     0xIITL00

     where:

     –*II* is the two-digit card instance number in hexadecimal (from **ioscan** command).

     –*T* is the largest SCSI ID number.

     –*L* is the LUN number.

     –*00* is reserved.

Major Number

To determine the value for *majornum*, enter one of the following commands:

- If the robotics is attached to an HP-PB adapter:

  lsdev -d spt

- If the robotics is attached to an HSC or PCI adapter:

  lsdev -d sctl

The output should resemble the following. The assigned number may differ from those displayed in this example:

```
      CharacterBlock DriverClass
HP-PB75-1 sptspt
HSC or PCI203-1sctlctl
```

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). For example, if the robotics is attached to an HP-PB adapter card, the **ioscan** should resemble the following:

```
Class     I      H/W PathDriverS/W StateH/W TypeDescription
ext_bus   12     10/4/4scsi3CLAIMEDINTERRACEHP 28696-Wide
                          SCSI ID=7
target    6      10/4/4.0targetCLAIMEDDEVICE
spt       0      10/4/4.6.0sptCLAIMEDDEVICEHP C6280-7000
```

Given the information in this example, the **mknod** command would be:

/usr/sbin/mknod /dev/sji *dev_filename* **c 75 0x0c6000**

In the case of an HSC or PCI adapter, the **ioscan** output varies from the previous output in one of two ways.

If the *schgr* driver is configured on the system, it appears associated with the library. The **ioscan** output line resembles:

```
Class     I      H/W PathDriverS/W StateH/W TypeDescription
spt       0      10/4/4.6.0schgrCLAIMEDDEVICEHP 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 PathDriverS/W StateH/W TypeDescription
unknown   -1     10/4/4.6.0schgrUNCLAIMEDDEVICEHP 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

The following is an example of the output from this command:

```
scsidev@0.1.0:HP      C1194F          0.14|Autochanger
(Jukebox), /dev/rac/c0t1d0

scsidev@0.2.0:Quantum DLT4000         CC37|Tape,
```

```
/dev/rmt/c0t2d0BESTnb
scsidev@0.3.0:Quantum DLT4000         CC37|Tape,
/dev/rmt/c0t3d0BESTnb
scsidev@0.4.0:Quantum DLT4000         CC37|Tape,
/dev/rmt/c0t4d0BESTnb
scsidev@0.5.0:Quantum DLT4000         CC37|Tape,
/dev/rmt/c0t5d0BESTnb
```

How to Test the Autochanger Connection

To test the autochanger connection, run the **jbexercise** program with at least two pieces of "scratch" media loaded in the first and last slots of the autochanger. The devices should be empty, with the doors open. For complete instructions, refer to the **jbexercise** man page.

When running **jbexercise**, specify the control port and the autochanger model type. The control port for SCSI autochanger models is typically */dev/sjid1u1* (or */dev/sjid1u2* for a second autochanger). To run the **jbexercise** program on an Exabyte 10i autochanger, use the following command:

```
jbexercise -c /dev/sjid1u1 -m EXABYTE_10I
```

"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 StorEdge EBS */nsr/logs/messages* file. If this error message appears, do the following:

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 from *daemon.notice* to *local 7.notice*.
3. Apply and save the changes to the Log Default notification.
4. Enter the following 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. 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 StorEdge 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.

Note: On UNIX systems, when you use the **nsrjb** command with the **-f** *devicename* attribute and the device name contains backslashes, use the backslash "escape" character with the backslash. A device with the name *rd=sn1:*\\.*Tape0* would be entered as *rd=sn1:*\\\\.*Tape0*.

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 StorEdge 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 a Sun StorEdge EBS motif-based graphical interface is being run (**nwbackup**, **nwrecover**, **nwarchive**, or **nwretrieve**) 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 execute the following command:

```
export LC_CTYPE=C.iso88591
```

The Sun StorEdge EBS motif-based graphical interfaces 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 StorEdge 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 StorEdge EBS storage node software running on the IRIX platform.

IRIX Displays “No Space Left in Tape” Message

If an IRIX Sun StorEdge 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 StorEdge 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.

Appendix F: Mac OS X Support

This appendix covers the following topics:

- "Support for Mac OS X" on page 869
- "Mac OS X Backup Considerations" on page 870
- "Recovering Files and Directories on Mac OS X" on page 873

Support for Mac OS X

This section describes Sun StorEdge EBS client support for the Mac OS X platform. Mac hosts can be set up as Sun StorEdge EBS clients by using the NetWorker ClientPak for Mac OS X. Mac hosts that are set up as Sun StorEdge EBS clients can be backed up and restored by using any supported Sun StorEdge EBS server on Solaris. Currently the Sun StorEdge EBS Server for Mac OS X is not supported.

Mac OS X Metadata Support

The Sun StorEdge 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 or later)
- Access Control Lists (Mac OS X 10.4 or later)

Supported File Systems

The Sun StorEdge EBS client for Mac OS X software supports the following file systems:

- HFS+ (including journaled)
- HFS
- UFS

Known Limitations

The following are known limitations for this release of the NetWorker ClientPak:

- There is no graphical user interface, however there is full support of the command-line interface.
- Sun StorEdge EBS for Mac OS X is not localized for non-English environments.

Mac OS X Backup Considerations

Use this section to help plan successful backups for Sun StorEdge EBS clients on the Mac OS X platform.

Scheduling a Sun StorEdge EBS Client Backup on Mac OS X

This section provides information on configuring backups for a Sun StorEdge 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. For information about creating a Client resource, see "Task 7: Create a Backup Client Resource" on page 61.
2. Select one of the following directives from the Directive attribute list:
   - Mac OS Standard Directives
   - Mac OS with Compression Directives

   For more information about Mac OS directives, see "Preconfigured Global Directive Resources" on page 192.
3. Click OK.

For more information about scheduling a backup, see "Scheduled Backups" on page 53.

Backing Up Mac OS X Server's Open Directory for Disaster Recovery

This section describes how to backup the Mac OS X Server's Open Directory. Open Directory contains system configuration information that is essential for disaster recovery. For more information about disaster recovery, refer to the *NetWorker Disaster Recovery Guide*.

The Sun StorEdge 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, use the **savepnpc** script to automatically export and backup the required Open Directory database files. 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 StorEdge EBS Client resource.

   For more information about enabling the Client resource to use the **savepnpc** command, see "Using the savepnpc Command with a Customized Backup Program" on page 94.
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 the following 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: Custom savepnpc Script for Mac OS X

A Mac OS X Sun StorEdge EBS client that belongs to the Default group will have a */nsr/res/Default.res* script with the following 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.

F

Performing an Ad Hoc Backup on Mac OS X

Ad hoc backups for Mac OS X clients must be performed from the command-prompt. To perform an ad hoc 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 first Sun StorEdge EBS server defined in the */nsr/res/servers* file. To specify an alternative Sun StorEdge EBS server, use the **save** command with the **-s** *Sun StorEdge EBS_server* option.

Recovering Files and Directories on Mac OS X

The following sections provide information on recovering individual files and directories from a Sun StorEdge EBS client on Mac OS X:

- "Task 1: Browse Backed-Up Mac OS X Data" on page 873
- "Task 2: Recover Individual Files or Directories" on page 873

Each task in this section uses the Sun StorEdge EBS **recover** command. For more information about this command, refer to the *NetWorker Command Reference Guide* or the UNIX man pages.

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 StorEdge EBS server by using the following command:

   $ recover

   By default, the **recover** command contacts the first Sun StorEdge EBS server defined in the */nsr/res/servers* file. To specify an alternative Sun StorEdge EBS server, use the **recover** command with the **-s** *Sun StorEdge EBS_server* option.

2. At the recover prompt, browse backed-up Mac OS X data 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, 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 entering the following command:

   recover> **recover**

F



Important: 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*.

Glossary

This glossary provides definitions for terms used in this guide.

Access Control List (ACL)	A list that specifies the permissions assigned to a specific file or directory.
ACL	See *Access Control List (ACL)*.
active group	A Sun StorEdge EBS backup group that has its Autostart attribute enabled.
ad hoc backup	A backup that a user requests from the client. The user specifies the files, file systems, and directories to backup.
administrator	The person who installs, configures, and maintains the Sun StorEdge EBS software.
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 logical server (HP TruCluster), 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.

Application Specific Module (ASM)	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 StorEdge EBS directive used to compress files.
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 StorEdge 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 volume* or a *clone volume*.
ASM	See *Application Specific Module (ASM)*.
attribute	A feature (property) of a *resource*.
authorization code	A unique code that unlocks the Sun StorEdge 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.
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 volume* or a *clone volume*.

base enabler code — See *enabler code*.

bootstrap — A save set that is essential for Sun StorEdge EBS disaster recovery. A bootstrap is composed of two components that reside on the Sun StorEdge EBS server: the *media database*, and the *resource database*. The bootstrap also provides information that is essential for recovering *client file indexes*.

browse policy — A policy that determines how long backed up data is available for quick access. See also *retention policy*.

canned report — A preconfigured report that can be tailored by the user.

carousel — See *autochanger*.

client — A computer, workstation, or fileserver whose data can be backed up or recovered. See also *Client resource*.

client file index — A database that tracks every database object, file, or *file system* that is backed up. The Sun StorEdge EBS server maintains a single client index file for each client.

client-initiated backup — See *ad hoc 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 an *archive volume* or a *backup volume*.

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.
Common Internet File System (CIF)	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 StorEdge EBS processes use to perform backup and recovery sessions through a firewall.
Console server	Sun StorEdge EBS servers and clients are managed from the Sun StorEdge EBS Console server. The Console server also provides reporting and monitoring capabilities for all Sun StorEdge EBS 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 StorEdge EBS software.
daemon	A process that runs in the background and performs a specified operation at predefined times or in response to certain events.
DAR	See *Direct Access Restore (DAR)*.
Data Management Application (DMA)	Initiates a backup or recovery request through the NDMP connection.
data retention policy	See *retention policy*.
Data Server Agent (DSA)	An *agent save* or *agent recover* between a Sun StorEdge EBS server and a non-Sun StorEdge EBS client. An example of a DSA is an NDMP host that generates proprietary save data and sends that data to a Sun StorEdge EBS storage device to have a save set associated to it.
Data Service Provider (DSP)	Controls access to disk storage during an NDMP back up.

datawheel — See *autochanger*.

datazone — A group of computers administered by a Sun StorEdge EBS server.

DDS — See *Dynamic Drive Sharing (DDS)*.

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*.

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 StorEdge EBS libraries.

DFS — See *Distributed File System (DFS)*.

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

Direct Access Restore (DAR) — 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.

Distributed File System (DFS) — A Microsoft Windows add-on that allows you to create a logical directory of shared directories that span multiple machines across a network.

DMA — See *Data Management Application (DMA)*.

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 also *device*.
DSA	See *Data Server Agent (DSA)*.
