



Sun StorEdge Enterprise Backup Software™ 7.1

Administrator's Guide

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

Part No. 817-3216-10
September 2003

© 2003, LEGATO Systems, Inc. All rights reserved.

This product or document is distributed under licenses restricting its use, copying, distribution, and decompilation. No part of this product or 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 other countries, exclusively licensed through X/Open Company, Ltd.

Sun, Sun Microsystems, the Sun logo, AnswerBook2, docs.sun.com, Solaris, and StorEdge Enterprise Backup are trademarks, registered trademarks, or service marks of Sun Microsystems, Inc. in the U.S. and other countries. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the U.S. and 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 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.

Federal Acquisitions: Commercial Software—Government Users Subject to Standard License Terms and Conditions.

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 LEGATO Systems, Inc., 2350 W. El Camino Real, Mountain View, CA 94040 Etats-Unis. Tous droits réservés.

Ce produit ou document est 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 en a. 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, Solaris, et StorEdge Enterprise Backup sont des marques de fabrique ou des marques déposées, ou marques de service, 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 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 licence non exclusive de Xerox sur l'interface d'utilisation graphique Xerox, cette licence couvrant également les licenciés 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'ETAT" 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.



Contents

Preface xlvii

Before You Read This Book xlvii

Using UNIX Commands xlvii

Typographic Conventions xlviii

Shell Prompts xlviii

Related Documentation xlix

Accessing Sun Documentation Online xlix

Sun Welcomes Your Comments xlix

1. Introduction 1

Sun StorEdge EBS Storage Management 2

 Cross-Platform Management 3

 Performance 4

 Optional Sun StorEdge EBS Software Additions 4

 Ease of Use 4

 Scalability 5

How the Sun StorEdge EBS Software Works 5

 Sun StorEdge EBS Resources and Attributes 7

 Sun StorEdge EBS Daemons and Programs 9

The Master Sun StorEdge EBS Server Process	10
How the Sun StorEdge EBS Software Backs Up Data	11
How the Sun StorEdge EBS Software Recovers Data	13
Sun StorEdge EBS Operations with Storage Nodes	15
Recovery Operations with Storage Nodes	16
Sun StorEdge EBS Product Family	17
Server Software Products	17
Client Software Products	18
Other Sun StorEdge EBS-Related Products	19
Sun StorEdge EBS Documentation	20
Documentation Suite CD-ROM	20
Printed Sun StorEdge EBS and Legato NetWorker Product Manuals	20
Administrator's Guide	20
Installation Guide	21
Release Supplement	21
Fixed Bugs List	22
Online Help	22
Man Pages	22
Disaster Recovery Guide	23
Performance Tuning Guide	23
Error Message Guide	23
Legato Compatibility Guides	23
Command Reference Guide	24

2. Getting Started 25

Preparing to Use the Sun StorEdge EBS Software	25
Operating Sun StorEdge EBS Software for the First Time	26
Operating Tools	26

▼	To Start the Sun StorEdge EBS Administrator Program	28
▼	To Label and Mount a Volume	28
	Backing Up Data	29
▼	To Enable a Scheduled Backup	29
▼	To Override a Scheduled Backup	30
▼	To Customize a Backup	30
▼	To Add More Clients	32
	Verifying Operations	32
▼	To Monitor and Verify Data during a Backup	32
▼	To View the Group Control Details Dialog Box	33
▼	To Verify Backup Servers	33
▼	To Verify That the Sun StorEdge EBS Processes Are Running	33
	How to View a Message	34
	Customizing and Configuring the Sun StorEdge EBS Software	35
	Preconfigured Settings	35
	Configuration Planning	35
	Sun StorEdge EBS Administrator Interfaces	38
	Administrator Privileges	39
	TCP/IP Network Requirements	39
	Sun StorEdge EBS Administrator Program	39
	Sun StorEdge EBS Administrator Graphical Interface	39
	Customizing the Sun StorEdge EBS Administrator Program Interface	42
	Foreign Languages in the Interface	43
	Sun StorEdge EBS Resources	44
	Sun StorEdge EBS Character-Based Interface	44
	Command Line Interfaces	44
	Sun StorEdge EBS Manual Backup, Recovery, and Archiving Interfaces	45

The nwbackup Program	45
The nwrecover Program	46
The nwarchive Program	48
The nwretrieve Program	49
Rehosting a Sun StorEdge EBS Server	51

3. Server Management 53

Sun StorEdge EBS Server Connectivity	53
▼ To Establish a Server Connection	54
▼ To Close the Connection	54
Client/Server Communication Configuration	55
Server Configuration	55
▼ To Set Up the Server	55
▼ To Determine the Current User, Privileges, and User Group Membership	56
Performance Controls	57
Parallelism	57
Target Sessions	58
Server Access Management	59
Administrator Privileges	59
Example: Using the Administrator Attribute	61
Sun StorEdge EBS User Groups	61
Preconfigured User Groups	61
Administrators	62
Users	62
Customizing Privileges Associated with Sun StorEdge EBS Users	62
▼ To Edit a User Group	64
Notifications	65
Preconfigured Notifications	65

Customizing Notifications	66
Events	67
Actions	67
Priorities	68
▼ To Create a Notification	68
▼ To Edit a Notification	69
▼ To Delete a Notification	69
Reports	70
Routine Data Movement Operations Reports	70
Sun StorEdge EBS Server Bootstrap Report	70
Storage Management Application Reports	71
Generating Program-Readable Reports	71
Sun StorEdge EBS Server Statistics and Diagnostic Reports	74
Message Log Files	74
Monitoring and Reporting Changes to Sun StorEdge EBS Resources	74
▼ To Enable the Monitor RAP Attribute	74
Firewall Support	75
Configuring Firewall Support	75
▼ To Configure Firewall Support by Using Default Settings	76
Restricting the Port Range	76
▼ To Determine Minimum Port Range Required	76
▼ To Change Sun StorEdge EBS Default Settings	80
▼ To Complete the Configuration of Firewall Support	81
Sample rpcinfo Output	81
Troubleshooting Port Configuration Problems	83
Storage Nodes	83
Prerequisites	84

Licensing	84
Storage Node Configuration	85
Timeout Attribute for Storage Node Operations	85
Storage Node Affinity	86
Bootstrap Backup on a Storage Node	87
Storage Nodes and Cloning	87
Storage Node Troubleshooting Tips	87
Online Indexes	88
Characteristics of the Online Indexes	89
Automated Index Activities	90
Checking Online Indexes	90
▼ To View Information about the Indexes	90
Index Save Sets	92
▼ To View Information about a Client's Save Sets	92
Save Set Details	93
▼ To View Save Set Details	93
Location of the Client File Index	94
▼ To Designate the Client File Index Location	94
▼ To Move a Client File Index	95
Managing the Size of the Online Indexes	96
Reducing Client File Index Size	96
Reducing Media Database Size	97
▼ To Remove the Oldest Save Set Cycles	98
Removing Volume-Based Entries from the Online Indexes	98
Removing Client File Index Entries	98
Removing Client File Index and Media Database Entries	99
▼ To Remove Online Index Entries	99

▼	To Compress the Media Database	100
	Server Maintenance Tasks	100
	Message Log Management	100
▼	To Set the Startup Script to Trim Log Files	101
4.	Client Operations	103
	Sun StorEdge EBS Client	104
	Client Configuration	104
	Client Name	105
	Save Sets	105
	Predefined Save Sets	106
	The All Save Set	107
	SYSTEM and VSS SYSTEM Save Sets	108
▼	To Create a Client	110
▼	To Edit a Client	112
▼	To Delete a Client	113
▼	To Restore a Deleted Client	113
	The Client ID	113
▼	To Access the Client ID	114
▼	To Change a Client Name	114
	Editing a Client's nsrla.res Database	115
▼	To Edit the nsrla.res Database	115
▼	To Display Hidden Attributes in the nsradm Program	116
	Multiple Clients from the Same Computer	116
	Using Unique Combinations of Clients and Save Sets	117
	Client and Save Set Combination	118
	Save Set Status Values	118
▼	To View the Status of a Save Set	119

▼ To Change the Status of a Save Set	119
Browse and Retention Policies	120
How the Sun StorEdge EBS Server Uses Browse and Retention Policies	121
How the Browse Policy Works	122
Example: One-Week Browse Policy	122
Example: Two-Week Browse Policy	123
How the Retention Policy Works	124
Example: Three-Week Retention Policy	125
Managing the Data Life Cycle	127
Assigning Multiple Policies to a Single Client	129
Example: Assigning Different Policies for an Identical Client	129
Preconfigured Policies	129
▼ To Create a Policy	130
▼ To Edit a Policy	131
▼ To Delete a Policy	131
Snapshot Policies	132
Browse and Retention Policies for Manual Backups	132
Browse and Retention Save Set Policy Modification	133
Example: Changing Browse and Retention Policies with nsrmm	134
▼ To Modify Manual Backup Policies	134
Reports on Browse and Retention Policies for Save Sets	135
Granularity	135
Overriding Browse and Retention Policies for Save Sets	135
Backup Command Customization	136
Using the save Command with a Customized Backup Program	137
▼ To Create a Custom Backup Script by Using the save Program	138
Example: The save Backup Command	138
Using the savepnpc Command with a Customized Backup Program	140

Timeout Attribute	142
▼ To Customize the savenpc Command for Multiple Groups	143
Message Logging by the savenpc Command	143
Directives	144
Local and Global Directives	144
Application Specific Modules	145
Password ASMs	148
Important Precautions When Using rawasm	148
Directive Format	149
Directory Specifications	149
Application Specific Module Specifications	150
Save Environment Keywords	151
Preconfigured Directive Resources	151
▼ To Create a Directive Resource	153
▼ To Edit a Directive Resource	153
▼ To Delete a Directive Resource	153
Adding or Changing the Sun StorEdge EBS Servers for a Client	154
▼ To Edit the servers File	154
Client Priority	155
Logical Volume Backup	156
Sun StorEdge EBS Client Licensing	157
The nsrlic Program	157
Client Licensing Error Messages	158
Legato NetWorker ClientPak for Mac OS X	159
Mac Legacy Files Supported	159
Supported Filesystems	159
Known Limitations	159
Backing Up a Sun StorEdge EBS Client on Mac OS X	160

Performing Backups of Data on Mac OS X Clients	160
MAC OS X Required Directives	161
Backing Up the NetInfo Database for Disaster Recovery	162
Restarting the Sun StorEdge EBS Client on Mac OS X	163
▼ To Shut Down the Sun StorEdge EBS Client Daemon on Mac OS X	163
▼ To Start the Sun StorEdge EBS Client Daemon on Mac OS X	163
Recovering Individual Files and Directories	164
5. Device and Media Management	165
Pools	166
How the Sun StorEdge EBS Server Uses Pools	166
Sun StorEdge EBS Pool Types	167
Preconfigured Pools	167
Default Pool	168
Example: Using Pool Configurations to Sort Data	168
How the Sun StorEdge EBS Server Uses Pool Criteria to Sort Data	169
Directing Client Indexes and Bootstrap to a Separate Pool	170
Meeting the Criteria for More Than One Pool Configuration	171
Not Meeting the Criteria for Any Pool	172
Configuring Pools	173
Incremental Backups	173
Manual Backups	173
Clone Data	174
Archive Data	174
Migration Data	175
Using Storage Devices and Pool Configuration to Sort Data	175
Using Different Media	175
Using a Specific Device for Backup Data	175

▼	To Create a Pool	176
	Expression Matching of Save Sets to Pools	177
	Auto Media Verification	178
▼	To Edit a Pool	179
▼	To Delete a Pool	179
▼	To Create an Archive Pool	180
	Volume Labeling	180
	How the Sun StorEdge EBS Server Uses Label Templates	180
	Preconfigured Label Templates	181
▼	To Complete Label Template Attributes	182
	Label Template Naming Strategies	184
	Labeling Tips	185
▼	To Create a Label Template	186
▼	To Edit a Label Template	186
▼	To Delete a Label Template	187
	Devices	187
	Device Support by Sun StorEdge EBS Edition	187
▼	To View Supported Devices and Media Types	188
	Configuration of the Sun StorEdge Network Foundation HBA / Driver	189
	Sample /usr/sbin/nsr/lus_add_fp_devs Script Output	190
	Nonrewinding Tape Device Usage	192
	Block Size Mode	192
▼	To Add a Device	192
▼	To Edit a Device	193
▼	To Delete a Device	194
	Device Parameter Environment Variables	194
	NSR_DEV_BLOCK_SIZE_MEDIA_TYPE	195

NSR_DEV_TAPE_FILE_SIZE_MEDIA_TYPE	195
NSR_DEV_DEFAULT_CAPACITY_MEDIA_TYPE	196
NSR_DEV_LOAD_TIME_MEDIA_TYPE	196
NSR_DEV_LOAD_POLL_INTERVAL_MEDIA_TYPE	196
NSR_DEV_LOAD_TRY_LIMIT_MEDIA_TYPE	196
Common Device Interface	197
Dynamic Drive Sharing	198
Introduction to DDS	198
Benefits of DDS	199
DDS Configuration	199
Block Size Compatibility	200
Block Size Incompatibility between UNIX and Microsoft Windows	201
SCSI Reserve/Release Setting	201
Enabling DDS within NDMP	201
DDS on NDMP Nodes in a SAN Environment	202
DDS Attributes in the Devices Dialog Box	203
Hardware ID Attribute	203
Shared Devices Attribute	203
▼ To Create a Device and Enable DDS	204
▼ To Stop Sharing a Device through the Devices Window	204
▼ To Change the Idle Device Timeout Attribute	204
High Availability and DDS	205
The nsrjb Program High Availability Limitations	205
Example: Host Crash Requires User Intervention	205
DDS Licensing	206
Sun StorEdge EBS DiskBackup	206
File Type Device	206

Advanced File Type Device	207
Maximum Number of Save Sets for Advanced File Types	208
Differences between File Type and Advanced File Type Devices	209
Advanced File Type Device Concurrent Operations	211
Differences in the Cloning Process	212
Example: Manual Cloning with Advanced File Type Device	212
Changes in How Sun StorEdge EBS Administrator Displays	213
Propagating Changes from Primary to Read-Only	215
Configuration Support for Sun StorEdge EBS DiskBackup	216
Reclaiming Space for Sun StorEdge EBS DiskBackup	218
Configuring Sun StorEdge EBS DiskBackup	218
Configuring a CIFS-Mounted Advanced File Type Device on a Microsoft Windows Storage Node	221
▼ To Configure a CIFS-Mounted Advanced File Type Device on a Microsoft Windows Storage Node	222
Sun StorEdge EBS DiskBackup Licensing	224
Sun StorEdge EBS DiskBackup Licensing Error Messages	225
Media Handling Error Resolution	225
▼ To Limit Repeated Device Errors	226
Auto Media Management	226
Storage Nodes	228
▼ To Configure a Storage Node	228
Administrator Privileges	229
Installing the Software on a New Storage Node	229
Installing an Update Only on the Sun StorEdge EBS Server	229
Installing an Update on an Existing Storage Node	229
▼ To Configure Autochangers and Silos as Storage Node Remote Devices	230
▼ To Set Timeouts for Storage Node Remote Devices	230

Dedicated Storage Node	231
Storage Management Operations	231
How the Sun StorEdge EBS Server Uses Volume Labeling	231
How the Sun StorEdge EBS Server Uses Pools	232
How the Sun StorEdge EBS Server Selects and Mounts a Volume during a Backup	232
How the Sun StorEdge EBS Server Selects a Volume for Relabeling	233
Volume Status	234
Read-Only Mode	235
Sun StorEdge EBS Volume Operations	236
▼ To View Volume Information	236
▼ To View Save Set Information	237
▼ To Mount a Volume	238
▼ To Label and Mount a Volume in One Operation	239
▼ To Label without Mounting	240
▼ To Unmount a Volume	240
▼ To Find a Volume Name	241
▼ To Change the Mode of a Volume	241
Marking a Volume as Full for Off-Site Storage	242
Removing Bad Volumes from the Online Indexes	242
Maximum Number of Save Sets Per Volume	243
Save Set Staging	244
Staging Data from a File Type or adv_file Device	245
Staging and Cloning from the Command Line	245
▼ To Create a Staging Policy	246
▼ To Edit a Staging Policy	248
▼ To Delete a Staging Policy	248

6. Backup Groups and Schedules 249

Overview of Sun StorEdge EBS Scheduling 249

Group Resource 250

Schedule Resource 250

Backup Groups 250

How the Sun StorEdge EBS Server Uses Groups 251

Example: Using Groups to Back Up Multiple Clients 251

Preconfigured Groups 252

Key Group Attributes 253

Aborted Group Backups 253

▼ To Create a Group 254

▼ To Set the Group Backup Time Interval 255

▼ To Force an Incremental Backup 256

▼ To Edit a Group 256

▼ To Delete a Group 257

Group Backup Management 257

▼ To Start a Group Immediately 258

▼ To Stop a Group Immediately 258

▼ To Manually Restart a Group 259

▼ To View Group Control Details 259

▼ To Preview a Group Backup 260

Savegroup Completion Report 261

Bootstrap Generation and Printout 261

▼ To Save the Bootstrap to a File 261

▼ To Send the Bootstrap File by E-mail 262

System Console Log 262

Backing Up Open Files 262

Backing Up Open Files with VSS	263
Schedules	263
How the Sun StorEdge EBS Server Uses Schedules	264
Preconfigured Sun StorEdge EBS Schedules	265
Backup Cycles	266
Example: Weekly Backup Cycle	266
Example: Default Schedule with Multiple Clients	266
Example: Staggered Weekly Schedules for Multiple Clients	267
Schedule Planning Considerations	268
Scheduling Large Client Filesystems	269
Key Schedule Attributes	269
Resource Configuration Order for Schedules	271
▼ To Create a Schedule	271
▼ To Edit a Schedule	272
▼ To Delete a Schedule	272
Overriding a Client's Regular Backup Schedule	272
Backup Levels	273
How Sun StorEdge EBS Backup Levels Work	274
Example: Backup Levels (Part 1)	274
Example: Backup Levels (Part 2)	275
Example: Backup Levels (Part 3)	276
How the Sun StorEdge EBS Server Uses Backup Levels	276
Online Indexes and Backup Levels	277
Backup Levels for the Sun StorEdge EBS Online Indexes	278
Save Set Consolidation	278
How Save Set Consolidation Works	278
Requirements	279
Save Set Consolidation Usage	279

Limitations of Save Set Consolidation	280
Invoking Save Set Consolidation	281
▼ To Schedule Save Set Consolidation	281
▼ To Invoke Save Set Consolidation from the Command Prompt	281
▼ To Direct Data from a Consolidation Backup to a Specific Pool	282
Example: Consolidation of Level 1 and Full Backups	282
7. Backup and Recovery	285
Manual Data Backups	285
▼ To Start a Manual Backup	286
Backing Up Data Manually from the Command Prompt	287
Data Recovery	287
▼ To Recover Data	288
Recovering Data from the Command Prompt	290
Recovering Data from a Previous Backup	290
▼ To Change the Browse Time	291
Resolving Recovered Data Conflicts	291
Relocating Recovered Data	292
▼ To Relocate Recovered Data	292
▼ To View Versions of a Directory or File	293
▼ To Search for a Recoverable File	293
Viewing Details of Recoverable Files	293
▼ To View Details of Recoverable Files	293
Viewing the Volumes Required for Data Recovery	294
▼ To View Volumes Required for Data Recovery	294
Disaster Recovery	294
Directed Recovery	294
Uses of Directed Recovery	297

Access Requirements for Directed Recoveries	297
Browsing Files Before Performing a Directed Recovery	300
Performing a Directed Recovery	301
Using the recover Command in Directed Recoveries	301
Example: recover Command	302
Directed Recovery of Specific Files and Directories	302
Save Set Recovery	303
About Save Set Recovery	303
▼ To Recover an Entire Save Set	304
▼ To Recover Individual Files or Directories from a Save Set	305
▼ To View Volumes Required for Save Set Recovery	306
▼ To Relocate Recovered Save Set Data	306
▼ To Resolve Recovered Save Set Data Conflicts	307
▼ To Recover a Save Set from a Volume to the Sun StorEdge EBS Server	308
▼ To Recover One File Directly from a Volume	309
Comparison of Recovery Methods	309
Index-Based Recovery	309
Advantages of Index-Based Recovery	309
Disadvantages of Index-Based Recovery	310
Save Set Recovery	310
Advantages of Save Set Recovery	310
Disadvantages of Save Set Recovery	311
Deciding between Index-Based and Save Set Recovery	311
Recovering Online Indexes	312
▼ To Recover Client File Indexes	313
Recovering Client File Indexes from a Specific Time and Date	313
Example: Recovering a Client File Index from a Specific Time and Date	314

▼	To Restore a Save Set Entry in the Online Indexes	314
	Client File Index Only	314
	Client File Index and Media Database	315
	The scanner Program	316
	Importing Save Set Information with the scanner Program	317
▼	To Create a New Client Based on the Client ID	318
8.	Data Archiving	321
	Overview of Sun StorEdge EBS Archiving	321
	Archive Requirements	322
	How the Sun StorEdge EBS Server Archives Data	323
	Indexed and Nonindexed Archiving	324
	Indexed Archiving (Indexed Archive Pool, PC Archive Pool)	324
	Nonindexed Archiving (Archive Pool)	325
	Permissions for Archiving	325
	Enabling Archive Services for the Client	325
	Enabling or Restricting Archive Access	326
	Enabling Public Archive Access	326
	About Archive Pools	326
	Preconfigured Indexed Archive Pool and PC Archive Pool	327
	Preconfigured Archive Pool	327
	Creating Custom Archive Pools	327
	Archiving Data Procedures	328
	Enabling Archive Services for a Sun StorEdge EBS Client	328
	Manually Archiving Data	328
▼	To Request a Manual Archive from a Sun StorEdge EBS Client	329
	Manually Archiving from the Command Prompt	330
	Scheduling Archives	331

- ▼ To Schedule an Archive 331
- ▼ To Edit an Archive Request 332
- ▼ To Delete an Archive Request 333
- Retrieving and Recovering Archived Data 333
 - Retrieval Permissions 333
 - ▼ To Retrieve Nonindexed Archived Data 334
 - ▼ To Recover Indexed Archive Data 335
- Archive Request Management 335
 - ▼ To Start a Scheduled Archive at Any Time 336
 - ▼ To Stop a Scheduled Archive While in Progress 336
 - ▼ To Disable a Scheduled Archive 336
 - ▼ To View Details of a Scheduled Archive 337

9. Cloning 339

- Volume Cloning 339
 - How Volume Cloning Works 339
 - Volume Cloning Requirements 340
 - Volume Cloning Performance 340
 - Cloning Compared to Duplication of Volumes 341
 - Problems with Volume Duplication When Using a Copy Operation 341
 - Volume Cloning and Data Tracking Information 342
 - Changing the Mode of a Volume Clone to Recyclable 342
 - Purging a Clone Volume 342
 - Deleting a Clone Volume 343
 - Volume Cloning and Recovery 343
 - Creating Clone Volumes 343
 - ▼ To Create a Clone Volume 344

▼	To View Clone Volume Details	345
	Save Set Cloning	345
	Overview	345
	Retention Policies	346
	Performance Levels	346
	Online Index Entries	346
	Devices	347
	How Save Set Cloning Works	347
▼	To Manually Clone a Save Set	347
	Automatically Cloning a Save Set	350
	Examples of Cloning Save Sets	350
	Example: Manually Cloning a Save Set	351
	Example: Automatically Cloning One Client	352
	Example: Automatically Cloning Several Clients	353
	Recovering Cloned Save Sets	355
	Other Information and Procedures	356
▼	To Clone Archived Data	356
	Clone Storage Node Affinity	357
10.	Autochangers and Silos	359
	Autochanger Installation and Configuration	360
▼	To View Autochangers	360
▼	To Add an Autochanger	361
	Example: Configuring a SCSI Autochanger	362
	Example: Configuring an SJI Autochanger	364
	Configuring Autochangers on a Storage Node	367
▼	To Test the Autochanger Connection	367
	Testing Device Configuration Using the jbverify Program	368

Requirements for Running the jbverify Program	368
▼ To Test an Autochanger	369
▼ To Test a Stand-Alone Device	370
▼ To Test a Local Device	370
Testing Tapes with the jbverify Program	370
Configuring Autochanger Attributes	371
Key Autochanger Attributes	371
▼ To Add a Nonshared Device to an Autochanger	371
▼ To Edit Autochanger Attributes	373
▼ To Delete an Autochanger	374
▼ To Reset an Autochanger	374
Remote Autochanger Management	374
▼ To Add a Remote Autochanger Device	375
Adjusting Sleep Periods for Autochanger Tasks	375
▼ To Adjust Sleep Periods for Autochanger Tasks	376
Editing Autochanger Configurations	376
Using the jbedit Program	377
Limitations	377
Dynamic Drive Sharing	378
▼ To Enable DDS When Configuring an Autochanger	378
Sample jbconfig Output	379
▼ To Stop Sharing a Device within an Autochanger	382
▼ To Change the Idle Device Timeout Attribute	383
Sharing Autochangers among Sun StorEdge EBS Hosts	384
Autochanger Sharing Licensing and Availability	384
How Autochanger Sharing Works	385
Autochanger Sharing Configuration	385

Example: The jbconfig Program Additions	386
Device Naming during Configuration	386
Host Device Selection for Recoveries	386
Example: Data Recovery Using a Shared Autochanger	387
Host Device Selection for Cloning	388
Labeling, Loading, and Unloading Volumes on Shared Autochanger Devices	388
Managing Volumes on Server Devices	389
Managing Volumes on Storage Node Devices	389
Automatic Unmounting of Volumes	390
Setting Compatible Block Size for UNIX and Microsoft Windows	390
▼ To Determine the Allowable Block Size	390
Solutions	391
Autochanger Operations	391
▼ To Designate Available Autochanger Slots	392
Maintaining Autochanger Devices	393
Service Mode	393
▼ To Configure an Autochanger Resource for Cleaning	394
▼ To Use a Nondefault Slot for the Cleaning Cartridge	395
▼ To Replace Cleaning Cartridges	397
TapeAlert	397
Auto Media Management with Autochangers	399
How the Sun StorEdge EBS Server Selects a Volume for Writing	400
Mounting Autochanger Volumes	400
▼ To Mount an Autochanger Volume	401
▼ To Mount an Uninventoried Volume	401
▼ To Unmount an Autochanger Volume	402
▼ To Label an Autochanger Volume	402

Verifying the Label When Volume Is Unloaded	403
Labeling Volumes with Barcodes	404
▼ To Label a Volume with Barcodes	404
Barcode Labeling Tips	405
▼ To Use Unmatching Volume and Barcode Labels	406
Inventorying Autochanger Volumes	407
▼ To Inventory Autochanger Volumes	408
Using Autochangers with a Volume Import/Export Capability	409
▼ To Deposit a Volume by Using the Import/Export Feature	410
▼ To Withdraw a Volume by Using the Import/Export Feature	410
Tips for Using Autochangers	411
Autochanger Notifications	411
Recycling Compared to Adding More Volumes	412
Using Pools with Autochangers	412
Device Calibration	413
▼ To Disable the Element Status	413
Resetting an Autochanger	413
▼ To Reset an Autochanger	414
Optical Autochangers	414
Silos	415
How Sun StorEdge EBS Software Interacts with a Silo	415
▼ To Install a Silo	415
Silo Device Naming Conventions	417
StorageTek	417
IBM 3494	417
EMASS/Grau	418
▼ To Set up DAS to Work with Sun StorEdge EBS Software	418

▼	To Define a Silo Device	419
▼	To Release a Silo Device	419
	Silo Device Cleaning	420
▼	To Set Environment Variables for EMASS/ Grau and StorageTek Silos	420
	Media Management in a Silo	421
	Silo Slot Numbers	422
	Using Silos with a Volume Import/Export Capability	422
	Mounting and Unmounting a Silo Volume	423
	Labeling a Silo Volume	423
	Allocating a Silo Volume	424
▼	To Allocate a Silo Volume	425
	Deallocating a Silo Volume	425
▼	To Deallocate a Silo Volume	426
	Inventorying a Silo	426
11.	Sun StorEdge Enterprise Backup Software Cluster Support	429
	Cluster Terminology	429
	Clustering	430
	Cluster-Aware and Cluster-Unaware Sun StorEdge EBS Servers	431
	Sun Cluster Environment	432
	Backing Up Data in a Cluster Environment	432
▼	To Configure a Scheduled Save from a Physical Cluster Client	433
▼	To Configure a Scheduled Save from a Virtual Cluster Client	433
	Performing a Manual Backup from a Virtual Cluster Client	435
	Configuring an External Client to a Virtual Server	435
	Recovering Data in a Cluster Environment	435
▼	To Recover Data from a Private Disk Belonging to a Physical Client	435
▼	To Recover Data from a Shared Disk Belonging to a Virtual Client	436

Tracking Scheduled Saves	437
▼ To Force Save Sets to Use a Specific Index	437
Path-Ownership Rules	438
▼ To Check Path-Ownership Rules	438
Overriding Path-Ownership Rules	439
Using Autochangers and Stand-Alone Tape Devices	440
Configurations for Autochangers and Stand-Alone Tape Devices	440
Configuration Option 1	440
Configuration Option 2	443
Configuration Option 3	444
DDS within a Cluster	446
High Availability and DDS	446
Example: Host Crash Requires User Intervention	446
Controlling the Robotics Inside the Cluster Configuration	447
Using the jbconfig Program to Control the Robotics	448
Controlling the Robotics from a Stand-Alone Host	452
Controlling Robotics from a Sun StorEdge EBS Server on Another Cluster	453
Sun StorEdge EBS Cluster Licensing	454
Enabling the Highly Available Sun StorEdge EBS Server	455
Enabling Cluster and Virtual Clients	455
Cluster Clients	456
Virtual Clients	456
Licensing Requirements for Backing Up Cluster Clients	456
12. NDMP Support for the Sun StorEdge EBS Server	457
Overview of NDMP	457
NDMP Connection Support	458
NDMP Connection Software and Installation Requirements	458

Licensing	458
Supported NDMP Configurations	458
Local Backup	459
Three-Party Backup	459
Remote Backup	460
Remote Backup with SnapImage Software	460
Remote Backups with the DinoStor TapeServer	462
General Requirements for and Limitations of NDMP Support	463
Sun StorEdge EBS NDMP Requirements	463
Limitations of the Sun StorEdge EBS Software with NDMP Support	464
Protocol Limitations	465
Important Notes on Index Processing and Memory Requirements	465
Configuring the Sun StorEdge EBS Server for NDMP Operations	466
Configuring Autochanger Support for NDMP	467
Obtaining Information about NDMP Autochangers	468
Using the inquire Command with NDMP Devices	468
▼ To Configure an NDMP Autochanger	469
Sharing Autochangers among NDMP Clients	471
Running nsrjb to Load Tapes	472
Silo Support	473
▼ To Configure an NDMP Device Resource	473
▼ To Configure an NDMP Client Resource	474
Specifying Application Information	475
Additional Considerations for NDMP Operations	476
Force Incremental and Interval Attribute Settings	476
Group Parallelism Setting	476
Performing NDMP Backups	477
Prerequisites	477

Performing an NDMP Backup from the Command Prompt	478
▼ To Perform a Three-Party Backup	479
Performing DAR Backups	480
NDMP Cloning	481
NDMP Index and Bootstrap Cloning	481
Performing NDMP Recoveries	482
Performing a Destructive Save Set Recovery	482
▼ To Perform a Recovery from the Sun StorEdge EBS Administrator Program	483
▼ To Perform a Recovery from the Command Prompt	483
Performing a Directory-Level Recovery	484
▼ To Perform a Directory-Level Recovery from the Sun StorEdge EBS Administrator Program	484
▼ To Perform a Directory-Level Recovery from the Command Prompt	485
Performing a File-by-File Recovery	485
▼ To Perform a File-by-File Recovery from the Sun StorEdge EBS Recover Program	485
▼ To Perform a File-by-File Recovery from the Command Prompt	486
Performing a Directed Recovery	486
▼ To Perform a Directed Save Set Recovery	487
▼ To Perform a Directed Index Recovery	488
Vendor-Specific Information	489
Filers That Act as Data Servers	489
Network Appliance	490
Features and Limitations	490
Required NDMP Version When Running DAR on a NetApp Host	492
NetApp Zoning Requirements for DDS in a SAN Environment	493
▼ To Find the Device Pathname for a NetApp Host	493

Application Information Variables	495
EMC Celerra	496
EMC Celerra Considerations with NDMP	496
▼ To Find the Device Pathname for an EMC Celerra Host	498
Application Information Variables	499
EMC IP4700	499
▼ To Find the Device Pathname for an EMC IP4700 Host	499
Application Information Variables	500
Auspex	500
Auspex-Specific Considerations with NDMP	500
▼ To Find the Device Pathname for an Auspex Host	502
Application Information Variables	503
Procom NetFORCE	503
▼ To Find the Device Pathname for a Procom NetFORCE Host	503
Post-Configuration Requirements for Procom NetFORCE Host	504
Mirapoint	504
Requirements	504
Device Pathname for a Mirapoint Host	504
BlueArc	505
Filers That Act as Tape Servers Only	506
DinoStor TapeServer	506
Configuring the DinoStor TapeServer	506
▼ To Specify Device Handles for a DinoStor TapeServer	506
▼ To Perform a Remote Backup with a DinoStor TapeServer	507
SnapImage Module	507
Finding the Device Pathname for a SnapImage Module Device	507
▼ To Perform a Remote Backup with the SnapImage Software	509

13. SNMP Module 511

SNMP Traps 511

Configuring Sun StorEdge EBS SNMP Notifications 512

Command Line Options for nsrtrap 512

▼ To Modify the Preconfigured Sun StorEdge EBS SNMP Notification 513

▼ To Create Sun StorEdge EBS SNMP Notifications 514

Configuring SNMP Management Software 515

Sun StorEdge EBS SMI Network Management Private Enterprise Code 515

Receiving Traps in the SNMP Network Management Software 516

14. DiskXtender Data Manager Filesystem Support 517

Supported Configurations 517

Operating Systems 518

Filesystems 518

Sun StorEdge EBS Configurations 518

Path Information 518

Permissions 519

DiskXtender Data Manager Filesystem Overview 519

File Data in a DXDM Filesystem 519

User Access to File Data 520

Backup of DiskXtender Data Manager Filesystems 522

▼ To Backup a DXDM Filesystem 523

Excluding the DMAPI Attributes File 523

▼ To Exclude a DMAPI Attributes File 523

Aborted Backups 523

▼ To View Existing DMAPI Processes 524

▼ To Remove a Defunct DXULDMLIB Process 524

Sun StorEdge EBS Archive Feature 525

Recovery of DiskXtender Data Manager Filesystems	525
Initiating a Recovery	525
Recovered Files List	526
Restoring Deleted Files and Previous File Versions	526
▼ To Restore a File Version	526
Filesystem Synchronization	527
Automatic Synchronization	528
Manual Synchronization	528
A. Notes and Troubleshooting Tips	531
Before You Call Technical Support	531
Determining the Version of Sun StorEdge EBS Software Running on a Client	532
▼ To Determine the Client Software Version Using the Sun StorEdge EBS Administrator Program	532
▼ To Determine the Client Software Version Using nsradmin	532
The nsr_support Program	533
Sun StorEdge EBS Backup and Recovery	534
Multiple nsrexecd Processes	534
Checking the Sun StorEdge EBS Daemons	534
Unsynchronized Client Clock Errors	535
Backups Fail to Stop	536
Save Set Cutoffs for 5.x Clients	536
Memory Usage When Browsing Large Save Sets	537
Licensing-Based Errors with Release Clients Earlier Than 5.0	537
Check Failure of Client File Indexes	538
No Notification of Client File Index Size Growth	538
Media Position Errors Encountered When Auto Media Verify Is Enabled	538

PACKET RECEIVE BUFFER and NO ECB Counters Increase	539
The scanner Program Marks a Volume Read-Only	539
Index Recovery to a Different Location Fails	540
Potential Cause for Client Alias Problems	540
Illegal Characters to Avoid in Configurations	541
The scanner Program Requests an Entry for Record Size	541
Failed Recover Operation Directly after New Installation	542
Client File Index Is Missing Message	543
Recovering Files from an Interrupted Backup	543
Backup of a New Client Defaults to a Level Full	543
Renamed Clients Cannot Recover Old Backups	544
The savegrp Command Line Program Cannot Be Stopped from the Sun StorEdge EBS Administrator Program	544
▼ To Stop the savegrp Command Line Program	544
Disk Label Errors	545
Cannot Print Bootstrap Information	545
Server Index Not Forced	545
Copy Violation	546
Xview Errors	547
Limited Frame Buffer Capabilities May Cause nwadmin to Display Incorrectly	547
Converting Sparse Files to Fully Allocated Files	549
Cannot Manage Sun StorEdge EBS User Groups from Sun StorEdge EBS Release 6.x or Earlier	549
The mminfo -N Command is Case Sensitive Regarding Save Set Names	549
Client/Server Communications	550
▼ To Troubleshoot IP Errors	550
▼ To Set Up a Host Table	551

Troubleshooting Hostname Resolution Problems	552
▼ To Set the IP Name Search Order	553
▼ To Use ping to Verify Network Connections	553
▼ To Use rpcinfo to Verify That Sessions Can Be Established	554
Replacement for Functional Use of portmapper and rpcbind	555
▼ To Verify Firmware for Switches and Routers	556
Naming Requirements	556
Binding to Server Errors	557
Saving Remote Filesystems	557
Remote Recover Access Rights	558
Sun StorEdge EBS Archive and Retrieve	559
Remote Archive Request from Server Fails	559
Multiple Save Sets Appear as a Single Archive Save Set	559
Cloned Archives Do Not Appear in Sun StorEdge EBS Retrieve Program	559
Wrong Archive Pool Is Selected	560
Second Archive Request Does Not Execute	560
Command Line Archive Does Not Start Immediately	560
Filenames with Carriage Return not Allowed	560
Storage Nodes	561
Storage Node Affinity Errors	561
Storage Node Timeout Errors	561
Devices and Autochangers	562
Additional Attributes in the Jukebox Resource	562
Maintenance Commands	563
▼ To Test the Device Driver Installation	564
The lusdebug Command Fails	565

The jbxercise Command Fails	566
Autodetected SCSI Jukebox Option Causes Server to Hang	567
X11 Libraries Required for jconfig	568
Autochanger Inventory Problems	568
Destination Component Full Messages	569
Tapes Do Not Fill to Capacity	569
Server Cannot Access Autochanger Control Port	571
Nonrewinding Device Requirement	571
ATL 2640 Libraries	571
Depositing Volumes to Slots	572
Depositing Volumes to Slots on a Solaris Server or Storage Node	572
Using the -o Command Option for nsrjb	573
Scanner Command Behaves Differently with Adv_file Type Device	573
Misleading Error Message Displayed If Server and Remote Storage Node Clocks Are Not Synchronized	574
Sleep Times Required for TZ89 Drive Types	574
Message Displayed When CDI Enabled on NDMP or Disk File Type Device	575
Barcode Attribute for the Sun L280 Library	575
Diagnostic Tools	575
Vendor-Specific ASC and ASCQ Codes Supported in Sun StorEdge EBS Software	576
Sun StorEdge EBS Locale and Code Set Support	577
Locale Support	577
Code Set Support	578
Resource Database Notes	578
Viewing Resources	578
Restoring Old Resource Files	579

Error Logged If Old and New Resource Files Exist 579

Repairing Resource Database Corruption 580

B. UNIX and Linux Platform-Specific Notes 581

Solaris 581

Sun StorEdge EBS Executables Not Found in Expected Location for Solaris Client 581