DSA save set	Save sets of an NDMP client that are backed up to non-NDMP tape device.
DSP	See *Data Service Provider (DSP)*.
Dynamic Drive Sharing (DDS)	A feature that allows Sun StorEdge EBS software to recognize shared drives.
Sun User SCSI (LUS)	A proprietary device driver that sends arbitrary SCSI commands to an autochanger.
enabler code	A code that allow you to run Sun StorEdge 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.
failover	Relocating a cluster resource to its redundant (backup) component, either because of a hardware or software failure or for administrative purposes.
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*.
Generic Services Toolkit (GST)	Provides the underlying framework of the Console server.
groom	To removing the original files from a disk after an archive.
group	A client or group of clients configured to back up files to the Sun StorEdge EBS server at a designated time of day.
GST	See *Generic Services Toolkit (GST)*.
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 ID	A serial number that uniquely identifies a computer.

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 StorEdge 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.

JAR See *Java Archive (JAR)*.

Java A high-level programming language. The *Java Virtual Machine (JVM)* enables the same, unmodified, Java program to run on most computer operating systems.

Java Archive (JAR) A file that contains compressed components needed for a Java applet or application.

Java Plug-in A JVM that can be used by a web browser to run Java applets.

Java Virtual Machine (JVM) An execution environment for interpreting the Java programming language. Each operating system runs a unique JVM to interpret Java code.

jukebox See *autochanger*.

JVM See *Java Virtual Machine (JVM)*.

label A Sun StorEdge EBS assigned label that uniquely identifies a volume. Templates can be used to define label parameters.

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.

License Manager (LLM) — An application that provides a central point for managing product licenses.

LLM — See *License Manager (LLM)*.

local cluster client — A Sun StorEdge 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 StorEdge EBS device used in the integration of Sun StorEdge EBS software with SmartMedia or AlphaStor. SmartMedia or AlphaStor assigns the logical device to an actual device.

logical server — A term used by HP TruCluster to denote a cluster server. Also known as an agent (Sun), a package (HP-UX), and a virtual server (Microsoft).

low-water mark — The percentage of disk space filled that, when reached, automatically stops the staging process.

LUS See *Sun User SCSI (LUS)*.

managed application A program that can be monitored and/or administered from the Console server.

manual backup See *ad hoc backup*.

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 StorEdge EBS server manages. See also *volume*.

media index See *media database*.

member A physical host, such as the HP TruCluster Server, 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.

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.

NDMP See *Network Data Management Protocol (NDMP)*.

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 data streams 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 *autochanger*.

Network Data Management Protocol (NDMP) — A TCP/IP-based protocol that specifies how heterogeneous network components communicate for the purposes of backup and recovery.

Network File System (NFS) — A client/server application that allows users to access shared files stored on networked computers of different types.

Sun StorEdge EBS administrator — A user who can add to or change the configuration of the Sun StorEdge EBS server, media devices, and libraries. Sun StorEdge EBS administrators must have their usernames included in the Sun StorEdge EBS server Administrator list.

Sun StorEdge EBS client — See *client*.

Sun StorEdge EBS Console server — See *Console server*.

Sun StorEdge EBS Management Console — See *Console server*.

Sun StorEdge EBS server — The computer running the Sun StorEdge EBS server software, which contains the *online indexes* and provides backup and recovery services to the clients on the same network.

Sun StorEdge EBS storage node — See *storage node*.

NFS — See *Network File System (NFS)*.

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 StorEdge EBS administrator about important Sun StorEdge EBS events.

nsrd — The master Sun StorEdge EBS server process.

nsrhost — The logical hostname of the Sun StorEdge EBS server.

online indexes — Databases on the Sun StorEdge EBS server that contain information about client backups (*client file index*) and backup volumes (*media database*).

operator — The person who monitors the server status, loads backup volumes into storage devices, and executes day-to-day Sun StorEdge EBS tasks.

override — A backup level that takes place instead of the scheduled one.

package — A term used by HP-UX to denote a cluster server. Also known as an agent (Sun), logical server (HP TruCluster), and 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 also *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 StorEdge EBS host that is involved in an authentication process with another Sun StorEdge EBS host.

permanent enabler A code that allows the software to be authorized for permanent use. See also *temporary enabler* and *authorization code*.

physical cluster client A Sun StorEdge 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 Sun technology that provides point-in-time snapshots of data. Sun StorEdge 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.

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.

Remote Procedure Call (RPC) — The protocol that the Sun StorEdge EBS server uses to perform client requests over a network.

requestor — A VSS-aware application that creates and destroys a *shadow copy*. Sun StorEdge EBS software is a requestor.

resource — A component that describes the Sun StorEdge EBS server or its clients. Clients, devices, schedules, groups, and policies are all Sun StorEdge 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 also *browse policy*.

retrieve — To locate and recover archived files and directories.

retry mechanism — The action Sun StorEdge EBS software performs when client operations fail. This situation might occur when the rate of transmission is either low or nonexistent.

root — 1. (UNIX only) The UNIX superuser account.

2. (Microsoft Windows and UNIX) The highest level of the system directory structure.

RPC — See *Remote Procedure Call (RPC)*.

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 also *level*.

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 StorEdge EBS server* and *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 also *Volume Shadow Copy Service (VSS)*.

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.

Simple Network Management Protocol (SNMP)	A protocol used to send messages to the administrator about Sun StorEdge EBS events.
skip	A backup *level* in which designated files are not backed up.
snapshot	A point-in-time copy of data created during an instant backup.
snapshot policy	A Sun StorEdge EBS resource used to configure backups using the Sun StorEdge EBS PowerSnap Module software.
SNMP	See *Simple Network Management Protocol (SNMP)*.
ssid	See *save set ID*.
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, if the *Sun StorEdge EBS.clustersvr* file is missing, the Sun StorEdge EBS server starts in noncluster (stand-alone) mode.
stand-alone device	A storage device that contains a single drive for backing up data.
storage device	The hardware that reads and writes data during backup, recovery, or other Sun StorEdge EBS operations.
storage node	A storage device physically attached to another computer whose backup operations are controlled by the Sun StorEdge EBS server.
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.
TCP/IP	See *Transmission Control Protocol/Internet Protocol (TCP/IP)*.

temporary enabler — A code that allows operation of the software for an additional 45 days beyond the evaluation period. See also *enabler code* and *permanent enabler*.

Transmission Control Protocol/Internet Protocol (TCP/IP) — 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.

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 StorEdge 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 alias — The username seen by the Sun StorEdge EBS server when a Console user connects to the Sun StorEdge EBS server.

user groups — Used to assign users to access control groups and to configure the privileges associated with those groups.

versions — The date-stamped collection of available backups for any single file.

virtual cluster client — A Sun StorEdge 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), a logical server (HP TruCluster), and 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 StorEdge 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 also *label*.

volume pool — See *pool*.

Volume Shadow Copy Service (VSS) — Microsoft technology that creates a point-in-time snapshot of a disk volume. NetWorker software backs up data from the snapshot. This allows applications to continue to write data during the backup operation, and ensures that open files are not omitted

VSS — See *Volume Shadow Copy Service (VSS)*.