▼ To Set the Executable Path 582

Linux 582

Configure Linux Operating System to Detect SCSI Devices 582

Linux Journaled Filesystem Support 583

HP-UX 583

Autochanger Installation on an HP-UX System 583

Selecting SCSI Addresses for the Autochanger 583

Installing the SCSI Pass-Through Driver 584

Installing the Autochanger Hardware 588

Major Number 590

Minor Number 590

▼ To Test the Device Driver and Device File Installation 591

▼ To Test the Autochanger Connection 592

“Hardware Address Not Found” Message 592

Errors from Unsupported Media in HP Tape Drives 593

Unloading Tape Drives on an HP-UX Storage Node 594

Symbolic Link Entries in the fstab File 594

AIX 595

▼ To Activate Full Diagnostic Core Dumps on AIX 595

STK-9840 Drives Attached to AIX 595

IRIX 596

IRIX Displays “No Space Left in Tape” Message 596

Glossary 597

Index 609

Tables

TABLE 1-1	Sun StorEdge EBS Resources	8
TABLE 1-2	Server Daemons and Programs	9
TABLE 1-3	Client Daemons and Programs	10
TABLE 1-4	Additional Sun StorEdge EBS Products	19
TABLE 1-5	Online Help Menu Options	22
TABLE 2-1	Sun StorEdge EBS Files and Directories	27
TABLE 2-2	Sun StorEdge EBS Processes	34
TABLE 2-3	Configuration Information	36
TABLE 2-4	Message Area	40
TABLE 2-5	Sun StorEdge EBS Administrator Toolbar Functions	41
TABLE 2-6	Polling Interval Options	43
TABLE 3-1	Maximum Parallelism Levels	58
TABLE 3-2	User Group Privileges	62
TABLE 3-3	Preconfigured Notifications	65
TABLE 3-4	Events	67
TABLE 3-5	Actions	67
TABLE 3-6	Priorities	68
TABLE 3-7	Storage Management Report Programs	71
TABLE 3-8	Options for the nsrports Command	81

TABLE 3-9	Sample of rpcinfo Output	82
TABLE 3-10	Clients Window Description	91
TABLE 3-11	Save Sets Window Description	92
TABLE 3-12	Instances Window Description	94
TABLE 4-1	Components in the All Save Set	107
TABLE 4-2	Components in the SYSTEM Save Sets	108
TABLE 4-3	Components in the VSS SYSTEM Save Sets	108
TABLE 4-4	Save Set Status Values	118
TABLE 4-5	Sun StorEdge EBS ASMs	145
TABLE 4-6	Save Environment Keywords	151
TABLE 4-7	Preconfigured Directives	152
TABLE 4-8	Mac OS X Files or Directories to Exclude from Backup	161
TABLE 5-1	Preconfigured Pools	167
TABLE 5-2	Example Attributes for Sending Bootstrap to a Separate Pool	170
TABLE 5-3	Sun StorEdge EBS Hierarchy for Resolving Pool Conflicts	171
TABLE 5-4	Key Label Template Attributes	182
TABLE 5-5	Examples of Number Sequences for Volume Labels	184
TABLE 5-6	Using Label Template Components	185
TABLE 5-7	Maximum Devices	188
TABLE 5-8	Shared Devices Attributes	203
TABLE 5-9	Advanced File Type and File Type Device Differences	209
TABLE 5-10	Devices Window Display	211
TABLE 5-11	Concurrent Operations with Advanced File Type Device	212
TABLE 5-12	Attributes Requiring Manual Propagation	216
TABLE 5-13	Volume Modes	235
TABLE 5-14	Volumes Window	237
TABLE 5-15	Volumes Window with Save Set Information	238
TABLE 6-1	Group Attributes	253

TABLE 6-2	Preconfigured Sun StorEdge EBS Schedules	265
TABLE 6-3	Key Schedule Attributes	270
TABLE 6-4	Backup Levels	273
TABLE 6-5	Advantages and Disadvantages of Backup Levels	277
TABLE 6-6	Regular File and Index Backup Levels	278
TABLE 7-1	Naming Conflict Resolution Options	292
TABLE 7-2	Access Permissions Used in Directed Recoveries	298
TABLE 7-3	Recover Options	301
TABLE 7-4	Naming Conflict Resolution Options	307
TABLE 7-5	Common scanner Options	316
TABLE 9-1	Save Set Status Settings	349
TABLE 9-2	Accounting Incremental Schedule	354
TABLE 9-3	Accounting Full-Clone Schedule	354
TABLE 10-1	The jverify Command Line Options for Devices	368
TABLE 10-2	Autochanger Attributes	371
TABLE 10-3	Autochanger Resource Sleep Attributes	376
TABLE 10-4	Common jbedit Options	377
TABLE 10-5	Notification Messages	395
TABLE 10-6	TapeAlert Severity	398
TABLE 10-7	EMASS/Grau and StorageTek Silo Environment Variables	420
TABLE 12-1	Vendor-Specific Information	463
TABLE 12-2	Configuring Resources for NDMP	467
TABLE 12-3	Command Line Options for nsrndmp_save	478
TABLE 12-4	Command Line Options for nsrndmp_recover	483
TABLE 12-5	Command Line Options for nsrndmp_recover	488
TABLE 12-6	NetApp Application Information Variables	495
TABLE 12-7	EMC Celerra Variables	499
TABLE 12-8	EMC IP4700 Variables	500

TABLE 12-9	Auspex Variables	503
TABLE 13-1	Command Line Options for nsrtrap	512
TABLE 14-1	DXDM Filesystem Data Types Included in a Sun StorEdge EBS Backup	522
TABLE A-1	Boot-Time File Locations	558
TABLE A-2	Additional Jukebox Attributes	562
TABLE A-3	Autochanger Maintenance Commands	563
TABLE A-4	Solaris Language Environment Settings for Supported Code Sets	578

Figures

FIGURE 1-1	Data Backup Flow	7
FIGURE 1-2	Daemons and Programs during a Scheduled Backup	11
FIGURE 1-3	Daemons and Programs during a Recover Session	14
FIGURE 1-4	Daemons and Programs during a Save Session with a Storage Node	16
FIGURE 1-5	Daemons and Programs during a Recover Session with a Storage Node	17
FIGURE 2-1	Resource Configuration Order	31
FIGURE 2-2	Administrator Program Interface	40
FIGURE 2-3	nwbackup Program	46
FIGURE 2-4	nwrecover Program	47
FIGURE 2-5	nwarchive Program	49
FIGURE 2-6	nwretrieve Program	50
FIGURE 3-1	Indexes Dialog Box	91
FIGURE 3-2	Instances Window	93
FIGURE 4-1	One-Week Browse Policy	123
FIGURE 4-2	Two-Week Browse Policy	124
FIGURE 4-3	One-Week Backup Cycle and Three-Week Retention Policy	126
FIGURE 4-4	Data Life Cycle in the Client Index and the Media Database	128
FIGURE 5-1	Using Pool Configurations to Sort Data	169
FIGURE 5-2	Labeling a Volume Using a Label Template	181
FIGURE 5-3	Dynamic Drive Sharing	200

FIGURE 5-4	DDS with NDMP	202
FIGURE 5-5	Device Display with Advanced File Type Device	213
FIGURE 5-6	Volume Display with Advanced File Type Device	214
FIGURE 5-7	Save Set Display with Advanced File Type Device	215
FIGURE 5-8	Sample Sun StorEdge EBS DiskBackup Configuration	218
FIGURE 5-9	CIFS-Mounted Advanced File Type Device	223
FIGURE 6-1	How Sun StorEdge EBS Groups Are Used to Back Up Multiple Clients	252
FIGURE 6-2	Weekly Backup Cycle	266
FIGURE 6-3	Using the Sun StorEdge EBS Default Schedule for Multiple Clients	267
FIGURE 6-4	Staggered Weekly Schedules for Multiple Clients	268
FIGURE 6-5	Backups for October 1 through October 7	275
FIGURE 6-6	Backups for October 1 through October 14	275
FIGURE 6-7	Backups for October 1 through October 15	276
FIGURE 6-8	Consolidated Backup	283
FIGURE 7-1	Browsing the Client File Index, Followed by Directed Recovery from a Remote Client	296
FIGURE 8-1	Overview of Archive Operation	324
FIGURE 9-1	Clone Volume Window	344
FIGURE 9-2	Save Set Clone Dialog Box	348
FIGURE 9-3	Automatic Cloning Scheme	353
FIGURE 10-1	Autochanger Device Pathnames	373
FIGURE 10-2	How Autochanger Sharing Works	385
FIGURE 10-3	Data Recovery Using a Shared Autochanger	387
FIGURE 11-1	Basic Cluster Configuration	431
FIGURE 11-2	Configuration Option 1	441
FIGURE 11-3	Configuration Option 2	443
FIGURE 11-4	Configuration Option 3	445
FIGURE 11-5	DDS Enabled with Robotics Controlled in the Cluster	448
FIGURE 11-6	DDS Enabled with the Robotics Controlled Outside the Cluster	453

FIGURE 11-7	Robotics Control between Clusters	454
FIGURE 12-1	NDMP Local Configuration	459
FIGURE 12-2	NDMP Three-Party Configuration	460
FIGURE 12-3	Performing a Remote NAS NDMP Backup to a Sun StorEdge EBS Server	461
FIGURE 12-4	Performing an NDMP Backup to an SnapImage Module Server	462
FIGURE 12-5	Remote Backups with the DinoStor TapeServer	462
FIGURE 14-1	Prototypical DXDM Installation	521

Preface

The *Sun StorEdge Enterprise Backup 7.1 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 7.1 Installation Guide*
- *Sun StorEdge Enterprise Backup 7.1 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

Typeface	Meaning	Examples
AaBbCc123	The names of commands, files, and directories; on-screen computer output	Edit your <code>.login</code> file. Use <code>ls -a</code> to list all files. % You have mail.
AaBbCc123	What you type, when contrasted with on-screen computer output	% su Password:
<i>AaBbCc123</i>	Book titles, new words or terms, words to be emphasized. Replace command-line variables with real names or values.	Read Chapter 6 in the <i>User's Guide</i> . These are called <i>class</i> options. You <i>must</i> be superuser to do this. To delete a file, type <code>rm filename</code> .

Shell Prompts

Shell	Prompt
C shell	<i>machine-name%</i>
C shell superuser	<i>machine-name#</i>
Bourne shell and Korn shell	\$
Bourne shell and Korn shell superuser	#

Related Documentation

Application	Title	Part Number
Installation	<i>Sun StorEdge Enterprise Backup Software 7.1 Installation Guide</i>	817-3217-10
Release Notes	<i>Sun StorEdge Enterprise Backup Software 7.1 Release Notes</i>	817-3218-10

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 (817-3216-10) of your document in the subject line of your email.

Introduction

The Sun StorEdge Enterprise Backup Software client/server environment provides the ability to protect your enterprise against the loss of valuable data. In a network environment, where the amount of data grows rapidly as computers are added to the network, the need to protect data becomes crucial. The Sun StorEdge Enterprise Backup Software product gives you the power and flexibility to meet such a challenge.

The Sun StorEdge EBS software is based upon a client-server architecture. All supported versions of Sun StorEdge EBS servers, storage nodes and clients are interoperable with each other, so you can have one network-wide backup and recovery solution for your entire network.

Note – For further details on supported versions of the Sun StorEdge EBS software, refer to the *Legato Compatibility Guides* available on the Legato web site at www.legato.com.

This chapter introduces the Sun StorEdge EBS software and provides background information, such as terms and key concepts, required to use and understand the product. This chapter also includes a summary of the printed and electronic documentation available to you. Subsequent chapters in this guide provide specific information on Sun StorEdge EBS interfaces, features, and methods for saving and recovering your data.

This chapter does not contain installation information. For instructions about installing the software, refer to the *Sun StorEdge Enterprise Backup Software Installation Guide*. The information contained in this Administrator's Guide assumes that the Sun StorEdge EBS software has already been successfully installed on your Sun StorEdge EBS server and clients.

The main sections in this chapter are:

- “Sun StorEdge EBS Storage Management” on page 2
- “How the Sun StorEdge EBS Software Works” on page 5

- [“How the Sun StorEdge EBS Software Backs Up Data” on page 11](#)
 - [“How the Sun StorEdge EBS Software Recovers Data” on page 13](#)
 - [“Sun StorEdge EBS Operations with Storage Nodes” on page 15](#)
 - [“Sun StorEdge EBS Product Family” on page 17](#)
 - [“Sun StorEdge EBS Documentation” on page 20](#)
-

Sun StorEdge EBS Storage Management

The Sun StorEdge EBS product is a network storage management application that backs up and recovers data across an entire network of computers. The software protects an organization’s valuable data by automatically backing up, storing, and indexing each individual data item (disk partition, directory, or file), enabling you to recover any particular data item, or even the complete contents of a client. You can also recover older versions of existing data items.

As the network and amount of data grow, Sun StorEdge EBS software provides the capacity and performance to handle the load. It offers a variety of features that give you the flexibility to sort your backed-up data by groups of specific clients, or onto specified storage volumes. It also provides regular backup status information regarding your data, media, and devices. Because the Sun StorEdge EBS index files that track the backed-up data and media can grow large over time, old and outdated entries are removed, based on your criteria, to keep the index size manageable.

Sun StorEdge EBS software features include a storage management application that directs high performance to a wide range of storage devices. The Sun StorEdge EBS software product is made up of the following components:

- Sun StorEdge EBS client software, which is a collection of processes and programs installed on the computers that contain data to be managed.
- Sun StorEdge EBS server software, which is a collection of processes and programs installed on the computer that performs the Sun StorEdge EBS services, including data backup, recovery, and archiving.
- Sun StorEdge EBS storage node software, which is a collection of processes and programs installed on computers that control their attached storage devices during Sun StorEdge EBS operations, including backup, recovery, and archive.

You can configure Sun StorEdge EBS software to back up data to storage devices that are directly connected to the Sun StorEdge EBS server, or you can create a Sun StorEdge EBS [storage node](#), which is a separate computer with one or more storage devices connected. The Sun StorEdge EBS server maintains the resources and files that keep track of Sun StorEdge EBS data. The storage node computer controls the storage devices and media.

You can direct backup, archive, and migration data to specific collections of volumes, according to particular data characteristics. As data is written to storage media, the Sun StorEdge EBS server creates a database to track the location of all managed data. The Sun StorEdge EBS server monitors data status and automatically rotates data through the tracking system as it reaches different levels of obsolescence, as defined by you.

You and your users can browse an index of recoverable files, create reports that describe the status of data or media, and recover data to a specified point in time. In response to a data recovery request, the Sun StorEdge EBS storage management system locates the volume that contains the requested data and either directs a device to mount the appropriate volume for recovery, or sends a message to the operator to request the volume by name. In this way, the Sun StorEdge EBS system manages all storage volume operations.

Sun StorEdge EBS software is managed and operated by using a combination of the following three methods:

- By using the Sun StorEdge EBS Administrator graphical interface to manage and configure scheduled backups, add clients and devices, and configure the Sun StorEdge EBS server. Using the Sun StorEdge EBS Administrator program, you can configure and perform backup and recovery tasks for any Sun StorEdge EBS server, client, or storage device on the network.
- By using the Sun StorEdge EBS Backup, Restore, Archive, and Retrieve graphical interfaces to manually back up, recover, and archive data between the client computer and the Sun StorEdge EBS server. Sun StorEdge EBS servers need to have these programs installed along with the Administrator program because the server acts as a client of itself.
- By running Sun StorEdge EBS command line utilities. The software requires that certain tasks, such as autochanger and silo device configuration, be done with command line utilities.

Note – The term "autochanger" refers to a variety of backup devices: autoloader, carousel, datawheel, jukebox, library, and near-line storage.

Cross-Platform Management

Sun StorEdge EBS software is a cross-platform storage management product. You can administer a Sun StorEdge EBS server from a computer running Microsoft Windows, UNIX, or Novell NetWare if the computer has the Sun StorEdge EBS client software installed on it and is connected to a network. Additionally, a Sun StorEdge EBS client on one platform can back up its data to a Sun StorEdge EBS

server on a different platform. However, to back up clients of other operating systems, the Sun StorEdge EBS server must have a ClientPak[®] enabler for that type of client; see your Installation Guide for further information.

You can direct and conduct administration services for any client or server on the network that has the Sun StorEdge EBS software installed. By displaying only the options that are relevant for that computer, the software frees you from the need to know the configuration of individual computers before you start a monitoring session.

Performance

Standard Sun StorEdge EBS performance features include:

- Server parallelism, which allows several *save streams* to flow to the server or storage node at the same time.
- Multiplexing, which allows more than one save stream to write to the same device at the same time.
- Client parallelism, which allows the client to send more than one save stream at a time.
- Session management, which allows you to control the number of save streams per device to maximize the performance of each device.
- Backup service to file devices and optional subsequent staging to near-line or offline volumes.

Optional Sun StorEdge EBS Software Additions

Optional Sun StorEdge EBS software additions include:

- Sun StorEdge EBS Autochanger Module
- Networker Silo Software Module
- Sun StorEdge EBS Archive Module
- Networker Database Modules (for backing up several types of databases)
- Sun StorEdge EBS SNMP (Simple Network Management Protocol)
- Legato GEMS SmartMedia

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 computer on the network.
- Grant permission to provide the capability for recovery of one client's data to another client computer.
- Obtain immediate answers to questions by accessing Online Help and 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 on the Sun web site 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 Legato NetWorker Module client software to back up databases and other non-filesystem 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.

How the Sun StorEdge EBS Software Works

Sun StorEdge EBS software works as a client/server system consisting of Sun StorEdge EBS servers, clients, storage nodes, storage devices, and storage volumes. The Sun StorEdge EBS *server* is a computer that receives backup data from Sun StorEdge EBS *clients*, indexes each data item, and sends the data to a storage device. Data from a single client computer backed up onto storage media is known as a *save set*. A save set typically comprises all of the backed-up data from a single filesystem or local disk. When the Sun StorEdge EBS server receives one or more client save

sets, the data is written to a storage volume (usually a tape). The Sun StorEdge EBS server also indexes the contents of each individual volume to facilitate future recovery.

The Sun StorEdge EBS *client* is any computer that has the Sun StorEdge EBS client software installed, and that backs up its save sets to a Sun StorEdge EBS server. A Sun StorEdge EBS client can be a workstation or a dedicated server. It can also have a different platform or operating system from the Sun StorEdge EBS server used for backing up its data. A Sun StorEdge EBS server is always a client of itself, and Sun StorEdge EBS servers can also back up to other Sun StorEdge EBS servers as part of a hierarchical backup structure.

The index of backed-up data items is called the *client file index*. The index of volumes is called the *media database*. Collectively, these indexes are called the *online indexes*, and they form the backbone of the Sun StorEdge EBS software's ability to locate and recover data.

When you submit a request to view or browse backup data, the Sun StorEdge EBS server looks to the client file index for a record of each backed-up data item. In cases where the data item has been backed up several times over a specified period, the client file index provides this information, as well as the option of recovering different versions of the same data item.

When you request recovery of a selected data item, the Sun StorEdge EBS server searches the media database to locate the specific volume or volumes containing the data items you want.

- If the storage device is a *stand-alone device* (a single-drive device that requires volumes to be manually inserted), the server requests that you insert the appropriate volume into the device.
- If the storage device is an *autochanger*, the server instructs the autochanger to locate the volume in the appropriate autochanger slot.

Once you select the data you want to recover and the Sun StorEdge EBS server locates the appropriate volume, the data is written back to the original client location; or in certain cases, such as *directed recovery*, to other selected client locations. In cases where recovered data might overwrite data already on the client, the Sun StorEdge EBS server alerts you and provides several options for solving the conflict. You can also configure in advance the action to take when recovery conflicts occur.

[FIGURE 1-1 on page 7](#) shows the Sun StorEdge EBS data flow, starting with the Sun StorEdge EBS client and ending with the storage volume.




FIGURE 1-1 Data Backup Flow

Sun StorEdge EBS Resources and Attributes

Sun StorEdge EBS performance and functionality are configured through Sun StorEdge EBS *resources*. Resources are the building blocks used by the Sun StorEdge EBS server to license, configure, and operate Sun StorEdge EBS clients and servers to back up, store, and recover data. Sun StorEdge EBS resources are represented in the Administrator program by various windows.

Most resources can be created, edited, and deleted by administrators with the proper permissions. Only members of groups with the Configure Sun StorEdge EBS privilege can update resource configurations; for more information, see [“Sun StorEdge EBS User Groups” on page 61](#).

The dialog boxes for these resources contain the attributes for creating and editing resources. Sun StorEdge EBS *attributes* are the individual components of a resource. Attributes are represented by the various text boxes, drop-down lists, and so on that are contained in a resource dialog box. For example, a client would be a Sun StorEdge EBS resource, and the client’s name would be an attribute of that resource.

The Sun StorEdge EBS product comes packaged with preconfigured and default resource settings. However, you can modify these resources to optimize the performance of your server. Many of the features and commands found in the Sun StorEdge EBS interface, and much of the information contained in this Administrator's Guide, pertain to resource creation and configuration.

[TABLE 1-1](#) lists Sun StorEdge EBS resources and their descriptions.

TABLE 1-1 Sun StorEdge EBS Resources

Resource	Action
Autochanger	Describes a device that can mount and write to multiple volumes automatically. See Device in this table. The term "autochanger" refers to a variety of backup devices: autoloader, carousel, datawheel, jukebox, library, and near-line storage.
Client	Identifies a Sun StorEdge EBS client, including data items such as filesystems, disk volumes, directories, or files that will be backed up, and assigns specific backup instructions to the data.
Device	Describes what is connected to the server or storage node.
Directive	Defines file processing activities that are performed during a backup of a client.
Group	Names and defines a unique collection of Client resources that participate in a group backup operation. Clients can follow different backup schedules within the group operation, but all clients in the group start the operation at the same time.
Label Template	Provides instructions for the Sun StorEdge EBS server to follow whenever the storage management application needs to create a label for a volume. Each pool is associated with a specific label template.
Notification	Names and defines a specific Sun StorEdge EBS event and an associated response, such as a message logged to a log file or sent to an administrator via e-mail.
Policy	Names and defines a time period that can be associated with administrator-specified data lifecycle policies.
Pool	Names a collection of volumes and specifies the acceptance criteria that backup data must meet to be written to this named collection.
Registration	Defines the features enabled for a particular Sun StorEdge EBS server.
Schedule	Names and defines a backup schedule that can be assigned to a Client or Group resource. The schedule dictates the days that the data is to be backed up and the level of each backup.
Server	Identifies the local Sun StorEdge EBS server characteristics. Only one Server resource exists per Sun StorEdge EBS server.
Staging	Names and defines the Sun StorEdge EBS policy for moving data from one storage medium to another, and for removing the data from its original location.
User Group	Allows Sun StorEdge EBS Administrators to assign users to specific groups and to define user privileges, such as access to Sun StorEdge EBS configuration and operation tasks.

Many of these resources also exist as attributes of other resources. These interdependent relationships affect the order in which you implement the configuration. Certain component resources must be defined before they can be listed as attributes of other resources. For example, to apply a custom Label Template resource to a pool, you must create the label template before configuring the pool.

For more detailed information on accessing and using Sun StorEdge EBS resources and attributes in the Sun StorEdge EBS Administrator program, see [“Sun StorEdge EBS Administrator Interfaces” on page 38](#).

Sun StorEdge EBS Daemons and Programs

Sun StorEdge EBS client/server technology uses a network remote procedure call (RPC) protocol to back up data. Sun StorEdge EBS software consists of several daemons and programs that oversee the backup and recovery processes, as well as storage management client configurations, a client file index, and a media database.

The Sun StorEdge EBS environment calls upon several daemons and programs when a request for a backup is received. The daemons coordinate the tasks associated with a backup or recovery, and record information about what was backed up, as well as the media that contains the backed-up data.

[TABLE 1-2](#) provides a description of the Sun StorEdge EBS server and storage node daemons and programs. The Sun StorEdge EBS man pages contain further details about the Sun StorEdge EBS daemons and programs.

TABLE 1-2 Server Daemons and Programs (1 of 2)

Daemon or Program	Function
ansrd	Monitors an active <code>recover</code> session. This agent process is spawned by <code>nsrd</code> in response to a <code>recover</code> session.
nsrck	Checks the consistency of the client file index. It is invoked by <code>nsrd</code> whenever the consistency of the client file index needs to be confirmed.
nsrd	Provides an RPC-based <code>save</code> and <code>recover</code> service to Sun StorEdge EBS clients. This is the master Sun StorEdge EBS daemon. For details regarding this daemon, see “The Master Sun StorEdge EBS Server Process” on page 10 .
nsrexec	Monitors the progress of individual save sets; agent process invoked by the <code>savegrp</code> program.
nsrim	Manages the server’s client file index; invoked at the end of an instance of the <code>savegrp</code> program. Note that <code>nsrim</code> is invoked only once within a twenty-four hour period, regardless of how many times <code>savegrp</code> is run in that same period.

TABLE 1-2 Server Daemons and Programs (2 of 2)

Daemon or Program	Function
nsrindexd	Provides a method for inserting entries in the client file index based on information passed by the save program.
nsrmmd	Provides device support, generates mount requests, and <i>multiplexes</i> save set data during a multiclient backup. The nsrd daemon can start several nsrmmd daemons, one for each enabled device. An additional nsrmmd daemon is started during pending mount requests. This is the media multiplexor daemon.
nsrmmdbd	Provides media and save set database management services to the local nsrd and nsrmmd daemon, and records entries in the media database. This is the media management database daemon.
savegrp	Runs a group of Sun StorEdge EBS clients through the save process.

The Master Sun StorEdge EBS Server Process

The nsrd process is the master Sun StorEdge EBS server daemon, responsible for the following tasks:

- Starting other daemons.
- Allocating media daemons on server and storage node computers.
- Authorizing backup and recover services for the client.
- Contacting clients for scheduled backups.
- Maintaining Sun StorEdge EBS configuration information.
- Monitoring backup and recover sessions.
- Maintaining server statistics and message logs.

The nsrd Sun StorEdge EBS server daemon calls on the nsrexecd Sun StorEdge EBS client daemon and several client-side programs when it receives a scheduled or manual backup request. The ansrd agent starts on the Sun StorEdge EBS server computer to monitor the progress of the backup session.

TABLE 1-3 describes the client-side daemons and programs.

TABLE 1-3 Client Daemons and Programs (1 of 2)

Daemons or Programs	Function
nsrexecd	Authenticates the Sun StorEdge EBS server's remote execution request and executes the save and savefs programs on the client.
recover	Browses the Sun StorEdge EBS server's client file index and restores the specified file to primary disk storage.

TABLE 1-3 Client Daemons and Programs (2 of 2)

Daemons or Programs	Function
save	Sends specified files in a multiplexed data stream to the Sun StorEdge EBS server for backup to media by <code>nsrmmd</code> and entry in the client file indexes and media database by <code>nsrindexd</code> .
savefs	Saves a filesystem (using the <code>save</code> program) to a Sun StorEdge EBS server; a level-based system is used to save only those files which have been modified since some previous save (a partial save).

How the Sun StorEdge EBS Software Backs Up Data

When you configure a *backup group* on the Sun StorEdge EBS server, you schedule a start time for the backup group. [FIGURE 1-2 on page 11](#) shows how all the Sun StorEdge EBS client and server daemons and programs interact during a scheduled backup.




FIGURE 1-2 Daemons and Programs during a Scheduled Backup

The following is the sequence of events that occurs during a scheduled backup:

1. The `nsrd` server daemon starts the server's `savegrp` program for the backup group at the scheduled time.
2. The `savegrp` program queries all Client resources configured on the Sun StorEdge EBS server to determine the following:
 - Which clients configured on the server are members of the scheduled group.
 - What level of backup (`save`) to perform.
 - How many save sets to run concurrently, as determined by the parallelism value set on the Sun StorEdge EBS server.
 - When the most recent backup of the group occurred.

If any of this information is not available on the Sun StorEdge EBS server, `savegrp` sends a request to the `nsrexecd` client-side daemon to run `savefs` on each client assigned to the backup group to gather the necessary details.

3. The `savefs` program tells `savegrp` which objects to back up for the client. After `savegrp` receives information about the objects to back up, `savegrp` assembles a work list for the server. The work list specifies the order in which clients are contacted for backup. The order of the work list is determined by the Client Priority attribute in the Client resource. The client with the lowest value in the Client Priority attribute is contacted first.
4. If problems were encountered with the client file index during the previous backup session, `nsrd` invokes the `nsrck` program to check the consistency and state of the Sun StorEdge EBS server's client file indexes and media database. Then `nsrd` starts the `nsrindexd` client file index insertion daemon.
5. The `savegrp` program contacts the first client on the server's work list. The client's `nsrexecd` is invoked and starts a `save` session of the first save set listed on the server's work list. The `save` program passes to `nsrd` all save criteria, such as group, client, save sets, storage node affinity, and level of the save data. With this information, `nsrd` determines the pool of volumes that will store the data and forwards the information to the appropriate media daemon, on either the Sun StorEdge EBS server or a storage node controlled by the Sun StorEdge EBS server.

As a result, the `nsrmmmd` media daemon:

- Sends a message to the console of the Sun StorEdge EBS server, requesting a mount of the media assigned to the volume pool indicated by `nsrd`.
 - Writes the data sent by `save` to storage media.
 - Forwards storage information to `nsrmmdbd` for recording in the Sun StorEdge EBS server's media database.
6. Any time there is a lull in `save` set activity from the client, the Sun StorEdge EBS server attempts to find another `save` set in the group to keep the process moving. The `savegrp` program attempts to concurrently back up as many `save` sets as possible, up to the limit set by the `Parallelism` attribute in the Sun StorEdge EBS server's configuration.
 7. The `savegrp` program repeats the process for each item on the server's work list until all clients in the group are backed up, at which point `nsrim` is invoked if it has not already been run within the last 24 hours. If the Sun StorEdge EBS server is part of the group being backed up, the Sun StorEdge EBS server's *bootstrap* file is backed up. The bootstrap file contains essential information for the Sun StorEdge EBS disaster recovery procedures. When the bootstrap backup is completed, a bootstrap printout is sent to the default printer configured for the Sun StorEdge EBS server. Keep the bootstrap printout in a safe place in case you need to restore the Sun StorEdge EBS server.
 8. The final results of the `savegrp` execution are sent back to the server and are included in the `savegroup` completion report, which is sent by e-mail to the root user.

How the Sun StorEdge EBS Software Recovers Data

FIGURE 1-3 shows how the Sun StorEdge EBS server and client daemons and programs interact while recovering data to a Sun StorEdge EBS client.




FIGURE 1-3 Daemons and Programs during a Recover Session

The following is the sequence of events that occurs during a recover session:

1. When the Sun StorEdge EBS server receives a **recover** request from a client, the server's **nsrd** daemon contacts the server's **nsrmmd** media daemon.
2. The **nsrmmd** daemon contacts the server's **nsrmmdbd** media database daemon to determine which media contains the save set requested by **recover**.
3. After the save set's media location is obtained, **nsrmmd** issues a mount request, the media is positioned to the beginning of the save set, and the save set stored on the mounted media is passed to **nsrmmd**.
4. The media daemon forwards the save set to the client's **recover** program, which restores the data to the client's filesystem.
5. When the server's **nsrmmdbd** media database daemon cannot locate the required volumes, or when there are no enabled devices on the Sun StorEdge EBS storage node or server, the **recover** request fails. The Sun StorEdge EBS software displays the following message on the client computer that requested the **recover**:

NSR server *client*: no matching devices on *server or storage node*

The Sun StorEdge EBS software also displays the following message for the Sun StorEdge EBS server:

```
media notice: no matching devices on server or storage node for recover  
by client client_name
```

```
media notice: enable or check device(s) on server or storage node
```

Sun StorEdge EBS Operations with Storage Nodes

A [storage node](#) is a computer that contains Sun StorEdge EBS client and media management daemons and programs. A storage node is connected to one or more storage devices that are used in Sun StorEdge EBS operations, such as backup and recovery. Media management daemons on the storage node computer read and write data on storage volumes.

Backup and recovery operations with storage nodes are very similar to backup and recovery operations on the Sun StorEdge EBS server. The major difference is where the data resides.

A storage node is controlled by a Sun StorEdge EBS server (or multiple Sun StorEdge EBS servers). Backup and recovery operations occur as follows:

1. The Sun StorEdge EBS server's nsrd master daemon starts an instance of nsrmmmd (media management daemon) on the storage node.
2. The Sun StorEdge EBS server routes appropriate data to each media management daemon according to the Storage Node attribute in the Client resource.
3. After data is written to the storage devices connected to a storage node, the Sun StorEdge EBS server tracks the location of the data.
4. On the Sun StorEdge EBS server:
 - The nsrindexd daemon writes entries in the client file indexes to track the location of each file in a save set.
 - The nsrmmdbd daemon writes entries in the media database to track the location of each save set on the media.

[FIGURE 1-4 on page 16](#) shows the client's save process contacting the storage node's nsrmmmd process.




FIGURE 1-4 Daemons and Programs during a Save Session with a Storage Node

Recovery Operations with Storage Nodes

Recovery operations that use a storage node are similar to the recovery operations that do not use a storage node (as described in [“How the Sun StorEdge EBS Software Recovers Data” on page 13](#)), *except* that when the Sun StorEdge EBS server receives a request from the client’s `recover` program, it contacts the `nsrmmmd` daemon on the storage node. The daemon reads the data from media and sends the data to the Sun StorEdge EBS client, where the `recover` program rebuilds the data.

[FIGURE 1-5 on page 17](#) shows an example of how a `recover` process works in a Sun StorEdge EBS setup that includes a storage node.




FIGURE 1-5 Daemons and Programs during a Recover Session with a Storage Node

Sun StorEdge EBS Product Family

Sun StorEdge EBS has a large product family of server, client, database, and related backup-and-recover software.

Server Software Products

All Sun StorEdge EBS clients, as well as Sun StorEdge EBS servers—which can function as their own clients and back up their own data—can be backed up by Sun StorEdge EBS server software.

(32-bit and 64-bit versions) Sun StorEdge EBS server products are available in four distinct versions offering different levels of support:

- Sun StorEdge EBS Workgroup Edition enables you to back up the server and up to seven clients. Sun StorEdge EBS Workgroup Edition does not include support for additional clients or optional software features, such as Sun StorEdge EBS

Archive. If you decide at a later date to purchase a higher level of Sun StorEdge EBS software, all of the data backed up by Sun StorEdge EBS Workgroup Edition is recoverable by any level of Sun StorEdge EBS software you purchase.

- Sun StorEdge EBS Business Edition is the same as the Workgroup Edition, except that it includes a license to use an autochanger (1-26 slots) as part of its base enabler, allows backing up a two-node cluster client, and allows NDMP connections.
- Sun StorEdge EBS Network Edition enables you to back up the server and an unlimited number of client connections. You can upgrade this edition to include support for autochangers, as well as optional software features.
- Sun StorEdge EBS Power Edition is a server software product that has been specially tuned to protect and provide backup for environments that support Very Large Database (VLDB) or large filesystem applications (in the terabyte range).

Client Software Products

Some Sun StorEdge EBS client software products are included with the server products. Refer to your *Sun StorEdge EBS Installation Guide* for any clients packaged with your particular server software and instructions on how to install them. Support for platform-specific clients is also sold separately as Legato NetWorker ClientPak software. Sun supports clients for the following platforms and operating systems:

- UNIX: Solaris, SunOS, HP-UX, HP Tru64, AIX, Linux, DYNIX/ptx, and IRIX
- PC Desktops: Microsoft Windows NT 4.0 Workstation, Microsoft Windows 2000 Workstation, Microsoft Windows XP Professional (32-bit and 64-bit versions)
- Microsoft Windows NT 4.0 Server, Microsoft Windows 2000 Server and Microsoft Windows Server 2003 (32-bit and 64-bit versions)
- Novell NetWare/IntraNetWare

For complete information on Sun StorEdge EBS client software, refer to the documentation in your Sun StorEdge EBS software media kit.

Other Sun StorEdge EBS-Related Products

[TABLE 1-4](#) lists other products offered by Sun StorEdge EBS to enhance your storage management needs:

TABLE 1-4 Additional Sun StorEdge EBS Products

Product	Description	For More Information
Sun StorEdge EBS Autochanger Module	Enables the Sun StorEdge EBS software to connect to one or more autochangers.	See Chapter 10 .
NetWorker Silo Software Module	Enables the Sun StorEdge EBS software to connect to one or more silos.	See Chapter 10 .
Sun StorEdge EBS Archive feature	Provides a sophisticated method to archive valuable data for long-term storage.	See Chapter 8 .
Sun StorEdge EBS SNMP Module	Provides communication of Sun StorEdge EBS event notifications to SNMP-compliant network management stations.	See Chapter 13 .
Legato NetWorker Modules	Enables reliable backup and restore capabilities for the most widely used databases, including Microsoft Exchange Server and SQL Server, Oracle, Informix, Lotus Notes, SAP R/3, and Sybase.	Refer to the specific documentation for each module.

Sun continues to develop enhancements to the Sun StorEdge EBS product line, including additional server, client, database, and add-on products. For the latest product information, refer to the Sun web site (www.sun.com).

Sun StorEdge EBS Documentation

The Sun StorEdge EBS product comes with a complete suite of documentation to provide the information you need to optimally run the Sun StorEdge EBS software. The following sections briefly describe the Sun StorEdge EBS documentation products.

Note – Sun StorEdge EBS documentation is periodically updated between releases, and the most recent versions are posted on the Legato web site.

Documentation Suite CD-ROM

All Sun StorEdge EBS product manuals and guides are provided in PDF format on a separate Documentation Suite CD-ROM as part of the product media kit.

Printed Sun StorEdge EBS and Legato NetWorker Product Manuals

You can also obtain a *Sun StorEdge EBS Documentation Kit*, including printed copies of the Sun StorEdge EBS manuals, for a nominal charge. Contact your Sun StorEdge EBS sales representative for information on ordering documentation kits for a particular Sun product.

Administrator's Guide

This Administrator's Guide is your main resource for understanding how the Sun StorEdge EBS product works. It provides information on a wide variety of Sun StorEdge EBS topics, including:

- How the Sun StorEdge EBS software saves and recovers data.
- Configuring Sun StorEdge EBS servers and clients.
- Understanding and using Sun StorEdge EBS resources.
- Configuring backup devices and media, including autochangers and silos.

- Using specialized storage management tools and protocols, including:
 - Archiving
 - Cloning
 - Staging
 - Network Data Management Protocol (NDMP)
 - SNMP
- Troubleshooting

Installation Guide

The *Sun StorEdge Enterprise Backup Software Installation Guide* provides the following information:

- How to install the software components on a server, client, and designated storage node computers.
- How to configure backup devices.
- How to update or upgrade existing Sun StorEdge EBS software.
- How to remove the software components.

Review the *Installation Guide* before you install the Sun StorEdge EBS software. After the installation has completed, keep a printed copy of the *Installation Guide* in a safe place, so that it is available if you need to reinstall or remove the software later.

Note – The *Installation Guides* for the Legato NetWorker ClientPak for NetWare and UNIX client software provide instructions for installing the client software on other operating systems that are not supported as Sun StorEdge EBS servers, but are supported as Sun StorEdge EBS clients.

Release Supplement

The *Sun StorEdge Enterprise Backup Software Release Supplement* provides late-breaking information on software bugs, workarounds, and other items not found in the other documentation sources. Depending on the Sun StorEdge EBS release, the *Release Supplement* may also contain information on new and improved features that were not included in this *Administrator's Guide*, or other documentation.

Updated versions of the *Release Supplement* are available in the Document Library on the Legato web site (www.legato.com). Refer to the web site periodically to view the latest version of this and all other Legato documentation. Compare the dates printed on each version of the Release Supplement to ensure you have the latest one.

Fixed Bugs List

The *Legato NetWorker Fixed Bugs List* includes a table listing all escalations and requests for enhancements (RFEs) that have been fixed in the current release.