VSS component — A subordinate unit of a *writer*.

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.

Index

URL
See universal resource locator.

A

aborting a recover 766
ACPI
OnNow 852
recover considerations 853
scheduled backup considerations 853
Sun StorEdge EBS support 852
actions, notifications 546
Active Directory
authoritative restore 840–843
backup 837
domain controller 837
explained 837
recover 838
prerequisites 406
SYSTEM STATE save set 802, 837
ad hoc backups 79
adding
annotation 474
enabler code 501
folder 521
host 519
license 501
multiple hosts 523
user 492
addresses
server, changing 789
Administration window 34
administration window
opening 36
administrator
adding 531
default password 492
privileges 527, 531
resetting administrator password 495
user 492
Windows groups and Sun StorEdge EBS privileges 533
administrators
group 40
Advanced Configuration and Power Interface. *See* ACPI
Advanced File Type Device
concurrent access 272
concurrent operations 272
notification 268
retention policy 268
simultaneous recovery of save sets 272
aes encrypted data
implementing 75
recovering 386
affinity
clone storage node 258
AIX
Legato Cluster, using with 592
alert priority 472, 480, 485
alerts
lists priority, category, time, and message 475
Alias attribute 779
aliases, host sharing restriction 763
annotation
adding 474
attribute 471
icon 473
viewing 474
Annotation Details 455
ansrd program 47
application data
recovering on Sun StorEdge EBS client UNIX 873
application information field
Auspex 683
EMC 678–679
NetApp 673
application media pools, Removable Storage 703
application specific module (ASM) 197
archive
changing time of 237

request, starting automatically later 488
stopping request in progress 488
Archive attribute 159
Archive button 40
archiving 225
archive pools
errors 791
archive requests
creating 235
defined 235
deleting 238
disabling 242
editing 237
archive services
enabling 231
archive volume pool 226
cloning 245
display errors 791
compared to backups 225
copying an archive request 236
disable scheduled archive 489
enabling 228–229
grooming files 225
indexed 228
licensing 225
managing 242
manual 232
manually archiving data 234
naming archive requests, errors 791
NDMP 626
nonindexed 228
nsrarchive program considerations 791
permissions 228–230
pools
configuration 212
creating 216
types of pools 230
remote requests, failure 790
requesting an archive 487
requesting an archive's status 476
requirements 226
retrieving from client machine 238
save sets 225
multiple 790
retrieving 238–239
scheduled 235
scheduling 242
starting 242
stopping 242
Sun StorEdge EBS User program
toolbar function 41
time change 237
tracking entries 226
troubleshooting 790
viewing details 242
volumes, cloning 258
ASM (application specific module) 197
ASR save set 732, 807
ASR. *See* Automated System Recovery
attribute
Administrator 532
annotation 471
auth code 503
category 471
enabler code 501
Grooming
limitation 235
message 471
note 471
priority of managed event 471
server name 471
time 471
attributes
See also specific attribute name
audience, for documentation 21
auth code attribute 503
authentication
considerations for 572
nsrauth 571
oldauth 571
authorization code
entering 502
error message 793
authorizing software 502
auto media management
cloning, considerations 245
recyclable volumes 173
auto media verification
media position errors 761
autochangers
AIX considerations 867
attributes
new 772
Auspex support 682

autodetection 772
control ports access 774
destination component 773
EMC Celerra support 677
EMC IP4700 support 679
HP-UX considerations 857
installation, HP-UX considerations 858
inventories 772
maintenance commands 772
NDMP support 635–636
NetApp support 672
sharing
NDMP clients 638
sharing with NDMP support 638
STK-9840 867
troubleshooting 771–773
Automated System Recovery
ASR save set 732, 807
cluster database, special handling 741
COM+ database, special handling 741
disk quota database, special handling 741
documentation 732
FAT16 partitions 743
limitations 743
OEM recovery CD 744
overview 731
recovering a client 739
scheduled backup of ASR save set 734–735, 737
special considerations 743
Sun StorEdge EBS software support 732
usage overview 732
verifying client recovery 742
WMI database, special handling 741
Autorestart attribute 590
Autostart attribute 136

B

back ups, manual 79
backing up Console 78
Backup button 40
backup groups. *See* groups
backup levels
1-9 154
described 154
options 154
overriding 152
planning 155–157
types
consolidated 147, 154, 158, 163
full 154, 157–158, 764
incremental 154
level 158
skip 154
usage 157–158
backup schedules. *See* schedules
Backup Summary Over Time 442
Backup Type Summary 439
Backup Type Summary by Server over Time 444
backups
Backup and Recover Server service 541
balancing resources 134, 148
bootstrap 142
client-initiated 211
clusters
requirements 615
commands 88
example 90
for open files 143
savepnpc program 94
completed 142
consolidated 147
customization scripts 88
cycle, using levels 146
directives 185
failed 86
filesystems 149
force incremental 139
groups. *See* groups
hard links 86
incremental
pool for 210
large client filesystems 583
large filesystems 149
levels 158, 166
log file 86
managing 140
manual 66–67
policies 182
pool for 211
NDMP 626
consideration with Auspex 681

non-scheduled 67
nonscheduled 66–67
nsrndmp_save command, using 648, 653
online indexes 158
operations
stopping 760
operators group 40, 540
password protection 73
performance when cloning 244
permissions, backup operators group 540
pools 203
previewing 140
recoveries 758
RPC errors 767
save sets. *See* save sets
server, forced 768
sessions 529
SharePoint Portal Server 724
Sun StorEdge EBS User program
browse windows 41
time intervals, setting 139
troubleshooting 758
types of backups. *See* backup levels
volumes
clones 250
Windows NT registry 70
bar chart 431
basic reports 420
Managed Event 454
BlueArc
requirements 685
bootstrap
NDMP limitation 625
pools 207
printing 142
failure 768
boot-time file 786
browse policies
about 169–170
client file index growth 553
clones, storage nodes 258
defined 169
save set recoveries 397–398
usage 169
browse windows
described 41
toolbar 40
browser
unresponsive 794

C

capture of managed event 470
case sensitivity 463, 511
category attribute of managed event 471
Category Summary, Managed Event basic report 454
centralized
license management 33, 500
Certificate Server
recovery prerequisites 406
SYSTEM STATE save set 802
certified protocols 786
changing
name of folder 523
password 494
changing password, administrator 492
changing password, user 494
changing servers 42
characters
illegal 522–523
characters, not permitted 493, 522–523
chart formats 427
CHKDSK, running 67–68
cleaning Removable Storage libraries 720
client 567
of Sun StorEdge EBS server 78
save set 584
Sun StorEdge EBS
definition of 567
client alias, changing 763
client configuration wizard 62
client file index
backup level, pool for 210
browse policies 169–170, 560
checking 766
clones and storage nodes 258
cross-checking 557
defined 552
entries
adding 554
removing 554, 560, 562
growth 553
location, designating 558

managing size 560
operations
backups 158
checking 554
moving 559
recovery 368, 395
policies 560
pools 207
restoration 397
save sets
cycles, removing 561
entries 173
removing 560
size 553
management 560
notification 766
Client resource
NDMP 642
Client Retries attribute 137
Client Service port 825
Client Status over Time, Sun StorEdge EBS Backup Status Report 452
Client Summary by Server and Group, Sun StorEdge EBS Backup Status 449
Client Summary, Sun StorEdge EBS Backup Statistics Report 439
client-initiated backups
pool for 211
clients
aliases, problems 763
backup commands 88
client ID
creating new client 399
client / server communication errors 781
cloning
examples 251–253
configuration 568
defined 567
DHCP 540
DNS name resolution 540
editing 568
groups 134
large filesystems 149
manual backups 66
multiple 583
NDMP considerations 642
operations
archive retrieve 238
backing up renamed clients 765
backups 583
creating 568
editing 568
indexes, moving 559
installation 568
manual backups 66
recovery, failure 765
policies, multiple 178
priority 586
save sets 583–584
Solaris binary location 855
Sun StorEdge EBS User program 39
unsynchronized client clock 759
clnt_svc_port 825
clone pools, configuring 211
Clone Storage Nodes attribute 258
Clone Volumes dialog box 246
cloning
archives 258
display errors 791
automatic 250
defined 243
destination volume, defined 244
examples 251–254
manual 248
online indexes, storage nodes 258
recovery
save sets 257
volumes 257
save sets 247–248
automated 250
manually 251
methods 247
performing 248
source volume, defined 244
storage nodes
affinity 258
online indexes 258
volume duplication, comparison 245
volumes 246, 250
automatic 250
creating 246
details, viewing 247
process 246
requirements 244

Cluster Server
recovery prerequisites 406
SYSTEM STATE save set 803
clusters
backup requirements 614–615
backups 589, 591
clients 590
configuration options 600
data corruption 601
devices 600
jbconfig program considerations 602
licensing 614–615
MC / ServiceGuard 591
recovery types 594
storage node 599
tape libraries 600
virtual client backing up to a local storage node 592
code set support 777
COM+ Database
recovery prerequisites 406
SYSTEM STATE save set 802
command
export 512
gst 50