Online Help

The Sun StorEdge EBS Administrator program and other graphical user interfaces include an Online Help component. The Online Help contains specific information about the user interface, including context-sensitive Help in each dialog box that describes the purpose of the displayed dialog box and what information to enter in each text box. There are also Help topics that describe Sun StorEdge EBS features as well as topics that provide step-by-step procedures for performing doing Sun StorEdge EBS tasks with the user interface.

To access the Online Help, select the Help menus displayed in the graphical user interface. [TABLE 1-5 on page 22](#) lists the Online Help menu options.

TABLE 1-5 Online Help Menu Options

Option	Description
On Window	Provides information for attributes in the current window.
On Topic	Provides a list of Help topics related to Sun StorEdge EBS software.
On Help	Provides information about how to use the Online Help.
On Version	Provides information on the particular release of the Sun StorEdge EBS software that is installed on your system.

The Sun StorEdge EBS Administrator also provides Help buttons in most of the dialog boxes. To display the Help topic associated with a particular Sun StorEdge EBS feature, click the Help button in the dialog box related to the feature.

Man Pages

The man pages are a reference for the syntax and function of the Sun StorEdge EBS commands that are issued from the command prompt. To view man pages, make sure that the \$MANPATH environment variable includes the path where your Sun StorEdge EBS man pages are installed. By default, Sun StorEdge EBS man pages are installed in /usr/share/man.

Disaster Recovery Guide

The *Legato NetWorker Disaster Recovery Guide* is a multiplatform guide that provides essential information about how to use Sun StorEdge EBS products to recover data from Sun StorEdge EBS servers in case of a disaster. This guide also contains tips and information about preventive measures to safeguard data and prepare backup policies to guard against a disaster.

Review the *Disaster Recovery Guide* when the Sun StorEdge EBS software is first installed and configured so that you are prepared with strategies to protect your data and operations. Keep a printed copy of the guide with other disaster recovery information, so that it is immediately available in case of an emergency.

Performance Tuning Guide

The *Legato NetWorker Power Edition Performance Tuning Guide* discusses the enhancements available with Sun StorEdge EBS Power Edition. It explains how to configure and measure the performance of the Sun StorEdge EBS Power Edition server.

Error Message Guide

The *Legato NetWorker Error Message Guide* describes messages and troubleshooting procedures for problems that may be encountered while using the Sun StorEdge EBS software.

Legato Compatibility Guides

The *Legato Compatibility Guides* provide the latest information on software and devices compatible with Sun StorEdge EBS software. These guides are available on the Legato web site at www.legato.com

Command Reference Guide

A complete set of the Legato online man pages is available to reference the syntax and function of the Sun StorEdge EBS commands that are issued from the command prompt.

The *Legato Command Reference Guide* describes the commands and resources that are involved in running Legato products.

For instructions on how to use the guide, including tips and troubleshooting, view the readme file in the `DOCS/cmdndref` directory on the Documentation Suite CD-ROM.

Getting Started

This chapter explains how to begin using the Sun StorEdge Enterprise Backup Software software to perform storage management activities.

The main sections in this chapter are:

- [“Preparing to Use the Sun StorEdge EBS Software” on page 25](#)
- [“Operating Sun StorEdge EBS Software for the First Time” on page 26](#)
- [“Verifying Operations” on page 32](#)
- [“Customizing and Configuring the Sun StorEdge EBS Software” on page 35](#)
- [“Sun StorEdge EBS Administrator Interfaces” on page 38](#)
- [“Sun StorEdge EBS Manual Backup, Recovery, and Archiving Interfaces” on page 45](#)
- [“Rehosting a Sun StorEdge EBS Server” on page 51](#)

Preparing to Use the Sun StorEdge EBS Software

If you are new to using Sun StorEdge EBS software, complete the following steps:

1. Review the Sun StorEdge EBS documentation, particularly the *Installation Guide* and the most recent *Release Supplement*.
2. Install the Sun StorEdge EBS server and client software.
3. Perform the tasks listed in [“Operating Sun StorEdge EBS Software for the First Time” on page 26](#). These tasks include:
 - Starting the Sun StorEdge EBS Administrator program.
 - Labeling and mounting a volume.

- Backing up data.
- Adding clients.

Once you have completed these tasks, learn how to monitor and verify data during a backup as described in [“Verifying Operations” on page 32](#), then continue to customize and configure the core Sun StorEdge EBS functions.

- Learn when and how to perform procedures for the following tasks described in [Chapter 7](#):
 - Manual backup
 - Data recovery
 - Directed recovery
 - Save set level recovery
 - Data recovery by using the scanner program
- Review the discussion of Sun StorEdge EBS browse and retention policies and online indexes (client file index and media database). See [“Browse and Retention Policies” on page 120](#).
- Review disaster recovery procedures. see [“Disaster Recovery” on page 294](#) and refer to the *Legato NetWorker Disaster Recovery Guide*.

Reviewing and working with this material helps to develop an understanding of how Sun StorEdge EBS software works, and of the basic functions and features that can be customized later to optimize performance and reliability.

Operating Sun StorEdge EBS Software for the First Time

This section provides an overview of the tools used to operate the Sun StorEdge EBS software. It also presents a series of steps to run basic Sun StorEdge EBS functions. Whenever possible, the tasks in this section use default and preconfigured settings. Throughout this *Administrator's Guide*, you can learn how to customize Sun StorEdge EBS resources and attributes, and how to configure the Sun StorEdge EBS software for your particular needs.

Operating Tools

To begin using Sun StorEdge EBS software, you need to:

1. Connect to a server.

2. Mount a volume for the backups.
3. Enable a configuration.

Then you can start a backup.

The primary graphical interface to operate Sun StorEdge EBS software is the Administrator program (`nwadmin`). You might also use any combination of the following:

- Sun StorEdge EBS backup program (`nwbackup`)
- Sun StorEdge EBS recover program (`nwrecover`)
- Sun StorEdge EBS archive program (`nwarchive`)
- Sun StorEdge EBS retrieval program, for retrieving archived data (`nwretrieve`)
- Command line interfaces

The combination of these tools enables you to:

- Perform scheduled and manual backups.
- Verify the servers in use.
- Verify that the necessary processes are running on their respective computers.
- Verify the success of a backup.
- Read message logs.
- Perform recoveries.

For a full description of these programs and their features, see [“Sun StorEdge EBS Storage Management” on page 2](#) and [“Sun StorEdge EBS Resources and Attributes” on page 7](#).

[TABLE 2-1](#) lists directories included with the Sun StorEdge EBS software, located in the Sun StorEdge EBS installation directory, usually `/nsr`.

TABLE 2-1 Sun StorEdge EBS Files and Directories

Directory	Files
<code>mm</code>	Sun StorEdge EBS media database
<code>res</code>	Sun StorEdge EBS resource files: • <code>nsr1a.res</code> • <code>servers</code>
<code>res/nsrdb</code>	Sun StorEdge EBS resource database
<code>index/client_name/db</code>	Client file index database
<code>logs</code>	Sun StorEdge EBS log files

As described in [Chapter 1](#), you configure performance and functionality through Sun StorEdge EBS resources. If you do not want to create a customized configuration, you need only enable a group to begin backups according to the default configuration.

To learn how to customize Sun StorEdge EBS resource settings, such as Group, Client, or Pool, see [“To Customize a Backup” on page 30](#). For more detailed explanations, see [“Customizing and Configuring the Sun StorEdge EBS Software” on page 35](#) and in subsequent chapters.

▼ To Start the Sun StorEdge EBS Administrator Program

To start the Sun StorEdge EBS Administrator program and connect to a server, enter the following command at the prompt:

```
# nwadmin &
```

▼ To Label and Mount a Volume

To perform backups, the backup media must be configured and available for the Sun StorEdge EBS server to write data to it. Once you have started the Sun StorEdge EBS Administrator program, label and mount a blank media volume as follows:

1. Insert a volume into the stand-alone device drive or slot 1 of the autochanger.

Note – Backing up data to an autochanger requires that you *first* install and enable the Autochanger Software Module.

2. From the Media menu, select Devices. to open the Devices window.
3. In the Devices window, click Create.
4. Enter the name of the device, and select the media type.
5. Click Apply.
6. Click Label.
7. From the Label dialog box, select a pool for the Pool attribute.
Selecting a pool automatically chooses a label for the volume.

8. If you back up to a stand-alone device, select Mount After Labeling in the Label dialog box.
9. If you back up to a device in an autochanger, enter 1 in *both* the First Slot and Last Slot attributes in the Autochanger Label dialog box.
10. Click Apply.

The Sun StorEdge EBS software labels the volume. If the storage device is a stand-alone device, the Sun StorEdge EBS software also mounts the volume and displays the name and status of the volume in the Monitor window. For example:

/dev/rmt/0cbn	venus.001	labeling, done

Backing Up Data

Sun StorEdge EBS software contains preconfigured resources you can start using immediately; therefore, when the volume is successfully labeled, you are ready to perform a backup.

There are two types of backups in the Sun StorEdge EBS environment:

- **Scheduled backups** — These are configured in the Sun StorEdge EBS Administrator program and are assigned a specific time for the backup to run, usually when network usage is low. For example, the preconfigured Default group is scheduled to begin its backup at 3:33 a.m. You can override scheduled backups to perform the backup immediately, if necessary.
- **Manual backups** — These are performed by the user of the client computer using the `nwbackup` program.

To run a scheduled backup, you need only enable the default schedule.

Sun StorEdge EBS release 7.1 software can use [Volume Shadow Copy Service \(VSS\)](#) technology to back up data. 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. For detailed information about VSS, refer to the *Legato NetWorker Administrator's Guide, Microsoft Windows Version*.

▼ To Enable a Scheduled Backup

To configure the Sun StorEdge EBS server for a scheduled backup:

1. From the Customize menu, select Groups to open the Groups window.
2. In the Groups attribute, select the Default group.

3. For the Autostart attribute, select Enabled.
4. Click Apply.

The Sun StorEdge EBS server is now enabled to back up itself, starting at 3:33 a.m. You can test the scheduled backup by overriding the scheduled time and starting the backup immediately.

▼ To Override a Scheduled Backup

To override a scheduled backup time and start the backup immediately:

1. From the Customize menu, select Groups to open the Groups window.
2. In the Groups attribute, select the Default group.
3. For the Autostart attribute, select Start Now.
4. Click Apply.

The Sun StorEdge EBS server immediately backs up itself.

Note – The scheduled backup will still occur at 3:33 a.m. But, by starting it now you override the backup schedule so that you can see how the Sun StorEdge EBS server performs an automatic backup.

▼ To Customize a Backup

You can customize a backup schedule and other resources to suit an environment. Sun StorEdge EBS resources should be configured in a specified order. This is because some resources contain information that will appear in the Sun StorEdge EBS Administrator as attributes of other resources, and resources that are also attributes of other resources should be defined first. After this initial setup, you can then modify the resources at any point.

Because the Schedule and Policy resources do not depend on any other resources, customize these resources first. Then move on to the Group, Client, and Pool resources. [FIGURE 2-1 on page 31](#) illustrates the process for configuring Sun StorEdge EBS resources.




FIGURE 2-1 Resource Configuration Order

To customize backup resources:

1. Create a customized Schedule, which becomes an attribute of the Group resource. For details, see [“Schedules” on page 263](#).
2. Create a customized Browse or Retention Policy, which becomes an attribute of the Client resource. For details, see [“Browse and Retention Policies” on page 120](#).
3. Customize the Group resource, which becomes an attribute of the Client resource. For details, see [“Backup Groups” on page 250](#).
4. Customize the Client resource, which can be an attribute of the Pool resource. For details, see [“Client Configuration” on page 104](#).
5. Customize the Pool resource. For details, see [“Pools” on page 166](#).

In [FIGURE 2-1](#), resources that do not depend on other resources are considered independent resources; they are also selected as attributes in those resources that depend on them. Note that the Directive, Schedule, and Policy resources must be defined before they are available as attributes in the Client resource.

▼ To Add More Clients

Because every Sun StorEdge EBS server is also a client to itself, you already have one client by default. Once the first backup of the Sun StorEdge EBS server has occurred, you should add clients to the server.

To add clients to the Sun StorEdge EBS server:

1. From the Clients menu, select Client Setup to open the Clients window.
2. Click Create.
3. In the Name text box, enter the new client's hostname.
4. Accept the default settings for the other attributes, or configure as appropriate.
5. Click Apply.

When a new client is added, the client's name and icon appear in the Clients window.

6. Repeat these steps for each Sun StorEdge EBS client to be backed up by the Sun StorEdge EBS server.

Verifying Operations

This section explains how to verify a backup. It also explains how to verify the authorized servers and that the Sun StorEdge EBS processes are running, as well as how to read message logs.

▼ To Monitor and Verify Data during a Backup

To monitor and verify data during a backup:

1. From the Customize menu, select Groups to open the Groups window.
2. In the Groups attribute, select the Default group.
3. For the Autostart attribute, select Start now.
4. Click Apply.

In the Messages window of the main Sun StorEdge EBS Administrator window, the Sun StorEdge EBS server displays detailed messages about the progress of the backup.

▼ To View the Group Control Details Dialog Box

You can also view information about the backup in the Group Control Details dialog box.

To open the Group Control Detail dialog box from the main window of the Sun StorEdge EBS Administrator program:

1. Select Group Control.
2. Select Details.

The Sun StorEdge EBS Administrator program lists the successful, unsuccessful, and incomplete save sets in the Group Control Detail dialog box.

Note – If you have a stand-alone storage device, make sure that you leave the volume in the device; otherwise, the Sun StorEdge EBS software will not be able to begin the scheduled backup for the Default group at 3:33 a.m.

▼ To Verify Backup Servers

You can verify the servers that are authorized to back up a client computer by checking the client's `/nsr/res/servers` file, which lists the servers that are authorized to back up the computer.



Caution – If the `/nsr/res/servers` file contains no data (no servers are listed), any server can back up that computer, unless the client daemon (`nsrexecd`) is started using the `-S server` option.

▼ To Verify That the Sun StorEdge EBS Processes Are Running

To verify that the Sun StorEdge EBS processes are running:

1. Enter *one* of the following at the command prompt:
 - `# ps -ef | grep nsr`
 - `# ps -ax | grep nsr`
2. Depending on the type of installation, the processes listed in [TABLE 2-2 on page 34](#) should be running:

TABLE 2-2 Sun StorEdge EBS Processes

Type of Installation	Processes
Sun StorEdge EBS Server	nsrd — Sun StorEdge EBS server nsrexecd — Sun StorEdge EBS client (two processes should be running) nsrmmd — Sun StorEdge EBS media multiplexer (one process per device) nsrmmdbd — Sun StorEdge EBS media management database nsrindexd — Sun StorEdge EBS file index database
Sun StorEdge EBS Client	nsrexecd — Sun StorEdge EBS client
Sun StorEdge EBS Storage Node	nsrexecd — Sun StorEdge EBS client nsrmmd — Sun StorEdge EBS media multiplexer (one process per device)

How to View a Message

Messages appear in the Sun StorEdge EBS Administrator program and are also written to message logs.

To view a message:

- Messages relating to the Sun StorEdge EBS server to which you are currently connected appear re displayed in the Messages window of the Sun StorEdge EBS Administration program. You can also use the command line program, `nsrwatch`, to monitor these messages.
- In message logs:
 - `/nsr/logs/daemon.log` (messages logged during a backup, and also displayed in the Sessions window of the Sun StorEdge EBS Administrator program).
 - `/nsr/logs/messages` (messages logged after a backup is complete, and also displayed in the Messages window of the Sun StorEdge EBS Administrator program).

Customizing and Configuring the Sun StorEdge EBS Software

Once you have used some of the basic features of the Sun StorEdge EBS software, you can start to customize and configure Sun StorEdge EBS settings for an environment. The Sun StorEdge EBS software enables you to configure some or all of its settings. You can also mix customized and preconfigured settings in any combination.

Preconfigured Settings

You can fully use and evaluate Sun StorEdge EBS software with the default settings that are automatically assigned to resources and attributes (for example, the Default pool). Whenever you create a new resource, you always have the option of using default attributes and preconfigured resources.

Configuration Planning

Configuration planning helps to clarify the interdependencies and trade-offs in implementing an overall backup strategy. Careful configuration planning is the most critical factor in ensuring that backups run smoothly and recoveries proceed efficiently and without error.

To plan a configuration:

1. Draw a diagram of the networked computers to be used as Sun StorEdge EBS servers and clients. Identify potential network trouble spots, such as computers on the other side of a firewall or routers that might not be able to successfully back up to the Sun StorEdge EBS server.
2. Estimate the size of the average daily backup window for the enterprise (the period of time set aside for backing up data on the network). Identify special circumstances that might alter the size of the backup window, such as weekend activity, operator availability (if a stand-alone device is being used), or end-of-quarter requirements.
3. Identify the amount of data that must be backed up by each computer and estimate the total amount of backed-up data. Estimate how long backups will take, given the characteristics of the specific environment (for example, network bottleneck speeds, device speeds, and data availability).

4. Organize and prioritize data into clusters that can begin backups at the same time of day (so that they can be organized into Sun StorEdge EBS backup groups). This step is most appropriate if backups require load balancing or staggered schedules to accommodate the constraints of the backup's allotted time.
5. Decide how long different data should remain in the storage management system for users to browse and recover. The availability of data files might need to be weighed against the space required to maintain the file entries within the client file index.
6. Decide how data should be organized within the storage management system. By using Sun StorEdge EBS pools, you can send specific save sets of data to designated collections of volumes. For example, one pool may be designated for all client backups of an organization's accounting department.
7. Identify the trade-offs and select optimal values for the following performance features:
 - Number of devices to use for backups
 - Number of devices to reserve for recoveries
 - Number of data streams to be interleaved onto a single volume
 - Number of data streams to send to each Sun StorEdge EBS server and storage node

[TABLE 2-3](#) summarizes the kinds of decisions you need to make prior to configuring Sun StorEdge EBS software. Decisions are identified in terms of the information required to perform a scheduled backup, and how that information is expressed in a Sun StorEdge EBS configuration.

TABLE 2-3 Configuration Information (*1 of 3*)

Type of Needed Information	How the Information Is Expressed in Sun StorEdge EBS Software
When does this backup start?	When the server computer's clock is equal to the time designated as the starting time for a backup group, the Sun StorEdge EBS server starts the backup.
	Sun StorEdge EBS resource: Group Resource attribute: Start Time
Which clients participate in the backup group?	All clients that are members of an executing backup group are considered candidates for backup. Client data might or might not actually be backed up, depending on the client's schedule. All new clients are, by default, members of the preconfigured group named Default.
	Sun StorEdge EBS resource: Client Resource attribute: Group

TABLE 2-3 Configuration Information (2 of 3)

Type of Needed Information	How the Information Is Expressed in Sun StorEdge EBS Software
How is client data listed for backup?	The Sun StorEdge EBS administrator clusters files, directories, and/or filesystems into packages organized by client computer. Each unique package is a <i>Client resource</i> ; multiple Client resources can be configured for one client computer. A <i>storage node</i> represents the data backed up for a Client resource during a single backup session.
	Sun StorEdge EBS resource: Client Resource attribute: Save Set
Where is client data stored?	Client data is written to the pool of storage volumes that is associated with the client's group. Preconfigured groups have been assigned to preconfigured pools.
	Sun StorEdge EBS resource: Pool Resource attribute: Group, Client, Save Set, Level
Which backup level is scheduled for this client's data?	Sun StorEdge EBS schedules specify a client's backup level for the day (or whether the backup skips the client that day). Group attributes, if specified, supersede a client schedule. Preconfigured schedules are shipped with Sun StorEdge EBS software.
	Sun StorEdge EBS resource: Client Resource attribute: Schedule <i>or</i> Sun StorEdge EBS resource: Group Resource attribute: Schedule <i>or</i> Sun StorEdge EBS resource: Group Resource attribute: Level
How long will the inventory of backed-up data be available for browsing?	The length of time that client data can be browsed via the <code>nwbackup</code> program is called the <i>browse policy</i> . The browse policy is assigned by the administrator for each Client resource. Preconfigured time periods are shipped with Sun StorEdge EBS software.
	Sun StorEdge EBS resource: Client Resource attribute: Browse Policy
How long will the data be available for recovery?	Data on storage volumes is protected from automatic overwriting for a period that is defined by the <i>retention policy</i> . The retention policy is assigned by the administrator for each Client resource. Data is available for recovery indefinitely, until the storage volume is physically relabeled. Preconfigured time periods are shipped with Sun StorEdge EBS software.
	Sun StorEdge EBS resource: Client Resource attribute: Retention Policy

TABLE 2-3 Configuration Information (3 of 3)

Type of Needed Information	How the Information Is Expressed in Sun StorEdge EBS Software
Is additional data processing needed?	A Sun StorEdge EBS <i>directive</i> can be applied to client data to perform commonly requested processing (such as compression and encryption). Preconfigured Directive resources are shipped with Sun StorEdge EBS software.
	Sun StorEdge EBS resource: Client Resource attribute: Directive
Is additional program execution needed for this environment?	Sun StorEdge EBS software can accommodate a customized script supplied by the administrator. For example, an environment might require that the client computer shut down its database before performing a backup.
	Sun StorEdge EBS resource: Client Resource attribute: Backup Command
Which users need access to manually back up, browse, or recover data?	Sun StorEdge EBS software provides manual backup, browse, and recovery access to all users on the client computer. (To recover data, the user must have read permissions to the data.) Access can be broadened to users on other client computers as well.
	Sun StorEdge EBS resource: Client Resource attribute: Remote Access

Sun StorEdge EBS Administrator Interfaces

You can administer a Sun StorEdge EBS server software by the following methods:

- Using the graphical Sun StorEdge EBS Administrator program (nwadmin). For more information, see [“Sun StorEdge EBS Administrator Program” on page 39](#).
- Using the character-based Sun StorEdge EBS administration interface (nsradmin). For more information, see [“Sun StorEdge EBS Character-Based Interface” on page 44](#).
- Using the command line. For more information, see [“Command Line Interfaces” on page 44](#).