nsrexecd 828
nsrmmd 827
savepsm 78
command line
reporting 420, 466, 468
reporting program 467
Common Device Interface
SCSI command 307
Completion Data Retention 415
completion data retention 415
completion message retention 415
compression, data 77
computer damage, recovery from 499
configuration
IP issues 787
NDMP Device resource 640
Configure 422
configuring
Host reports 457, 460
reports 422
Sun StorEdge EBS Backup Statistics reports 436
Sun StorEdge EBS server and firewall 827
User Reports 458
connection
problem 794
refused 794
to Sun StorEdge EBS server through firewall 826
connection ports 821
connections, server
problem resolution 786
Console
enviroment variables 504
Console client
starting, after the first time 36
console configuration wizard 514
Console server
stopping and starting 49
Console software
as Sun StorEdge EBS client 78
HTTP service port 35
improving performance 504
License Manager 503
load on CPU 506
logging on 35
managing window 31
URL 35
Console window
opening 35
console window 32
consolidated backups 158, 163
consolidating save sets 163
nsrssc command 165
contacting server, problem with 794
Content Index Server
backing up 83
backup 83
defined 83
recover 83, 412
SYSTEM DB save set 83, 412, 805
control zone 515
copying host 520
corrupted database 794
CPU load 506
creating
folder 521
host 519
label templates 223

pools 213
staging policies 262
user 492
criteria for organizing hosts 516
critical priority 472, 480, 485
cross-checking online indexes 557
cross-platform
name resolution 540
CSV 465
cutting and pasting host 521

D

daemon.log 566
Daily Backup Status 451
Daily Save Set Summary 449
Daily Summary 440, 448
Daily / Monthly Save Set Statement 443
damaged computer 499
Dantz servers
monitoring 469
data
compression 77
encryption 74
encryption with aes ASM 75
encryption with pw2 ASM 76
life cycle
cloned data 245
managing 176
password protecting 74
recovery, name conflicts 375
relocating, errors 767
sorting
into pools 205, 208–209
to storage devices 212
verify 72
data compression 774
Data Retention dialog box 417
database
backing up 78
corrupted 794
corruption 794
delete failed 795
failed to store 795
fetch operation 794
database, fetch operation failure 794
db-output 798
dbsrvr9 799
dbstop outpu 798
DDS (dynamic drive sharing)
clustering 605
robotics control 611
configuring 109
debug level 505, 512
debug messages
logging stopped 797
default
location of
install log 798
password for administrator 492
default backup schedules 144
default pool 204, 209, 211
defining user 492
definition of
Sun StorEdge EBS client 567
delete operation failed 795
deleting
folder 522
host 520
label templates 224
license 502
managed event note 473
multiple hosts 520, 523
note 473
problem 795
staging policies 265
user 494
determining client service port 825
determining HTTP service port 825
device
lists of 476
related messages 476
Device dialog box 529
devices
device drivers
maintenance commands 772
disk label errors 768
filesystem
staging 261
labeling errors 768
nonrewinding 775
pools 212
troubleshooting 771
volumes, clones 250

DFS (Distributed File System)
backup 831
defined 831
junctions 831–832
not in SYSTEM save sets 806
recover 831, 835
DHCP (Dynamic Host Configuration Protocol)
database 84, 411
static IP address for Sun StorEdge EBS server 789
DHCP (dynamic host configuration protocol) 783
clients 540
dialog box
Data Retention 417
DinoStor TapeServer
configuration 686
directed recoveries 394
directed recovery
access 390
advantages 388
defined 387
NDMP 661
performing 380, 395
recover program 393
UNIX clients 395
use of 388
directives
ASM (application specific modules) 197
copying 187
creating 186
defined 185
deleting 187
editing 187
naming restrictions 762
preconfigured 192
disabling
managed event capture 470
disabling user datagram protocol 508
disallowed characters 493, 522–523
disappearing managed event 474, 797
disaster recovery 386
disaster recovery, SharePoint Portal Server 729
disk quota database
recovery prerequisites 407
SYSTEM DB save set 805
disk space
insufficient 794
disk space, gstd.log 505, 511–512, 798
display problem 795
displaying
annotation 474
reports 434
DNS
hostname alias, troubleshooting 763
DNS (Domain Name System) 783
host name determination 540
document view 434
domain controller
Active Directory, configured by 837
defined 837
encryption keys
not supported 844
recovering system state on 838
drag and drop, column 37
drill-down reports, Managed Event 456
DSA 619
DSA and NDMP Tape Server distinctions 621
duplicating host 520
dynamic addressing 540
dynamic host configuration protocol (DHCP) 783

E

ECB (Event Control Block) 761
editing
clients 568
folder 523
label templates 223
pools 215
staging policies 264
user 494
EMC, application information values 678
emergency priority 472, 480, 486
enabler code 501
enabler code, entering 501
enabler code, problem 795
enabler codes
entering 501
removing 501
enabling

connection to Sun StorEdge EBS through firewall 827
debug information 505, 512
JavaScript 800
software 501
encrypting data 74
aes 75
pw2 76
Encrypting File System
backup 844
directed recover 844
encryption keys 844
explained 844
not in SYSTEM save sets 806
recover 844
entering
authorization code 502
enabler code 501
license 501
enterprise
button 32
enterprise hierarchy
adding folder 521
adding host 519
copying folder 523
copying host 520
deleting folder 522
deleting host 520
folder 516
host 516
managing host 518
managing multiple hosts 523
moving folder 523
moving host 521
renaming folder 523
viewing 518
Enterprise Summary report 457
environment variable 511
GST_DEBUG 799
GST_MAXLOGSIZE 511, 799
GST_MAXLOGVERS 512, 799
setting 511–512
environment variables
format 565
NSR_DEV_BLOCK_SIZE_*MEDIA_TYPE* 299, 305
NSR_DEV_LOAD_POLL_INTERVAL_*MEDIA_TYPE* 306
NSR_DEV_LOAD_TIME_*MEDIA_TYPE* 306
NSR_DEV_LOAD_TRY_TIMEOUT_*MEDIA_TYPE* 306
NSR_DEV_TAPE_FILE SIZE_*MEDIA_TYPE* 305
NSR_DEV-DEFAULT_CAPACITY_*MEDIA_TYPE* 307
error
database store 795
relocating data 767
RPC (remote procedure call) 767
error messages 793–794
copy violation 768
destination component 773
disk label 768
environmental variables 792
hardware address not found 865
hostname permission denied 785
illegal record size 762
media verification 761
no client resource 763
no space left in tape 868
nwrecover program 764
print server 548
RPC errors 785
save sets 764
server, unavailable 785
xview 769
Event Control Block (ECB) 761
Event Details 455
event logs
backup 844
explained 844
not in SYSTEM save sets 806
recover 844
Event Time 453
Event Viewer 550
Events by Category 456
Events by Priority 456
Events by Server 456
example
organizing hosts 517
sorting managed event 37
export command 512

export formats 465
exporting
non-ASCII characters 466
reports 466

F

failed
operation 795
failed operation 794
failover, save group requirements 590
FAT16 partitions, ASR limitation 743
fetch operation 794
file conversion, sparse to fully-allocated 770
file handling
indicators 41
file index missing, message 765
File Manager 763
File Replication Service
backup 845
SYSTEM STATE save set 803
filemarks 774
filename, support for short names 85
files
compressing 77
encrypting 74–76
HOSTS 540
log 798
naming conflicts on recovery 375
open files, backing up 143
relocating during recovery 376
servers 585, 779
verify 72
filesystem devices 261
filesystems
backups, large 583
remote, saving 785
firewall 811
firewall support 881
configuration 816–817
packet filtering 811
firmware, verification 776
folder 516
adding 521
deleting 522
editing 523
Force Incremental attribute 150
force incremental attribute 139
formats
export 465
free media pools, Removable Storage 702
full backups 157

G

grooming 235
limitation 235
Grooming attribute 234
grooming files 225
group
backup, stopping 478
restarting backup 478
start immediately 478
viewing control details 479, 484
viewing status 475
Group Details report 449
Group resource
Group Status by Server, Sun StorEdge EBS Backup Status Report 452
Group Status over Time, Sun StorEdge EBS Backup Status Report 451
Group Summary
Sun StorEdge EBS Backup Statistics Report 438
Group Summary by Server
Sun StorEdge EBS Backup Statistics Report 440
Group Summary by Server, Sun StorEdge EBS Backup Statistics Report 440
Group Summary by Server, Sun StorEdge EBS Backup Status Report 449
Group Summary, Sun StorEdge EBS Backup Statistics Report 438
Group Summary, Sun StorEdge EBS Backup Status Report 448
groups 134
attributes 136
backup operators 540
backups
management 140
previewing 140
bootstrap 142
client policies, multiple 178
completed 142

containing bootstrap 762
defined 134
filesystems, large 149
force incremental, setting 139
listing of per server 475
naming restrictions 762
operations
copying 138
creating 56
deleting 138
editing 138
previewing 140
time intervals, setting 139
overview 133
Sun StorEdge EBS User program privileges 39
types 136
gst command 50
GST server process 799
GST_DEBUG 799
GST_MAXLOGSIZE 511, 799
GST_MAXLOGVERS 512, 799
gstd 799
environment variable 512
log 505, 512, 798
process 50
size of log 505, 512, 798
gstd service
defined 47
gstd.log file 798
gstmodconf command 523, 525

H

hard links, backup and recovery 86
hardware
upgrade 498
host
adding 519
copying 520
deleting 520
managing 518
moving 521
organization of 516
transfer affidavit 500
Host List 457
host name determination
DHCP clients 540
TCP/IP 540
Host Reports 418, 457
types and their configuration 457, 460
hostname alias 763
hostname file 524
HOSTS file 540
hosts table, setting up 782
HP MC/ServiceGuard
clusters 591
HP-UX
creating device files 862
installing autochangers 858, 862
Legato Cluster, using with 592
pass-through drivers 858–861
system actions 862
system shutdown 862
testing autochangers 865
HTML 465
HTTP service port 35
http_svc_port 825
hung browser 794

I

icon
annotation 473
priority 472
If 759
IIS (Internet Information Server)
defined 845
recovery prerequisites 406
SYSTEM STATE save set 802
import media pools, Removable Storage 702
Importing save set information to a different Sun StorEdge EBS server 399
importing volumes, Removable Storage 714
Inactivity Timeout attribute 137
incremental backup
NDMP limitation 628
incremental backups
pool for 210
Index Save Sets dialog box 555
indexed archives 228
individual user authentication 509
info priority 472, 480, 486
information, sorting table 37

inquire program 635, 857
install log 798
Interactive view 426
interface
  overview 31
Internet Information Server *See* IIS
Interval attribute 136
Invalid Object id, error message 795
ioscan program 858
IP
  errors
    hosts table 782
    troubleshooting 781
  name search order, setting 783
IP configuration, issues 787

J

Java Runtime Environment 468
JavaScript 800
jbconfig program
  clusters 602
  hanging 772
jbexercise program
  NDMP, not supported 626
JRE 35, 425–426, 467–468

L

label templates
  attributes 219–221
  components 222
  creating 223
  deleting 224
  editing 223
  naming restrictions 762
  naming strategies 222
  number sequences 221
  preconfigured 218
labeling
  tips 222
labels
  silos, volumes 217
launch button 35
LEGATO
  License Manager 503
Legato Cluster 592
level of debug 505, 512
Level Summary 439
Level Summary by Server over Time 443
license
  adding 501
  deleting 502
License allocation failed, error 796
License Manager 503
licenses 500
licensing
  archiving 225
  clusters 614–615
  copy violation 768
  entering authorization code 502
  License Manager 503
  modules 615
  Removable Storage 699
  Sun StorEdge EBS support for SPS 722
Linux, Legato Cluster, using with 592
list command 677
List Report, User 458
LMHOSTS, use with Sun StorEdge EBS 540
local backups with NDMP
  introduction 619
locale support 777
localized environments 423, 425
location of
  gst command 50
  install log 798
log file 511, 798
log files
  backup / recovery attempts 86
  cloning information 243
  management 564
  viewing 485
logging on
  to Console software 35
logging, events 550
login
  name 493
  password 493
lsdev program 863–864

M

Mac OS X 869

managed application 516
managed event
  define 470
  deleting note 473
  disabling capture of 470
  disappearing 474, 797
  polling 506
  priority 472
  response time
    modifying 506
  sorting example 37
Managed Event basic reports 454
Managed Event configuration
  parameters 453
Managed Event drill-down reports 456
Managed Event Reports 418
Managed Events button 32, 34
managed node 516
  adding 519
  copying 520
  deleting 520
  moving 521
managing
  host 518
  license 503
manual backups. *See* backups, manual 66
Manual Restart option 590
MC / ServiceGuard
  clusters 591
media allocation
  Removable Storage 713
media database
  backups 158
  cloned data 245
  clones and storage nodes 258
  compression 561, 564
  cross-checking 557
  entries, removing 562
  managing size 552, 561
  restoration 398
  retention policies 169–170
  save sets, entries 173
media deallocation
  Removable Storage 713
media label library (MLL) 712
media pools
  Removable Storage 702
media pools, Removable Storage
  application 703
  overview 702
  Sun StorEdge EBS 703
  viewing 703
media pools. *See* pools
media position errors 761
memory
  NDMP requirements 628
message attribute of managed event 471
message logs
  failed backup / recovery attempts 86
  management 564
  reducing size 564, 566
message retention 415
messages file
  Event Viewer 550
Microsoft Automated System Recovery. *See* Automated System Recovery
Microsoft Removable Storage. *See* Removable Storage
Microsoft SharePoint Portal Server. *See* SharePoint Portal Server
Microsoft Windows
  backup operators group 540
  Event Viewer 550
  Windows 2000
    databases
      not in SYSTEM save sets 805
      SYSTEM DB save set 804
  Windows Management Instrumentation
    SYSTEM DB save set 805
Mirapoint
  requirements 684
mminfo program
  reports 184
modifying
  log size 511
  response time 506
  user 494
modules
  NDMP 645
  NetWorker database modules 615
monitoring
  Dantz servers 469
  Retrospect servers 469
Monthly Client Statement 442

Monthly Group Statement 442
Monthly Summary 440, 448
  Sun StorEdge EBS Backup Statistics Report 440
moving host 521
multiple hosts
  adding or deleting 523
multiple hosts, deleting 520
multiplexing
  performance issues 27

N

name
  conflicts on recover 375
  folder, editing of 523
  resolution cross-platform 540
  servers, issues 788
name resolution 783
name servers, disabling 783
naming restrictions 762
NDMP (Network Data Management Protocol) 640
  application information variables 644
  archives 626
  attributes 642
  Auspex
    autochanger handle 682
  autochangers
    configuration 636
    list command 677
    sysconfig program 672
  backups
    command line options 648, 653
    parallel 626
  bootstrap limitation 625
  clients 642
    definitions 626
  configuration
    options 619
    Sun StorEdge EBSserver 633
  device pathname 635
  directed recoveries 661
  DSA 878
  DSP 878
  EMC Celerra
    autochanger handle 677
  EMC IP4700
    autochanger handle 679
  Force Incremental setting 646
  group Parallelism setting 646
  incremental backup limitation 628
  inquire program 635
  installation 618
  memory requirements 628
  NDMP data server 884
  NDMP host 884
  NDMP save set 884
  NDMP server 885
  NDMP service 885
  NDMP storage node 885
  NetApp
    autochanger handle 672
  non-NDMP device 886
  nsr resource attributes 628
  operations
    browsing 626
    recoveries 655
    recoveries, directory level 658
    recoveries, save set 656
  prerequisites 647
  protocol limitations 628
  save set operations
    consolidation 626
    recovery 628
  scanner program 627
  sharing autochangers 638
  silo support 640
  storage node support for NDMP client backup 651
  storage node support for NDMP clients 630
  Sun StorEdge EBS
    changes with NDMP 626
    changing location of temporary space directory 629
    configuring 633
    index database 628
    index processing and memory requirements 628
    storage node support for NDMP clients 630
    temp storage space requirements 629
NDMP connection support 618
Nested Mountpoints 408

network
DNS connection issues 786
IP configuration issues 787
TCP/IP certified 786
node
adding 519
copying 520
deleting 520
managed 516
moving 521
nonindexed
archives 228
note
attribute 471
deleting from managed event 473
notifications 551
defined 541
deleting 552
operations
customizing 545
preconfigured 541–544
printing 547
priorities 549
programs 546
SNMP
configuring 692
creating 694
modifying 693
nsrtrap 692
NSR_AVOID_ARCHIVE 159
nsr_getdate program 182–183
nsr_shutdown program 566, 761
nsradmin program
editing the nsrla.res database 570
starting 47
nsrarchive program 227
nsrauth authentication 571
nsrck program 395
nsrd daemon 566
nsrd service 566
defined 47
nsrexec daemon
multiple processes 758
troubleshooting 758
nsrexecd 50, 828
nsrexecd daemon 566
functionality 785
nsrexecd service 566
defined 48
nsrindexd service
defined 48
nsrjb program
troubleshooting, HP-UX 866
nsrla.res database 570
nsrlpr program 547
nsrmm program 183, 397, 560, 761
nsrmmd 827
nsrmmd daemon 306
NDMP
unsupported options 628
nsrmmd service
defined 48
NDMP
unsupported options 628
nsrmmdbd service
defined 48
nsrndmp_recover program 655–656, 658
nsrndmp_save program 651
backing up data 648, 653
nsrports program 821
nsrssc program 165
nsrtrap 692
command line options 631, 692
verbose mode 693
numeric order 37
nwarchive program 45, 234
nwbackup program 42, 67
nwrecover program 44
browse policy 169
nwretrieve program 46, 239, 241, 790

O

object, reference 795
OEM recovery CD, ASR limitation 744
oldauth authentication 571
online indexes
cross-checking 557
entries
checking 554, 766
removing 554, 562
information
refreshing 557
viewing 554

management
manual 552
size 560
moving 559
recovery
location 762
restoration 397
save sets, viewing 555
size considerations 553
volumes
removing 562
OnNow
Advanced Configuration and Power Interface (ACPI) 852
Sun StorEdge EBS support 852
recover considerations 853
scheduled backup considerations 853
open files, backing up with VSS 144
operation failed 794–795
optimizing Console 504
organizational criteria 516
organizational structure, labeling for 222
organizing
Sun StorEdge EBS Servers
example 517
Sun StorEdge EBS servers 516
out of memory 794

P

packet filtering 811
page cannot be displayed 795
parallelism
defined 528
maximum 528
performance 27
setting 529
password
changing 494
default for administrator 492
password protection 74
pathname restrictions 763
PDF report format 465
performance
cloning 244
features 27
parallelism 528
target sessions 529
Performance Counters
SYSTEM STATE save set 802
performance of Console software 504
permissions
Archive feature 228–230
backup operators group 540
permissions, setting user 493
physical host 884
pie chart 432
pie report format 431
ping 829
plot chart 432