Administrator Privileges

Only users included in the Sun StorEdge EBS Administrator list have permission to use and make changes to the Sun StorEdge EBS server. Users who do not have root privileges can start and use the Sun StorEdge EBS Administrator program, but are not able to make changes to the configuration. To learn more about adding administrators and administrator privileges, see [“Server Access Management” on page 59](#).

TCP/IP Network Requirements

The following are the requirements for TCP/IP network communication when using Sun StorEdge EBS software:

- The Sun StorEdge EBS server name must be added to the domain name system (DNS) file that contains the names and aliases of the servers (hosts) on your network, or to the local `hosts` file located in `/etc`.

Note – If DNS is used, reverse lookup must be correctly configured.

- If the Sun StorEdge EBS server is a dynamic host configuration protocol (DHCP) client, it must use a reserved address.
- If DHCP with dynamic addresses is used, DHCP must be synchronized with DNS.

Sun StorEdge EBS Administrator Program

The Sun StorEdge EBS Administrator program provides a graphical interface for configuring Sun StorEdge EBS servers and clients, assigning clients to servers, creating backup schedules, and managing other information.

The following sections describe the various windows, toolbars, and icons that help you perform these tasks.

Sun StorEdge EBS Administrator Graphical Interface

You can perform most Sun StorEdge EBS administration tasks in the Sun StorEdge EBS Administrator program. The Sun StorEdge EBS Administrator program contains the program menus, the Sun StorEdge EBS toolbar, and the windows shown in [FIGURE 2-2 on page 40](#).




FIGURE 2-2 Administrator Program Interface

Server Information

The server information area displays information regarding the server to which you are currently connected.

Message Area

[TABLE 2-4](#) lists the information displayed in the various message windows.

TABLE 2-4 Message Area (1 of 2)

Window	Description
Devices window	Displays information about the devices configured for the current server.
Sessions window	Displays progress messages. This information is also written to the <code>daemon.log</code> file in the <code>/nsr/logs</code> directory.

TABLE 2-4 Message Area (2 of 2)

Window	Description
Messages window	Displays completion messages. This information is also written to the <code>messages</code> file in the <code>/nsr/logs</code> directory.
Pending window	Displays information on why a backup cannot progress, along with requests to mount tapes and other conditions that require intervention.

The `nsrwatch` command line program also provides this information in a character-based format. For more information on the `nsrwatch` [program](#), refer to the man page for this command, or the *Legato Command Reference Guide*.

Sun StorEdge EBS Administrator Toolbar

The Sun StorEdge EBS Administrator program's toolbar enables you to perform common tasks by clicking its buttons.

[TABLE 2-5](#) describes the function of each button.

TABLE 2-5 Sun StorEdge EBS Administrator Toolbar Functions (1 of 2)

Button	Name	Function
	Change Server	Opens the Change Server dialog box to connect to a new Sun StorEdge EBS server.
	Mount	Mounts the device selected in the Devices window.
	Unmount	Unmounts the device selected in the Devices window.
	Label	Labels the device selected in the Devices window.
	Indexes	Opens the Indexes window, which displays information about the online indexes.

TABLE 2-5 Sun StorEdge EBS Administrator Toolbar Functions (2 of 2)

Button	Name	Function
	Volumes	Opens the Volumes window, used to manage volumes.
	Group Control	Opens the Group Control dialog box, used to monitor groups and manually start and stop group backups.

Customizing the Sun StorEdge EBS Administrator Program Interface

You can customize the look and functionality of the Sun StorEdge EBS Administrator program.

Hidden Attributes

By default, the Sun StorEdge EBS Administrator program leaves some attributes hidden from view in the resource windows. Hidden attributes provide additional functionality and control over Sun StorEdge EBS resources.

You can display hidden attributes by selecting Details from the View menu in a particular window.

▼ *To Change the Server Information Refresh Rate (Polling Interval)*

The *polling interval* determines how often (in seconds) data is refreshed in the Sun StorEdge EBS Administrator program. A low number refreshes the data in the status windows more frequently. However, refreshing the data more often can create more network traffic and server overhead. On slower networks, you might want to increase the interval between refreshes by increasing the polling interval value.

To change the polling interval:

1. Select Polling Interval from the Options menu to open the Change Polling dialog box.

2. Using the slide bars, select a interval for each polling interval setting.
[TABLE 2-6 on page 43](#) lists the polling interval options and their functions.

TABLE 2-6 Polling Interval Options

Option	Description
Polling Interval	The amount of time, in seconds, that the Sun StorEdge EBS Administrator waits before polling the Sun StorEdge EBS software for new information and updating its display.
(Minimized) Polling Multiplier	The number by which the Sun StorEdge EBS Administrator multiplies the Polling Interval to determine how often it will poll the Sun StorEdge EBS software when the Sun StorEdge EBS Administrator is minimized. For example, if the Polling Interval is 10 seconds and the Polling Multiplier is 6, then when the Sun StorEdge EBS Administrator is minimized, it will poll the Sun StorEdge EBS software every 60 seconds.

3. Click OK.

Because the Sun StorEdge EBS Administrator does not save the settings in the Change Polling dialog box, two environment variables can be used to set the default polling interval:

- `NWADMIN_POLL` is used to reset the default setting for the Polling Interval option. The value of this environment variable is in seconds, and allowed values are from 5 to 300. If this environment variable is not set, the default value is 10 seconds.
- `NWADMIN_POLL_MULT` is used to reset the default setting for the Polling Multiplier. If this environment variable is not set, the default is value is 6.

Foreign Languages in the Interface

Because Sun StorEdge EBS software uses a client/server model, and is available in several different languages, foreign language messages and other text strings may appear in the windows of the Sun StorEdge EBS Administrator and other programs. This can occur if you are connected to, or are remotely administering, a Sun StorEdge EBS product designed for a different language.

In cases where the foreign language is a double-byte language (for example, Japanese Kanji), the computer must meet these requirements:

- Have appropriate double-byte fonts installed.
- Be configured to use the fonts correctly.

Sun StorEdge EBS Resources

Each resource has a corresponding dialog box that opens by selecting menu commands. Because resource attributes are created and displayed dynamically in the Sun StorEdge EBS Administrator program, the same resource dialog box might show different attributes depending on the Sun StorEdge EBS server to which you are currently connected.

To find out more about each attribute, select On Window from the Help menu.

For more information on Sun StorEdge EBS resources, see [“Sun StorEdge EBS Resources and Attributes” on page 7](#).

Sun StorEdge EBS Character-Based Interface

You can use the Sun StorEdge EBS character-based interface (`nsradmin`) to perform the same configuration and management tasks as in the Sun StorEdge EBS Administrator program.

You can start this interface by entering the following at the command line:

```
# nsradmin
```

For more information about using the character-based interface, refer to the `nsradmin` man page or the *Legato Command Reference Guide*.

Command Line Interfaces

You can perform any Sun StorEdge EBS configuration or administration task (as well as the tasks available in the client-side backup, recover, archive, and retrieve programs) by entering Sun StorEdge EBS commands at the command prompt. For instructions on how to use these commands, refer to the *Legato Command Reference Guide* or the man pages for individual Sun StorEdge EBS commands.

Sun StorEdge EBS Manual Backup, Recovery, and Archiving Interfaces

Sun StorEdge EBS software uses separate graphical interface programs for performing each of the following tasks:

- Manual backups (nwbackup)
- Recoveries (nwrecover)
- Archiving (nwarchive)
- Retrieval of archived data (nwretrieve)

The features and functionality associated with these programs are described in detail in the following chapters:

- [Chapter 7](#)
- [Chapter 8](#)

The nwbackup Program

Use the nwbackup program to perform a manual backup initiated from the client, rather than from the server's scheduled group backup.

To start the nwbackup program, enter the following command at the prompt:

```
# nwbackup
```

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

```
# nwbackup -s server
```

[FIGURE 2-3 on page 46](#) displays the nwbackup program.



FIGURE 2-3 **nwbackup** Program

The nwrecover Program

Use the nwrecover program to administer and monitor Sun StorEdge EBS servers.

To start the nwrecover program, enter the following command at the prompt:

```
# nwrecover
```


To specify the Sun StorEdge EBS server to which you want to connect, use the `-s server` option. To specify the Sun StorEdge EBS client, use enter the `-c client` option:

```
# nwrecover -s server -c client
```

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

[FIGURE 2-4 on page 47](#) displays the `nwrecover` program.



FIGURE 2-4 `nwrecover` Program

The nwarchive Program

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

Note – Use of the `nwarchive` program requires the purchase of a license for the Sun StorEdge EBS Archive application. You must also enter that licensing information into the Sun StorEdge EBS server to fully enable the archiving functionality. See the Sun StorEdge EBS Roadmap for licensing information.

To start the `nwarchive` program, enter the following command at the prompt:

```
# nwarchive
```

To specify the Sun StorEdge EBS server to which you want to connect, enter the `-s server` option with `nwarchive`:

```
# nwarchive -s server
```

[FIGURE 2-5 on page 49](#) displays the `nwarchive` program.




FIGURE 2-5 **nwarchive** Program

For more information on Sun StorEdge EBS archiving, see [Chapter 8](#).

The nwretrieve Program

Use the `nwretrieve` program to retrieve archived files manually from a Sun StorEdge EBS server.

To start the `nwretrieve` program, enter the following command at the prompt:

```
# nwretrieve
```

To specify the Sun StorEdge EBS server to which you want to connect, enter the `-s server` option with `nwretrieve`:

```
# nwretrieve -s server
```

FIGURE 2-6 on page 50 displays the `nwretrieve` program.




FIGURE 2-6 `nwretrieve` Program

Rehosting a Sun StorEdge EBS Server

If you need to transfer the server software to another computer (or transfer the computer to another IP address), the authorization code is invalidated. You must get a *Host Transfer Affidavit* from Sun Customer Service, complete the form, and submit it to Sun. You will receive new authorization codes, which you must install within 15 days, or the Sun StorEdge EBS software becomes disabled.



Caution – Before starting the Sun StorEdge EBS software on the new server, you must remove the server software from the computer you transferred it from, or shut down that computer's `nsrd` daemon. If you start the `nsrd` daemon on another computer with the same enabler or authorization code, you will receive a copy violation error message and the software will become disabled.

Server Management

This chapter describes operations that are managed from the Sun StorEdge Enterprise Backup Software server.

The main sections in this chapter are:

- [“Sun StorEdge EBS Server Connectivity” on page 53](#)
- [“Client/Server Communication Configuration” on page 55](#)
- [“Server Configuration” on page 55](#)
- [“Performance Controls” on page 57](#)
- [“Server Access Management” on page 59](#)
- [“Sun StorEdge EBS User Groups” on page 61](#)
- [“Notifications” on page 65](#)
- [“Reports” on page 70](#)
- [“Firewall Support” on page 75](#)
- [“Storage Nodes” on page 83](#)
- [“Online Indexes” on page 88](#)
- [“Managing the Size of the Online Indexes” on page 96](#)
- [“Server Maintenance Tasks” on page 100](#)

Sun StorEdge EBS Server Connectivity

Sun StorEdge EBS software allows you to connect to and monitor multiple Sun StorEdge EBS servers. The Sun StorEdge EBS Administrator program automatically connects to the local Sun StorEdge EBS server, or to the client’s default backup server if the Sun StorEdge EBS Administrator program is launched on a computer with only the Sun StorEdge EBS client installed. You can change to another Sun

StorEdge EBS server at any time in order to configure that server or monitor its activities by following the instructions for [“To Establish a Server Connection”](#) on page 54.

Note – Unless you are a member of the Administrators Group or a user with the Configure Sun StorEdge EBS privilege, you cannot perform any configuration tasks or issue commands from the Sun StorEdge EBS Administrator program. Only a user with the Monitor Sun StorEdge EBS privilege can monitor the server’s activities. For more information about Sun StorEdge EBS Groups, see [“Sun StorEdge EBS User Groups”](#) on page 61.

▼ To Establish a Server Connection

To change the Sun StorEdge EBS server you are administrating:

4. From the Server menu, select Change to open the Change Server window.
5. From the Server list, select a server. If the server you want is not listed, do one of the following:
 - Click Update Server List to refresh the list.
 - Enter the name of the server in the Server attribute.
6. Click OK.

The Sun StorEdge EBS Administrator program displays information for the newly connected server.

▼ To Close the Connection

To quit the Sun StorEdge EBS Administrator program and terminate the server connection, select Exit from the File menu.

Client / Server Communication Configuration

Communication between the Sun StorEdge EBS server and its clients is defined by configuration attributes entered for the Server and Client resources. How you configure these attributes—such as the level of parallelism or who has administrator privileges—can affect the performance and security of your backups.

For more details about how to set up Sun StorEdge EBS clients in the Client resource, see [Chapter 4](#). To diagnose problems with network communications that affect Sun StorEdge EBS software performance, refer to the troubleshooting information in [“Client/Server Communications” on page 550](#).

Server Configuration

The following section describes how to configure the Sun StorEdge EBS server from the Sun StorEdge EBS Administrator program.

▼ To Set Up the Server

To set up a Sun StorEdge EBS server:

1. At the command prompt, enter the following command to start the Sun StorEdge EBS Administrator program:

```
# nwadmin &
```

2. From the Server menu, select Setup Server to open the Server dialog box.
3. From the View menu, select Details to display hidden attributes.
4. In the Parallelism attribute, enter the maximum number of save streams that the Sun StorEdge EBS server allows to arrive concurrently to the server. For more information on this feature, see [“Parallelism” on page 57](#).
5. For the Manual Saves attribute, select Enabled if you want to allow clients to perform manual backups.

6. For the Public Archives attribute, select Enabled if you want to allow a user to retrieve archived data belonging to another author.
7. For the Update Licenses attribute, select Yes if you want the server to resynchronize its externally managed licenses immediately.
8. In the Administrator attribute, specify all users that should have administrative privileges. For allowed formats, see [“Server Access Management” on page 59](#).

Note – You can also specify administrators by adding users to the Administrators Group. For more information, see [“Sun StorEdge EBS User Groups” on page 61](#).

9. For the Volume Priority attribute, select NearLine Priority.
10. For the Volume Priority attribute, select NearLine Priority.
11. Complete the Company Information attributes in order to complete the registration form and obtain an authorization code.

The product serial number is located on the Enabler Certificate you received when you bought your Sun StorEdge EBS software. Make sure you enter only the product serial number of the server, not that of any add-on product.
12. Click Apply.

▼ To Determine the Current User, Privileges, and User Group Membership

The current user is listed on the main window of the Sun StorEdge EBS Administrator program, in the server information area. It is also listed in the Server resource, along with the current user’s privileges and User Group membership.

To determine the user identification of the current user, as well as the user’s privileges and User Group membership:

1. In the Sun StorEdge EBS Administrator program, select Setup Server from the Server menu.
2. In the Server dialog box, select Details from the View menu to display hidden attributes.
3. Scroll to the bottom of the dialog box.

The following information is displayed:

- The User ID attribute identifies the current user.
- The Privileges attribute identifies the privileges that the current user has.
- The Member Of attribute identifies the User Groups to which the current user belongs.

For more information about User Groups and their associated privileges, see [“Sun StorEdge EBS User Groups” on page 61](#).

Performance Controls

You can increase backup speed by setting the Sun StorEdge EBS server to *multiplex* data on a storage device. The Parallelism and Target Sessions attributes enable data from more than one save set to be written to a single storage volume by using the devices to their maximum capabilities.

Note – By definition, save sets that are multiplexed must belong to the same pool of storage volumes.

Multiplexing optimizes and distributes the flow of data from multiple clients to all of the storage devices available to the Sun StorEdge EBS server.

Parallelism

The Sun StorEdge EBS *parallelism* feature enables the Sun StorEdge EBS server to back up data from many clients concurrently, and multiplex the data onto the same backup volume. The Parallelism attribute is located in the Server dialog box. For more information on server configuration and the Set Up Server dialog box, see [“To Set Up the Server” on page 55](#).

Parallelism 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 the 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 as described in [TABLE 3-1 on page 58](#). To learn more about storage nodes, see [“Storage Nodes” on page 83](#).

[TABLE 3-1](#) displays the maximum parallelism levels for this release of the Sun StorEdge EBS server product, including the maximum for each enabled storage node.

TABLE 3-1 Maximum Parallelism Levels

Sun StorEdge EBS Product	With Each Enabled Storage Node	Without Storage Nodes
Workgroup and Business Editions	Not applicable	Parallelism: 32
	Not applicable	Devices: 4
Network Edition	Parallelism: 32 Maximum = (32 x number of nodes) + 32	Parallelism: 32
	Devices: 16 Maximum = (16 x number of nodes) + 16	Devices: 16
Power Edition	Parallelism: 32 Maximum = (32 x number of nodes) + 64	Parallelism: 64
	Devices: 16 Maximum = (16 x number of nodes) + 32	Devices: 32

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 this 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.

To disable parallelism entirely, so that only one client can back up to the server at a time, set the parallelism value to 1.

Target Sessions

The Sun StorEdge EBS *target sessions* feature allows you to set the number of backup sessions accepted by an active device. The Target Sessions attribute is configured in the Device resource. 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 [“Devices” on page 187](#).

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.

Server Access Management

Many Sun StorEdge EBS server functions require the operator to have administrator privileges. All administrators that are authorized to configure the server must have their addresses included in the server administrator list. For more information, see [“Administrator Privileges” on page 59](#).

You must be a member of the Sun StorEdge EBS Administrator Group to add, modify, or delete Sun StorEdge EBS administrators. For more information about the Administrator Group, including how to add and delete members from this group, see [“Sun StorEdge EBS User Groups” on page 61](#).

Administrator Privileges

Although any user with the Monitor Sun StorEdge EBS privilege can view the server’s resources from a client computer, the following additional privileges are required to administer the server:

- To change security settings, the user must be a member of the Sun StorEdge EBS Administrator Group.
- To perform a scheduled backup of SYSTEM and VSS SYSTEM save sets, the user must be a member of the Sun StorEdge EBS Administrator Group.
- To add to or change the configuration of the server, media devices, and autochangers, the user must have the Configure Sun StorEdge EBS privilege.