policies 179
backups, manual 182
browse
about 169–170
data life cycle 176
defined 169
modifying 183
usage 169
clients 178
data life cycle 176
multiple 178
naming restrictions 762
overriding 182
planning 148
retention 173
about 169–170
data life cycle 176
defined 169
modifying 183
usage 169
volume relabeling 169
setting expiration 417
policy
deleting 179
polling 507
pool
consolidated backup 207
copying resource 215
Pool Type attribute, with archive pools 231
pools
archive 212
errors 791
archives, creating 216
archiving 216

bootstrap 207
client file index 207
clones 211
configuration 210
archive 212
clone 211
creating 213
criteria 205, 208
data sorting 205, 209
default 204, 209, 211
clone pool 211
defined 203
devices 212
editing 215
expression matching 207
incremental backups 210
manual backups 211
precedence 208
Removable Storage media 702
restrictions 205
save set consolidation 166
sorting data 212
volume labels 217
port
client service 825
HTTP service 35, 825
service 827
port range
troubleshooting 789
portmappers
nsrexecd daemon 785
registration 823
verifying 784
ports
configuration 821
connection
configuration 821
default settings, changing 820
restricting range 817
service, configuration 821
troubleshooting 823
POSIX hard links, problems recovering 86
postprocessing command 143
PostScript report format 465
Power Monitor service 852
precedence for pools 208
preconfigured notifications 551
preprocessing command 143
preventing unauthorized access 492, 829
printing
reports 468
priorities, notifications 549
priority
of managed event 472
symbol 472, 480, 486
Priority attribute 471
Priority Summary 454
privileges
administrator 527
problem
contacting server 794
process, stopping and restarting 50
processor load 506
Procom NetFORCE
requirements 684
Program not registered 796
program not registered 796
programming thread 507
programs. *See* specific program name 42
properties
user 494
Properties dialog box
Sun StorEdge EBS User program 42
protocols, certified 786
provider 746
ps command 800
pw2 encrypted data
implementing 76
recovering 387

R

rearranging enterprise hierarchy 521
rearranging information in table 37
Recover button 40
recover program
clusters, using with 595
media database 173
options 393
retention policy 173
recoveries 369
archives 238
Backup and Recover Server services 541
backup operators group 540

client file index 395
clients, renamed 765
clone volumes 257
cluster 594
conflict resolution 375
conflicts ??–374
directed 394
access 390
advantages 388
defined 387
recover program 393
usage 388
disaster-related 386
encrypted data 386
failed 86
failure 764
files, finding for recovery 378
hard links created by POSIX 86
index-based
advantages 368
interrupted backups, from 764
log file 86
methods, comparing 367
NDMP 655, 658
save set recoveries 656
planning 148
relocation 376
save set recovery 379
save sets
client file index 397
media database 398
Sun StorEdge EBS User program
browse windows 41
type, determination 367
volumes
required 378
recovery
aborting 766
ASR 739
Removable Storage database 698
SharePoint Portal Server 726
recycling volumes, Removable Storage 718
recycling. *See* volumes
reference object 795
refused connection 794
registering program 796
registry
backups 845
backward compatibility 846
explained 845
recover
prerequisites 407
REGISTRY save set 846
SYSTEM STATE save set 802, 845
relocating data 767
remote access 369
recoveries 780
Remote Access list 780
remote archives, failure 790
remote backups with NDMP TapeServer software
performing 689
remote filesystems, saving 785
Remote NDMP backups with DinoStor TapeServer
performing 687
remote NDMP backups with SnapImage software
performing 689
Removable Storage
adding media 715–716
cleaning devices 720
configuring a Sun StorEdge EBS server for 706
creating Sun StorEdge EBS resources for 708
database 698
database recovery 698
device cleaning 720
disabling a library 706
enabling a library 706
hardware support 700
importing volumes 714
inquire command 708–709
jbconfig command 708–709
jukebox resource on a Sun StorEdge EBS server 709
jukebox resource on a Sun StorEdge EBS storage node 710
libraries 698
licensing 699
limitations 701
logical slot numbers 711
media allocation 713
media deallocation 713
media label library 712
media pools 702

media pools, viewing 703
monitoring the work queue 707
OS support 699
overview 697
physical slot locations 711
recycling volumes 718
removing media 718
reset 719
Sun StorEdge EBS software support for 699
systems 698
uninstalling Sun StorEdge EBS software 720
usage scenarios 704
Removable Storage Manager
recovery prerequisites 407
SYSTEM DB save set 805
removing
folder 522
host 520
license 502
user 494
renaming folder 523
report
background processing 435
chart types 430
document mode 428
export formats 465
Group Details 449
interactive mode 426
restricted views 435
User List 458
reports
basic and drill-down 419
clones 250
command line reporting program 467
customized 421
daemon.log 566
date and time formats 423
Managed Event Basic 454
Managed Event Drill-Down 456
message logs
management 564
reducing size 564
save set policies 184
saved 462
viewing 434
Reports button 33, 35
requirements
Sun StorEdge EBS User groups 39
reset, Removable Storage 719
resetting administrator password 495
resolved events 474
resource
archive
changing archive time 237
request
copying 236
status 476
client
defining save set with 584
label template
copying 224
notification
copying 551
deleting 552
policy
deleting 179
pool
copying 215
directing data from consolidated backup 207
staging
copying a policy 265
user group
copying 539
creating 538
customizing privileges 535
deleting 539
editing 539
preconfigured 533
privileges 535
resources
configuring for NDMP 633
response time
managed event 506
restrictions on pathnames 763
retention policies 173
about 169–170, 173
clones, storage nodes 258
defined 169
usage 169
volumes, relabeling 169
retention, completion data 415
retention, completion message 415
retention, save set 416

retrieval
troubleshooting 790
Retrieve Status window 240
retrieving archives 238
retrieving save sets 238–239
Retrospect servers
monitoring 469
RPC (remote procedure call)
errors 767
rpcbind program 785
rpcinfo program 784, 822
RSM. *See* Removable Storage

S

save program 88, 182–183, 595
save set
All save set 65
defining with a client 584
save set consolidation 163
Save Set Data Retention
and Sun StorEdge EBS Backup Statistics Reports 437
Save Set Data Retention Policy
and configuration of Sun StorEdge EBS Backup Statistics Reports 437
Save Set Details 441, 445, 450
Save Set Details by Client 441, 450
Sun StorEdge EBS Backup Statistics Report 441
Save Set Details by Server over Time 443
Save Set Name 446
Save Set Output 450
save set retention 416
Save Set Summary 439
Sun StorEdge EBS Backup Statistics Report 439
save sets
archives, retrieving 238–239
backup commands 88
backups 584
client combination 584
client file index, entries 173
client priority 586
clones 247
cloning
automatically 252
described 247
manually 251
status 248
consolidation 163
files, large 163
limitations 164
NDMP 626
nsrssc command 165
pools 166
usage 163–164
defining 584
indexes, viewing 555
information, viewing 555
load balancing 583
media database, entries 173
multiplexing 530
policies
modification 183
reports 184
predefined 64
recoveries 379, 381–382
client file index 397
conflicts 384
media database 398
NDMP 628
online indexes 397
relocating 383
requirements 380
usage 379
volumes 382–383
recovery
compared to index-based recovery 367
REGISTRY, backward compatibility 846
staging. *See* staging
status
clone 248
retention policy 173
suspected 257
SYSTEM
manual backups 69
point-in-time recovery
command prompt 404
recover
command prompt 403
SYSTEM DB
backup levels 161
basic components 804

databases not included 805
optional components 805
recovery prerequisites 407
SYSTEM FILES
backup levels 160
basic components 803
SYSTEM STATE
backup levels 160
basic components 801
optional components 802
recovery prerequisites 406
VSS ASR DISK
backup levels 162
components 809
VSS OTHER
backup levels 162
components 809
VSS SYSTEM
manual backups 69
point-in-time recovery 404
VSS SYSTEM BOOT
backup levels 161
components 807
VSS SYSTEM FILESET
backup levels 162
components 808
VSS SYSTEM SERVICES
backup levels 162
components 809
VSS USER DATA
backup levels 162
components 809
Savegroup Completion Report 142
savegrp program
backup limitation 768
save set consolidation from command line 165
savepnpc program
message logging by 98
using with customized backup program 94
savepsm 78
savestream 203
scalability, Sun StorEdge EBS 28
SCANDISK, running 67–68
scanner program
NDMP 627
record size 762
recovering clone volumes 257
recovering save sets from volumes 176
recreating online indexes 382
recreating volume entries 382
retention policy 173
volumes, read-only 761
Schedule resource 133
schedules 144
attributes 150
backup cycle 146
balancing 148
copying 152
default 144, 147
deleting 151
editing 151, 165
large filesystems 149
load balancing 584
naming restrictions 762
overriding 152
overview 133
planning 148
preconfigured 144
staggered 147
usage 144
Schedules window 150
scheduling backups 78
SCSI address selection for HP-UX 858
securing Console 492
security 527
server
connecting through firewall 826
problem contacting 794
setting up 527
stopping and starting Console 49
web address 35
server name
attribute of managed event 471
Server Summary 454
Sun StorEdge EBS Backup Statistics Report 438
servers
address, changing 789
administrators
adding 531
backup
operators Group 540
client/server communication errors 781

clients
tasking rights 585
communication issues 540
DCHP 540
DNS name resolution 540
dynamic addressing 540
errors, binding to 785
file 585, 779
index
backup, failure 768
management 553
notifications
priorities 549
parallelism 528
ping command, testing 784
target sessions 529
testing 784
service
port for firewall support 827
port, client 825
port, HTTP 35, 825
service ports
configuration 821
services
Backup and Recover Server 541
described 47
Power Monitor service 852
stopping and starting 49
session management 27
sessions
lists of backup, recover, or browse sessions 475
set
expiration policies for data retention 417
set expiration policies for data retention 417
setting
data retention policies 417
environment variable 511–512
expiration policies 417
service port 827
target sessions 530
severity of managed event 471
shadow copy 746
SHAREPOINT 723
SharePoint Portal Server
backups
command prompt 725
manual 725
overview 724
scheduled 725
disaster recovery 729
overview 721
recovery
command prompt 728
overview 726
to a different server 728
using Sun StorEdge EBS User 727
required services 722
SHAREPOINT save set 723
Sun StorEdge EBS licensing 722
Sun StorEdge EBS software support 721
system requirements 722
SHAREPOINT save set 723
short filename support 85
silos
NDMP support 640
Simple Network Management Protocol. *See* SNMP
size of
gstd log 798
log file 511
sleep state
defined 852
SnapImage
autochanger handle 688
snapshot 745
snapshot policy
creating 181
SNMP (Simple Network Management Protocol) 691
configuring 692, 695
defined 691
notifications
configuring 692
creating 694
modifying 693
nsrtrap 692
traps 691
Software Administration button 33
Solaris
Legato Cluster, using with 592
troubleshooting 855
sorting managed event, example 37
sorting table 37
Source Client dialog box 380–384

sparse files 846
not in SYSTEM save sets 806
sparse files, converting 770
Special Handling dialog box 78
spoofing user datagram protocol 829
SPS. *See* SharePoint Portal Server
stacked bar report format 431
stacking bar chart 432
staging
defined 261
filesystem devices 261
policies
creating 262
deleting 265
editing 264
Start Time attribute 136
starting
server process 50
starting services 49
status
viewing
group status 475
status values, scanned-in 397
stopping Console server 50
stopping services 49
storage nodes
affinity
clones 258
problems 792
clones 258
clusters
virtual storage nodes, support 599
parallelism 528
timeouts 792
troubleshooting 791–792
virtual server 599
volumes, clones 250
Store Index Entries attribute, with archive pools 231
storing database failed 795
Sun 50
Sun StorEdge EBS
Backup and Recover Server 541
connect through firewall 827
Console software as client 78
firewall support 811
interfaces 42
nwarchive program 45
nwbackup program 42
nwrecover program 44
nwretrieve program 46
LMHOSTS 540
name resolution 540
NDMP
configuring 633
effects on 626
optimizing performance 27
Power Monitor service 852
remote access 369
scalability 28
server
connecting through firewall 826
organization of 516
target sessions 529
User program. *See* Sun StorEdge EBS User program
WINS 540
Sun StorEdge EBS Backup Statistics basic reports 438
Sun StorEdge EBS Backup Statistics drill-down reports 441
Sun StorEdge EBS Backup Statistics Reports 418
Sun StorEdge EBS Backup Statistics reports
types and their configuration 436
Sun StorEdge EBS Backup Status drill-down report 450
Sun StorEdge EBS Backup Status Reports 418
Sun StorEdge EBS Backup Status reports 444
Sun StorEdge EBS Client 50
Sun StorEdge EBS client
definition of 567
Sun StorEdge EBS DiskBackup
advanced file type concurrent operations 272
file type and advanced file type differences 281–284
notification 268
retention policy 268
Sun StorEdge EBS media pools, Removable Storage 703
Sun StorEdge EBS server
organizing 516
Sun StorEdge EBS User program

backing up registry, Windows NT 70
browse windows 41
changing servers 42
compression 77–78
connecting to a server 42
encryption 74, 78
manual backups 67
overview 39
password protection 74, 78
privileges 39
server connection 42
starting 39
toolbar 40
support
for firewall 811
suspected save sets 257
symptom of problem 793
sysconfig command 672
System File Protection
backup 803, 808
defined 803
explained 803, 808
recover 803, 808
SYSTEM FILES save set 803
system media pools, Removable Storage 702
system standby 852
SYSVOL
SYSTEM STATE save set 803

T

tables
display or hide columns in 38
multicolumn sorting 37
rearranging columns 37
sorting 37
target sessions 529
setting 530
TCP
See transmission control protocol.
TCP/IP
certification 786
changing Sun StorEdge EBS server address 789
DHCP clients 540
host name determination 540
troubleshooting hostname alias problems 763
technical support, troubleshooting information 755
temporary enabler code
expired 793, 796
temporary notes 471
Terminal Services Licensing
backup 847
defined 847
recover 847
prerequisites 407
SYSTEM DB save set 805
thread, programming 507
three-party NDMP backups
introduction 623
performing 649, 654
time attribute of managed event 471
time range 422
toolbars
User program 40
tracking
cloned data 245
online index information 553
transmission control protocol 825, 828
traps
categories 695
SNMP 691
troubleshooting 505, 512
aborted recover 766
AIX
STK-9840 867
archive pools 791
archive requests
naming 791
archives 790
display errors 791
multiple save sets 790
nsrarchive program 791
remote request failure 790
auto media verification 761
autochangers 771
AIX considerations 867
attributes 772
autodetected scsi errors 772
control port access 774
destination component 773

HP-UX considerations 862, 865–866
inventorying 772
maintenance 772
backups 758
backups levels 764
backups, stopping 760
bootstrap printing, failure 768
client file index
size growth 766
clients
alias 763
Solaris, location 855
unsynchronized errors 759
clones, display errors 791
connectivity issues 788
daemons 758
devices 771
maintenance 772
nonrewinding 775
disk label errors 768
display errors 791
DNS connection problems 786
DNS hostname alias 763
ECB counter 761
file conversion 770
firmware 776
hosts table 782–783
HP-UX
error messages 865
unloading drives 866
unsupported media 866
illegal characters 762
IP configuration issues 787
IP errors 781–782
IRIX 868
licensing, copy violation 768
name resolution 783
name server issues 788
name servers, disabling 783
nsrexec processes 758
nsrexecd program 785
online indexes 762
packet receive buffer 761
pathname restrictions 763
port range configuration 789
portmappers, verifying 784
ports 785, 823
recoveries 758
interrupted backups 764
new installation 764
online indexes 762
remote access 780
recovering POSIX hard links 86
remote filesystems, saving 785
renamed client backups 765
retrievals 790
routers 776
rpcinfo command 784
scanner program 761–762
server errors, binding to 785
server index 768
Solaris 855
storage nodes 791–792
technical support information 755
xview errors 769

U

uasm program 762
unauthorized access 492
universal resource locator 35
unresponsive browser 794
unsynchronized clock errors 759
URL
Console software 35
user
access 829
adding 492
administrator 492
authentication 508–509
deleted 464
deleting 494
editing 494
non-administrator 498
properties 494
views, granting permissions 493
views, restricting 497
user authentication 508
user data
recovering on Sun StorEdge EBS client
Mac OS X 873
user datagram protocol
and Console 828
disabling 508

ping 829
spoofing 829
user groups 533
creating 538
privileges 535
user interface
overview 31
setting preferences 38
User List Report 458
user privileges 39
User program. *See* Sun StorEdge EBS User program.
User Reports 419
types and their configuration 458
User Reports, configure 458

V

vanishing managed event 797
variable
case sensitivity 511
for gstd log 799
GST_DEBUG 799
GST_MAXLOGSIZE 511
GST_MAXLOGVERS 512, 799
setting 511–512
verification
of files 72
Sun StorEdge EBS User program
browse windows 41
Verify button 40
verifying ASR recovery 742
versions of gstd.log 512
view
document view 434
viewing
annotation 474
enterprise hierarchy 518
reports 434
volume pool
See also pools
archive 226
defined 203
Volume Shadow Copy Service
commands 752
controlling from Administration window 749
controlling from command-prompt 752
controlling from Sun StorEdge EBS client 751
limitations with Microsoft Exchange Server 370
limitations with Microsoft SQL Server 370
overview 745
writers 748
volumes
archive 258
client file index
removing 562
cloning 246
archive data 258
creating 246
performance 244
recovery 257
duplicating, problems with 245
labeling
how Sun StorEdge EBS uses 217
maximum size 222
tips 222
modes
types 333
nonrewinding 775
recoveries, required for 378
recycling 561
relabeling 328
removing 561
save sets
recoveries 382–383
verify 72
VSS. *See* Volume Shadow Copy Service

W

waiting priority 472, 480, 486
warning priority 472, 480, 486
web
browser
unresponsive 794
Console
server name 35
window
managing 31
WINS (Windows Internet Naming Service) 540

WINS (Windows Internet Naming Service)
database 84, 411
winworkr program. *See* Sun StorEdge EBS User program
wizard
client configuration 62
console configuration 514
writer
about 748

X

xview, errors 769

Z

zone, control 515