The Administrator attribute in the Server, which lists all members Administrator Group, must include the following entries, depending on the platform of the client trying to administer the Sun StorEdge EBS server:

- Microsoft Windows: user=Administrators, 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 users or user groups can be granted administrative privileges by adding their names to the Administrator attribute of the Server resource. User names should be listed in the form of name/value pairs, using the syntax:

name=value[,name=value, ...]

Where *name* can be one of the following:

- user
- group
- host
- domain
- domain_sid
- domaintype (either NIS or WINDOMAIN)

For example, to grant administrative privileges to a user named `jdoe` on the host named `jupiter`, enter the following for the Administrator 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 only a username, which allows that user to administer Sun StorEdge EBS software 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 to not compromise your enterprise security. Netgroup names can also be entered and must be preceded by an ampersand (&).

Example: Using the Administrator Attribute

The following example grants Sun StorEdge EBS administrative privileges to `root` from any host, to the user `operator` from the hosts *mars* and *jupiter*, and 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
```

Note – If you use the `nsradmin` interface to input these entries, they must be separated by commas, and multiple name=value pairs must be surrounded by quotation marks:

```
administrator: user=root, "user=operator,host=jupiter", "user=
operator,host=mars", &netadmins
```

Sun StorEdge EBS User Groups

Sun StorEdge EBS software includes an access control feature that is configured through the User Group resource. This feature allows Sun StorEdge EBS administrators to assign users to one of two Sun StorEdge EBS *user groups*. Each user group has a specific set of privileges associated with it, as defined in the Privileges attribute. Privileges include such things as permission to monitor Sun StorEdge EBS software, and to backup and recover local data.

Preconfigured User Groups

Sun StorEdge EBS software comes with two preconfigured user groups:

- Administrators
- Users

Administrators

Members of the Administrators group have permission to perform all Sun StorEdge EBS functionality. The `root@server` user is always a member of this group and cannot be removed from the group.

Note – 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 Associated with Sun StorEdge EBS Users

The privileges associated with the Users group can be customized to fit requirements for your Sun StorEdge EBS users. Privileges associated with the Administrators group cannot be edited.

[TABLE 3-2](#) lists these preconfigured privileges and their associated permissions.

TABLE 3-2 User Group Privileges (*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.

TABLE 3-2 User Group Privileges (2 of 3)

Privilege	Permissions
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.
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	Allows group members to perform the following: • Monitor Sun StorEdge EBS 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 backup 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 3-2 User Group Privileges (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 operating-system based user permissions. • 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 operating-system based user permissions. • User with the privilege still must be logged in as root (UNIX) or Administrator (Microsoft Windows) to run the <code>savegrp</code> command or perform NDMP backups. • To allow scheduled backups to operate correctly, the client's root user (UNIX) or Administrator (Microsoft Windows) has this privilege automatically.

▼ To Edit a User Group

Note – Only users who belong to the Administrators group, or users with the Change Security Settings privilege enabled, can edit a User Groups resource.

1. In the Sun StorEdge EBS Administrator program, select User Groups from the Customize menu.
2. Select the appropriate user group to be edited.
3. Edit the attributes for the group.
 - For the Administrators group, you can change the Comment attribute or the Users attribute; you cannot change the Privileges attribute.
 - For the Users group, all attributes can be changed.
4. Click Apply.

Notifications

A [notification](#) provides information about events that occur in a Sun StorEdge EBS environment. The information may be sent to an administrator via e-mail, logged to a file, or acted on by other operations, such as sending an SNMP message to an SNMP monitoring agent.

Notifications allow you to configure the events that you want 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 `/nsr/logs` directory. Depending on the notification being sent, the server might also display notifications in a Sun StorEdge EBS Administrator program status window.

Preconfigured Notifications

By default, Sun StorEdge EBS software is preconfigured to provide most of the notifications required to monitor Sun StorEdge EBS events. [TABLE 3-3](#) lists these preconfigured notifications and the associated actions performed by the Sun StorEdge EBS server.

TABLE 3-3 Preconfigured Notifications (1 of 2)

Notification	Action Performed by the Sun StorEdge EBS Server
Bootstrap	Sends the bootstrap information (important disaster recovery information) to the Sun StorEdge EBS server's default printer.
Cleaning cartridge expired	Sends e-mail to root: replace the expired cleaning cartridge.
Cleaning cartridge required	Sends e-mail to root: insert a cleaning cartridge.
Device cleaned	Sends e-mail to root: device cleaning is complete.
Device cleaning required	Sends e-mail to root: clean the device.
Device disabled	Sends e-mail to root: a device has been automatically disabled.
Filesystem full - recovering adv_file space	Launches the <code>nsrim</code> program to remove aborted and expired save sets. Used with advanced file type devices only.
Filesystem full - waiting for adv_file space	Sends e-mail to root: advanced file volume is full.

TABLE 3-3 Preconfigured Notifications (2 of 2)

Notification	Action Performed by the Sun StorEdge EBS Server
Index size	Sends e-mail to root: check the size of the client file index because it will soon exceed the space available.
Log default	Directs the UNIX syslog facility (/usr/bin/logger) to log information from the /nsr/logs/messages file into the appropriate system log file.
Migration attention	Sends e-mail to root: check the status of a migration operation.
Migration completion	Sends e-mail to root: a migration operation is complete.
Registration	Sends e-mail to root: check the registration status.
Savegroup completion	Sends e-mail 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 Sun StorEdge EBS SNMP module has been purchased and enabled).
Tape mount request 1	Requests media be mounted in a device and displays a pending message.
Tape mount request 2	Requests media be mounted in a device and displays a pending message.
Tape mount request 3	Sends e-mail to root requesting that the tape be mounted.

Customizing Notifications

Notifications require three elements:

- Events
- Actions
- Priorities

These elements are described in the following sections.

Events

TABLE 3-4 lists the type of event notifications that trigger an action.

TABLE 3-4 Events

Event	Description
Media	The media needs operator attention to mount or unmount backup volumes.
Savegroup	A backup group has completed backing up.
Index	The index needs attention.
Server	Other server events have occurred (for example, restarting the Sun StorEdge EBS server).
Registration	Product registration needs attention.
Device cleaned	The device has been cleaned.
Device cleaning required	The device requires cleaning.
Cleaning cartridge required	Mount the cleaning cartridge.
Cleaning cartridge expired	The cleaning cartridge has expired and needs replacing.

Note – You cannot edit the events for an existing notification. To change, you must delete the notification and re-create it with different settings.

Actions

TABLE 3-5 lists the actions that the server takes after an event notification occurs.

TABLE 3-5 Actions

Action	Description
logger	Uses the UNIX syslog facility (/usr/bin/logger) to log information or send messages.
lp	Prints the notification.
mail	Sends e-mail 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, you can use a third-party e-mail program rather than the UNIX mail command.

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 [“Sun StorEdge EBS User Groups” on page 61](#).

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; and the priority of the third request is Alert.

[TABLE 3-6](#) lists the priorities upon which notifications are based.

TABLE 3-6 Priorities

Priority	Description
Info	Supplies information about the server's status.
Notice	Presents important information.
Warning	Presents information about a nonfatal error.
Waiting	Indicates that the server is waiting for an operator to perform a routine task, such as mounting a backup volume.
Critical	Indicates that the server detected an error that should be fixed.
Alert	Indicates a severe condition that requires immediate attention.
Emergency	Indicates that a condition is imminent that might cause the Sun StorEdge EBS software to fail, unless corrected immediately.

Note – You cannot edit the priorities for an existing notification. To do so, you must delete the notification and re-create it with different settings.

▼ To Create a Notification

Only users who belong to the Administrators group, or users with the Change Security Settings privilege, can create notifications. For more information about privileges, see [“Sun StorEdge EBS User Groups” on page 61](#).

To create a notification:

1. In the Sun StorEdge EBS Administrator program, select Notifications from the Customize menu to open the Notifications window.
2. From the View menu, select Details.
3. Click Create.
4. In the Name attribute, enter a name for the notification.
5. Optional: In the Comment attribute, enter a description of the notification.
6. In the Action attribute, enter the name of the command you want the Sun StorEdge EBS server to use.

For example, to use the `nsrtrap` program to send a message to a network management console about a volume that needs mounting, enter:

```
/usr/sbin/nsr/nsrtrap -s 1 host
```

7. Complete the other attributes as necessary, then click Apply.

▼ To Edit a Notification

You cannot edit the name, events, or priorities for an existing notification. To change these, you must delete the notification and re-create it with different settings.

To edit a notification:

1. In the Sun StorEdge EBS Administrator program, select Customize from the Notifications menu.
2. From the list displayed in the Notifications attribute, select the notification to edit.
3. Edit the Action attribute, then click Apply.

Only users who belong to the Administrators group, or users with the Change Security Settings privilege, can edit the Action attribute. For more information about privileges, see [“Sun StorEdge EBS User Groups” on page 61](#).

▼ To Delete a Notification

Only users who belong to the Administrators group, or users with the Configure Sun StorEdge EBS privilege, can delete notifications. For more information about privileges, see [“Sun StorEdge EBS User Groups” on page 61](#).

You cannot delete any of the preconfigured Sun StorEdge EBS notifications.

To delete a notification:

1. In the Sun StorEdge EBS Administrator program, select Customize from the Notifications menu.
 2. From the list displayed in the Notifications attribute, select the notification to delete.
 3. Click Delete.
 4. Click OK to confirm the deletion.
 5. Click Apply.
-

Reports

The following sections describe various aspects of the Sun StorEdge EBS reporting function.

Routine Data Movement Operations Reports

The degree of success in the completion of scheduled group backups, group cloning, and archive operations is reported to you by the `savegrp` program through a Savegroup Completion Report. This report is the program triggered by the preconfigured notification Savegroup Completion. The report is sent as an e-mail to root and to the log file in `/nsr/logs/messages`.

The report consolidates the following information:

- Success or failure of each save set participating in the operation
- Operation's save date and time
- Bootstrap save set ID
- Bootstrap volume location (volume name, starting record number, and ending record number)

Sun StorEdge EBS Server Bootstrap Report

When the Sun StorEdge EBS server finishes a scheduled backup that includes the Sun StorEdge EBS server, it creates a *bootstrap* save set, which contains backups of the client file index, media database, and Sun StorEdge EBS configuration files. The

data in the bootstrap save set is required to re-create the Sun StorEdge EBS server in the event of a disaster. The bootstrap file is written to media and automatically printed to the default printer. This printed report is the action triggered by the preconfigured bootstrap notification.



Caution – Save the bootstrap report in a safe location. The bootstrap information is essential for recovery from a disaster. It provides a fast way to locate the starting point of a disaster recovery in the event of a server failure. For more information, refer to the *Legato NetWorker Disaster Recovery Guide*.

Storage Management Application Reports

[TABLE 3-7](#) lists the programs that the Sun StorEdge EBS software provides to query the contents of the storage management system. For more detailed descriptions of the most commonly used commands and options, refer to the *Legato Command Reference Guide*.

TABLE 3-7 Storage Management Report Programs

Program	Information Reported
mminfo	Contents and mode of the storage volumes and/or the identification numbers and status of the stored save sets.
mmlocate	User-defined location of storage volumes.
nsrinfo	Contents of the client file index.
nsrmm	Status of the storage devices known to the Sun StorEdge EBS software.

Generating Program-Readable Reports

The Sun StorEdge EBS software includes the ability to generate reports in formats that are easily parsed by scripts and computer programs. These program-readable formats can be generated from the `mminfo` and `nsrinfo` commands using the `-x` option. Two formats are available for output:

- Delimiter-separated values format.
- XML format

Delimiter-Separated Values Format

The delimiter-separated values format creates reports that can be easily parsed by separating each value with a delimiter character. You can select any delimiter character to separate the values. To generate a report with `mminfo` or `nsrinfo` by using a delimiter character, include the `-x` option, followed by `c`, followed by the appropriate delimiter character. For example:

```
% mminfo -xc, -s jupiter -c mars
```

This creates the following example program-readable report, using a comma to delimit the values:

```
volume,client,date-time,sum-size,level,name  
jupiter.012,mars,03/01/02,40 MB,full,/export/home  
jupiter.012,mars,03/01/02,51 MB,full,  
[...]
```

In the preceding example output, the heading (the first row) represents the name of each field, separated by the delimiter character. Scripts or programs can parse the first row to determine the names of each field, then parse the subsequent rows to determine the output values. This way, the same script can be used to parse the report, even when the fields included in the report change based on different command line options, or because of changes to the command output between releases of the Sun StorEdge EBS software.

XML Format

The XML format creates reports that use XML markup tags.

To generate a report in XML format with `mminfo` or `nsrinfo`, include the `-x` option, followed by `m`. For example:

```
% mminfo -xm -s jupiter -c mars
```

This creates the following example XML-formatted report:

```
<?xml version="1.0" encoding="UTF-8"?>
  <!DOCTYPE mminfo-tabular-report [
    <!ELEMENT result (field*)>
    <!ELEMENT field (annotation | barcode | browse-time |
    [...])>
    <!ELEMENT annotation (string-field)>
    <!ELEMENT barcode (string-field)>
    <!ELEMENT browse-time (time-field)>
    [...]
    <!ELEMENT boolean-field (#PCDATA)>
    <!ELEMENT bytes-field (#PCDATA)>
    <!ELEMENT boolean-field (#PCDATA)>
    <!ELEMENT id-field (#PCDATA)>
    <!ELEMENT kilobytes-field (#PCDATA)>
    <!ELEMENT level-field (#PCDATA)>
    <!ELEMENT number-field (#PCDATA)>
    <!ELEMENT string-field (#PCDATA)>
    <!ELEMENT time-field (#PCDATA)>
    [...]
  ]>

<mminfo-tabular-report>
<result>
<volume>jupiter.012</volume>
<client>mars</client>
<date-time>03/01/02</date-time>
<sum-size>40 MB</sum-size>
<level>full</level>
<name>/export/home</name>
</result>
<result>
<volume>jupiter.012</volume>
<client>mars</client>
<date-time>03/01/02</date-time>
<sum-size>51 MB</sum-size>
<level>full</level>
<name>/</name>
</result>
[...]
```

The opening section contains the document type definition (DTD), which defines each element that may be included in the report. The DTD section and actual report are separated by a blank line, which can be used by a script to determine when the DTD section ends and the report begins.

Sun StorEdge EBS Server Statistics and Diagnostic Reports

Messages that report on Sun StorEdge EBS diagnostics are displayed in the Sun StorEdge EBS administrator interface and are also contained in the Sun StorEdge EBS messages file (`/nsr/logs/messages`). These messages include warning and error conditions and notice of lost connections.

Message Log Files

The messages generated by the Sun StorEdge EBS server daemons (`nsrd`, `nsrindexd`, `nsrmmdbd`, and `nsrmmmd`) are contained in the Sun StorEdge EBS messages log and the `daemon.log` file, typically found in the `/nsr/logs` directory.

Monitoring and Reporting Changes to Sun StorEdge EBS Resources

The Monitor RAP (resource allocation protocol) attribute in the Server resource tracks the history of additions, deletions, or modifications to Sun StorEdge EBS resources and their attributes. These changes are recorded in the `rap.log` file, located in the `/nsr/logs` directory. The `rap.log` file lists the user name, the source computer, and the time of the modification. Sufficient information is logged in the `rap.log` file to enable an administrator to undo any changes.

Note – By default, the Monitor RAP option is disabled.

▼ To Enable the Monitor RAP Attribute

To enable the Monitor RAP attribute:

1. In the Sun StorEdge EBS Administrator program, select Server Setup from the Server menu.
2. From the View menu, select Details to display hidden attributes.
3. Set the Monitor RAP option to Enabled.
4. Click Apply.

Firewall Support

Firewall support enables you to back up 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 and accepts or rejects it based on user-defined rules.

- A *service port* is used by a Sun StorEdge EBS server or storage node to listen for client requests. The default range used for service ports is 7937 to 9936.
- A *connection port* is used by the Sun StorEdge EBS processes to connect to the server during backup and recovery sessions. The default range used for connection ports is 10001 to 30000.

This section explains how Sun StorEdge EBS software uses firewall ports, and provides guidelines on configuring these ports.

Configuring Firewall Support

This section provides guidelines on configuring the service and connection ports used by the Sun StorEdge EBS software. This section also provides details on the range of ports required by a server, client, and storage node.

To configure firewall support, do one of the following:

- Use the default settings. For more information, see [“To Configure Firewall Support by Using Default Settings” on page 76](#).
- Restrict the number of ports the Sun StorEdge EBS software uses. For more information, see [“Restricting the Port Range” on page 76](#).

If the configured service or connection port range is not large enough, the associated services and processes will not be able to communicate through the firewall. For more information, see [“Troubleshooting Port Configuration Problems” on page 83](#).

▼ To Configure Firewall Support by Using Default Settings

To configure the ports on the firewall by using the default settings:

1. Set the firewall rules to allow inbound and outbound TCP/UDP packets to the following ports:
 - For service ports, use the range of 7937-9936
 - For connection ports, use the range of 10001-30000



Caution – If the client computer or storage node is running Sun StorEdge EBS client software earlier than release 5.5, set the firewall rules to allow inbound and outbound packets from UDP 111 and 600-1023.

2. Ensure Network Address Translation (NAT) is disabled for each Sun StorEdge EBS host. To disable NAT, refer to the documentation accompanying the router.
3. Shut down and restart the Sun StorEdge EBS server.

Restricting the Port Range

Complete these three procedures when restricting the port range:

1. Determine the minimum port range required by each server, client and storage node. For more information, see [“To Determine Minimum Port Range Required” on page 76](#).
2. Change the default settings on the Sun StorEdge EBS server. For more information, see [“To Change Sun StorEdge EBS Default Settings” on page 80](#).
3. Complete the configuration of the firewall feature. For more information, see [“To Complete the Configuration of Firewall Support” on page 81](#).

▼ To Determine Minimum Port Range Required

To determine the minimum service and connection port range required by each Sun StorEdge EBS server, client, and storage node, see:

- [“Server Port Range Requirements” on page 77](#)
- [“Client Port Range Requirements” on page 78](#)
- [“Storage Node Port Range Requirements” on page 79](#)

To determine how the ports are being used on a specific computer, use the `netstat -a` or `rpcinfo -p` command.

Server Port Range Requirements

This section provides guidelines for determining the minimum number of service and connection ports required by a Sun StorEdge EBS server.

■ Service Ports

The minimum number of service ports that can be assigned to a Sun StorEdge EBS server is $10+4n$, where:

- 10 is the minimum number of services registered on the Sun StorEdge EBS server.

If the number of services registered on the server is greater than 10, `nsrexecd` will use one port and the additional 9 ports can be shared among the remaining services.

- n is the number of devices configured on the server.

Example: Setting the Server's Service Port Range

If your configuration included 4 tape devices and 15 registered services, you would need to allocate 31 service ports to the Sun StorEdge EBS server. In this case, you could set the service port range for the Sun StorEdge EBS server at 8001-8031.

■ Connection Ports

Typically, a server requires a minimum of 600 connection ports. The minimum number of connection ports that can be assigned to a Sun StorEdge EBS server is approximately the sum of the following:

- Number of:
 - a. Local and remote devices x 10.
 - b. Autochangers x 5.
 - c. Groups (or instances of `savegrp`), backing up to the server through the firewall x 5. For each `savegrp` instance that runs `nsrim`, an additional 3 ports will be required.
 - d. Sun StorEdge EBS interfaces to be run remotely (such as `nwadmin`, `nwbackup`, and `nwrecover`) + 6.
 - e. Noncritical messages being posted. The Sun StorEdge EBS server uses UDP ports for posting noncritical events. This number is dependent on the debug level and the number of errors. The larger the debug level, the more messages that are posted. To turn debugging off, set the debug level to zero.

- f. Number of mechanisms used for copy protection. The Sun StorEdge EBS Server uses UDP ports for copy protection. This number is dependent on the copy protection rules specified on the network.
- Parallelism value set in:
 - a. Server resource x 15.
 - b. Client or Server resource used during recovery sessions + 9.
 - c. Client or Server resource used during recovery sessions with a storage node + 9.
- Parallelism value assigned to a device x 6, if you are cloning from a storage node to a server.

Example: Setting the Server's Connection Port Range

If the server required 600 connection ports, you could set the connection port range for the Sun StorEdge EBS server at 10001-10600.

Based on this server-related example, you would set the firewall rules to allow inbound and outbound TCP/UDP packets from the following ports:

- Sun StorEdge EBS server service ports: 8001-8014
- Sun StorEdge EBS server connection ports: 10001-10600
- nsrexecd: 7937
- Sun portmapper: 7938

Client Port Range Requirements

This section provides guidelines for determining the minimum number of service and connection ports required by a Sun StorEdge EBS client.

- Service Ports

The Sun StorEdge EBS client services do not bind to service ports as specified by the service port range. You do not need to specify a service port range for a client.

The Sun StorEdge EBS client services, nsrexecd and Sun portmapper, will bind locally to ports 7937 and 7938, respectively.

- Connection Ports

Typically, a Sun StorEdge EBS client requires a minimum of 100 connection ports. The minimum number of connection ports that you can assign to a client is as follows:

- Parallelism value specified in the Client resource x 10
- Number of concurrent operations, such as `nwadmin`, `nwrecover`, and `nwbackup` x 10

Example: Setting the Client's Connection Port Range

If the client required 100 connection ports, you could set the connection port range for the Sun StorEdge EBS client at 10001-10100.

Based on this example, you would set the firewall rules to allow inbound and outbound TCP/UDP packets from the following ports:

- Sun StorEdge EBS client connection ports: 10001-10100
- `nsrexecd`: 7937
- Sun portmapper: 7938

Storage Node Port Range Requirements

This section provides guidelines for determining the minimum number of service and connection ports required by a Sun StorEdge EBS storage node.

■ Service Ports

The minimum number of service ports required by a storage node is equal to the number of tape devices attached to the storage node.

Example: Setting the Storage Node's Service Port Range

If your configuration included 4 tape devices, then you would need to allocate 4 service ports to the Sun StorEdge EBS storage node. In this case, you could set the service port range for the Sun StorEdge EBS storage node at 8001-8004.

■ Connection Ports

Typically, a storage node requires a minimum of 150 connection ports. The minimum number of connection ports that you can assign to a storage node is approximately the sum of the following:

- Number of local devices x 10

- Number of autochangers x 5
- If cloning from a server to a storage node, multiply the device parallelism x 6

Example: Setting the Storage Node's Connection Port Range

If the storage node required 150 connection ports, you could set the connection port range for the Sun StorEdge EBS storage node at 10001-10150.

Based on this example, you would set the firewall rules to allow inbound and outbound TCP/UDP packets from the following ports:

- Sun StorEdge EBS storage node service ports: 8001-8004
- Sun StorEdge EBS storage node connection ports: 10001-10150
- nsrexecd: 7937
- Sun portmapper: 7938

▼ To Change Sun StorEdge EBS Default Settings

You can configure the service and connection ports for the server, client, and storage nodes by changing the assigned default settings through the Sun StorEdge EBS Administrator program or the `nsrports` program.

- If the Sun StorEdge EBS Administrator program is used to configure the service and connection port ranges:
 1. From the Options menu, select Configure Ports.
 2. In the Configure Ports dialog box, enter the hostname or IP address of the computer to which the port configuration applies and click OK.
 3. In the Set Ports dialog box, enter a service ports range and a connection ports range.
 4. Click OK to save your changes.
 5. Click Cancel to exit the Configure Ports dialog box.
- If the `nsrports` program is used to configure the connection port and service port ranges, enter the following at the command prompt:

```
# nsrports -s server -a auth_server [-S|-C] range
```

[TABLE 3-8 on page 81](#) describes the options used with the `nsrports` command.

TABLE 3-8 Options for the nsrports Command

Option	Description
-S <i>server</i>	Specifies the system to contact.
-a <i>auth_server</i>	Connects to a remote system that is located on a platform other than the one on which the command is being executed.
-S	Sets the system's service ports range to the specified range.
-C	Sets the system's connection ports range to the specified range.

▼ To Complete the Configuration of Firewall Support

To complete this configuration process:

1. Set the firewall rules to allow inbound and outbound TCP/UDP packets within the port range set on the Sun StorEdge EBS server.
2. Set the firewall rules to allow inbound and outbound packets from ports 7937 and 7938.

These ports are used respectively by the Sun StorEdge EBS services nsrexecd and Sun portmapper.



Caution – If the release of Sun StorEdge EBS client software on the client computer or storage node is earlier than Sun StorEdge EBS release 5.5, the firewall rules must be set to allow inbound and outbound packets from UDP 111 and 600-1023.

3. Ensure NAT is disabled for each Sun StorEdge EBS host.

To disable NAT, refer to the documentation accompanying the router.

4. Shut down and restart the Sun StorEdge EBS server.

Sample rpcinfo Output

The `rpcinfo -p` command can be used to determine port allocations. The sample `rpcinfo` output shown in [TABLE 3-9](#) shows port allocations for five devices in addition to the ten services required by the Sun StorEdge EBS server. A single port,

7937, is reserved for the **nsrexecd** program. Three ports, 7938 to 7940, are shared among nine other services. Five ports, 7941 to 7945, are allocated for the five **nsrmmd** services. In total, nine ports are allocated for 15 services.

TABLE 3-9 Sample of rpcinfo Output

Program	Version Number	Protocol	Port	Program/ Daemon ^a
390113	1	tcp	7937	nsrexec
390103	2	tcp	7938	nsrd
390109	2	tcp	7938	nsrstat
390110	1	tcp	7938	nsrjb
390103	2	udp	7939	nsrd
390109	2	udp	7939	nsrstat
390110	1	udp	7939	nsrjb
390107	4	tcp	7940	nsrmdbd
390107	5	tcp	7940	nsrmdbd
390105	5	tcp	7940	nsrindexd
390104	105	tcp	7941	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 daemon name. For example, **nsrexecd** registers its daemon as the **nsrexec** daemon rather than the **nsrexecd** daemon.

Troubleshooting Port Configuration Problems

This section provides guidelines for troubleshooting common error messages that may 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:

1. Increase the service port range configured for the specified host name.
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 configured range on system *hostname*

The configured connection port range is not large enough for the Sun StorEdge EBS backup and recovery sessions to run. The current backup or recovery session has been aborted.

To resolve this problem:

1. Increase the connection port range configured for the specified host name.
2. Set the firewall rules to allow inbound and outbound TCP/UDP packets within the port range specified in [Step 1](#).

Storage Nodes

Storage nodes are host computers with attached *remote storage devices* (devices not physically attached to the controlling Sun StorEdge EBS server). 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. By using the Sun StorEdge EBS software, you can route client data 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, you can perform typical storage tasks, such as:

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

Note – Only users that have the Configure Sun StorEdge EBS privilege can add to or change the configuration of the Sun StorEdge EBS server, media devices, and autochangers. For more information, see [“Sun StorEdge EBS User Groups” on page 61](#).

The following sections contain information about storage nodes:

- [“Prerequisites” on page 84](#)
- [“Licensing” on page 84](#)
- [“Storage Node Configuration” on page 85](#)
- [“Timeout Attribute for Storage Node Operations” on page 85](#)
- [“Storage Node Affinity” on page 86](#)
- [“Storage Nodes and Cloning” on page 87](#)
- [“Storage Node Troubleshooting Tips” on page 87](#)

Prerequisites

To operate the Sun StorEdge EBS software with storage nodes, the following software must be installed on the storage node. Packages must be installed in the order listed here:

- Sun StorEdge EBS client software
- Device driver software for each device
- Sun StorEdge EBS storage node software

Licensing

You must 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 autochanger or silo connected to each storage node (included in the autochanger and silo Enabler Certificates).



Caution – If you configure and use the Sun StorEdge EBS storage node software on a remote computer in evaluation mode, and the evaluation period expires before you install the storage node enabler code, you will not be able to recover backed-up data to the storage node. To recover that data, you must move the remote volume to a local device and perform the recovery from there.

Storage Node Configuration

After the required software and enabler codes are installed, you must configure each storage node device. This procedure is similar to the tasks required for other devices, as described in [“Devices” on page 187](#). However, the storage node device name must follow this convention:

```
rd=storage_node_hostname:device_name
```

For example, `rd=jupiter.mars.com:/dev/rmt/0hbn`

For more details on configuring storage node devices, including those in autochangers and silos, see [“Storage Nodes” on page 228](#).

To run the `jbconfig` program, the root user must have the Configure Sun StorEdge EBS privilege. To run the `scanner` program, the user must have the Operate Sun StorEdge EBS privilege. For more information about Sun StorEdge EBS privileges and User Groups, see [“Sun StorEdge EBS User Groups” on page 61](#).



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

Timeout Attribute for Storage Node Operations

By specifying an attribute named `Nsrmmmd Control Timeout`, which is configured in the Server resource, you can configure the amount of time that a Sun StorEdge EBS server waits for a storage node request to be completed. You must have hidden attributes enabled in order to view this attribute.

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

The default value assigned to Nsrmmmd Control Timeout is 5 minutes; you can specify any value (in whole minutes) within a range of 2 to 30 minutes.

In previous releases of Sun StorEdge EBS software, this function was controlled through the NSR_MMDCONTROL environment variable. If nsrd initializes on the server and detects that a setting for NSR_MMDCONTROL exists, the following informational message is issued:

NSR_MMDCONTROL env variable is being ignored

use 'nsrmmmd control timeout' attribute instead

If you receive this message, do the following:

1. Shut down the Sun StorEdge EBS daemons.
2. Remove the environment setting for NSR_MMDCONTROL.
3. Restart the Sun StorEdge EBS daemons.
4. If necessary, start the Sun StorEdge EBS Administrator program and adjust the value of Nsrmmmd Control Timeout to the setting that was previously assigned to the NSR_MMDCONTROL variable, or to a setting that best meets your current requirements.

Storage Node Affinity

You can determine which Sun StorEdge EBS servers and storage nodes receive a client's data—known as *storage node affinity*—by entering their hostnames in the Storage Nodes attribute located in the Client dialog box. 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 you want the client to back up to that storage node, you must enter the name of the storage node in the client's Storage Nodes attribute *above* the default nsrserverhost. You can add storage node names 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. If you use `mmrecov` to recover a bootstrap save set, you must recover the data from a local device.

Storage Nodes and Cloning

Storage nodes can clone data in a manner similar to Sun StorEdge EBS servers. A hierarchical structure exists between the storage node and the Sun StorEdge EBS server. This determines where storage node data is routed to be cloned. The following list presents this hierarchy, from highest to lowest:

1. The storage node's Client resource contains the Clone Storage Nodes attribute, used to identify other storage nodes that this storage node can use for cloning its data. Enter the name of each storage node you want to be available. The storage node directs its data to the first storage node in the list that has an enabled device and is capable of receiving the data for cloning.

Unlike the Storage Nodes attribute, there is no default setting for the Clone Storage Nodes attribute. If there are no storage nodes listed, the Sun StorEdge EBS server's settings are used, as described in the steps that follow.

Note – The Cloning Storage Nodes attribute is only applicable for storage nodes. It is not used for clients that do not act as storage nodes.

2. If the Clone Storage Nodes attribute in the storage node's Client resource is not set, the Clone Storage Nodes attribute in the server's Client resource is used.
3. If the Clone Storage Node attribute in the server's Client resource is not set, the Storage Node attribute in the server's Client resource is used.

Storage Node Troubleshooting Tips

If a backup fails, and the following message is displayed:

no matching devices; check storage nodes, devices or pools

the problem may be related to storage node affinity.

Possible reasons for this include:

- There are no enabled devices on the storage nodes listed in the Storage Nodes attribute.
- The devices do not have volumes that match the pool required by the backup request.
- All devices are set to read-only.

For example, when the client has only one storage node in its Storage Node list, and all devices on that storage node are disabled, you must fix the problem and 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.
- Add another storage node to the list that has enabled devices meeting the pool restrictions.
- Set one of the devices to read / write.

Online Indexes

The Sun StorEdge EBS server tracks the files it backs up in two databases, which are stored on the server's local filesystem:

- 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
 - Filesystems in each save set

Unlike the client file indexes, there is only one media database per server.

The Sun StorEdge EBS server can automatically control the size of the client file indexes and media database according to the browse policies and retention policies you set. For more details about using browse and retention policies, see [“Managing the Data Life Cycle” on page 127](#).

The following sections provide information about managing online indexes:

- [“Characteristics of the Online Indexes” on page 89](#)
- [“Automated Index Activities” on page 90](#)
- [“Checking Online Indexes” on page 90](#)

- [“To View Information about the Indexes” on page 90](#)
- [“Index Save Sets” on page 92](#)
- [“Save Set Details” on page 93](#)
- [“Location of the Client File Index” on page 94](#)

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.



Caution – If the filesystem 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 filesystem.

Sun StorEdge EBS software uses browse and retention policies to manage the life cycle of data, and to automatically control the size of the client file index. For more information on policies, see [“Browse and Retention Policies” on page 120](#).

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 filesystem 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.

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” on page 312](#).

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, refer to the *Legato Command Reference Guide* or the man pages for these commands.

▼ To View Information about the Indexes

Information about the client file indexes can be used to:

- Determine the size of the indexes.
- Reclaim space used by the indexes.
- Determine the save sets associated with a client.

To view information about the indexes, select Index from the Clients menu to open the Indexes dialog box, which displays index information for all clients of the server. Index information for all clients of the server appears as shown in

[FIGURE 3-1 on page 91](#).




FIGURE 3-1 Indexes Dialog Box

TABLE 3-10 identifies the information displayed for each client in the Clients window of the Indexes dialog box.

TABLE 3-10 Clients Window Description

Column	Description
Client	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.
Used	Percentage of the index file's allocated space that is currently in use. • If there is no data in the index, 0% is displayed. • If there is data in the index, this value is always 100% because the Sun StorEdge EBS software no longer uses disk space reclamation.

Note – The Reclaim Space button is included for backward compatibility with Sun Solstice Backup software prior to release 6.0. For release 6.0 and later, this feature is no longer necessary.

Index Save Sets

The Save Sets window of the Indexes dialog box displays the save sets assigned to a particular client along with information, such as the save set name, size, and number of cycles. The Save Sets window also contains buttons for viewing details on a selected save set and removing old save set cycles.

▼ To View Information about a Client's Save Sets

To view the information about a client's save sets, select the client with save sets to be viewed in the Clients window. The save set information for that client appears in the Save Sets window of the Indexes dialog box.

[TABLE 3-11](#) identifies the information displayed in the Save Sets window for each save set associated with the selected client.

TABLE 3-11 Save Sets Window Description

Column	Description
Name	Name of the save set.
Size	Estimated amount of the index space used by the save set in the client file index.
Cycles	Number of backup cycles available for browsing. A cycle starts with a full backup and ends with the next full backup, and includes any incremental and level 1–9 backups that occur between full backups.

Note – The Remove Oldest Cycle button in the Save Sets Window can be used reduce the size of the client file index by deleting entries for the oldest version of the save set. For more information about reducing the size of the client file indexes, see [“Reducing Client File Index Size” on page 96](#).

Save Set Details


The Instances Details window provides detailed information about a specific save set. Refer to this information to determine how the resources are being used. For example, by seeing how large a save set is, you can plan the amount of disk space needed for the online indexes.

▼ To View Save Set Details

To view the save set details:

1. From the Clients menu, select Index to open the Indexes dialog box.
2. Select the client whose save sets you want to view.
3. Select a save set from the Save Sets window.
4. Click the Instances button to open the Instances window.

[FIGURE 3-2 on page 93](#) shows the Instances window.



The screenshot shows a window titled 'Instances' with a subtitle 'Instances of mercury:/file_sys1'. It contains a table with five columns: 'Id', 'Files', 'Size', 'Date', and 'Level'. The table lists 20 rows of data, each representing a save set. The 'Id' column contains long numeric strings. The 'Files' column shows the number of files. The 'Size' column shows the size in MB or KB. The 'Date' column shows the date in MM/DD/YY format. The 'Level' column shows the backup level, with 'full' for the first row and 'incr' for the others. At the bottom of the window are four buttons: 'Cancel', 'Print', 'Save', and 'Help'.

Id	Files	Size	Date	Level
1457624321	3246	2026 MB	01/06/01	full
1476954369	0	0 KB	01/07/01	incr
1520887810	33	115 MB	01/08/01	incr
1542986498	1666	1132 MB	01/09/01	incr
1565131779	4353	158 MB	01/10/01	incr
1587308545	6	793 KB	01/11/01	incr
1654035969	6	796 KB	01/15/01	incr
1676142849	19	819 KB	01/16/01	incr
1697791745	2440	770 MB	01/16/01	incr
1720165124	45	116 MB	01/18/01	incr
1742042116	20	31 MB	01/18/01	incr
1764190466	31	114 MB	01/19/01	incr
1830545153	33	117 MB	01/22/01	incr
1852603906	17	31 MB	01/23/01	incr
1874717186	7	31 MB	01/24/01	incr
1919092227	876	3529 MB	01/26/01	incr
1941079554	23	1855 MB	01/27/01	incr
1963197953	6	61 MB	01/28/01	incr
1985336834	60	2160 MB	01/29/01	incr

FIGURE 3-2 Instances Window

[TABLE 3-12](#) identifies the information displayed for each save set instance.

TABLE 3-12 Instances Window Description

Column	Description
ID	Unique identification number of the instance of the save set.
Files	Number of files backed up during that instance.
Size	Size of the backup.
Date	Date of the backup.
Level	Level of the backup (full, incr [incremental], or 1 to 9)

Location of the Client File Index

During the initial client setup, the Sun StorEdge EBS software normally designates a default location for the client's client file index on the Sun StorEdge EBS backup server. This default location is `/nsr/index/client_name`. However, there may be times when you 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.

▼ To Designate the Client File Index Location

Perform this procedure only when configuring a new Client resource or before the first backup of a new client occurs, and only if you want to use a nondefault location for the client's index. If you want to change the location of an existing client's file index, you must move that index. For more information, see [“To Move a Client File Index” on page 95](#).

To designate the client file index location:

1. In the Sun StorEdge EBS Administrator program, select Client Setup from the Clients menu.
2. From the Clients list, select the client.
3. From the View menu, select Details to display hidden attributes.
4. In the Index Path attribute, enter the full path of the directory where the client file index will be stored.
5. Click Apply.

6. Navigate to the directory where the client file index for the client resides and enter the following at the command prompt:

```
# nsrck -C client_name
```

If no problems are found when this command is run, then all future client file index information is saved to the designated location.

▼ To Move a Client File Index

In some cases, you might want to move a client's 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. Run the `nsrinfo client_name` command and redirect the output to a file, for example:

```
# nsrinfo client1 > /tmp/nsrinfo1
```

2. Shut down the Sun StorEdge EBS daemons.
3. Move the existing client file index to the new location by entering the following at the command prompt:

```
# mv previous_index_location/client_name new_index_location
```

For example:

```
# mv /nsr/index/jupiter /space/index
```

4. Start the Sun StorEdge EBS daemons.
5. In the Sun StorEdge EBS Administrator program, select Client Setup from the Clients menu to open the Clients window.
6. Select the client from the Clients list.
7. From the View menu, select Details to display hidden attributes.
8. In the Index Path attribute, enter the full path of the directory where the client file index will be stored.

9. Click Apply.
10. Run the `nsrinfo client_name` command again and redirect the output to a file.
For example:

```
# nsrinfo client1 > /tmp/nsrinfo2
```

11. Verify that the contents of the file created in [Step 10](#) are the same as the contents of the file created in [Step 1](#).

All future client file index information is saved to the new designated location.

Managing the Size of the Online Indexes

Over time, the size of the Sun StorEdge EBS server's online indexes can become prohibitively large. You can reduce the size of these indexes by using the solutions suggested in the following sections:

- [“Reducing Client File Index Size” on page 96](#)
- [“Reducing Media Database Size” on page 97](#)
- [“To Remove the Oldest Save Set Cycles” on page 98](#)
- [“Removing Volume-Based Entries from the Online Indexes” on page 98](#)
- [“To Compress the Media Database” on page 100](#)

Reducing Client File Index Size

You can reduce the size of the Sun StorEdge EBS server's client file indexes by using one or more of the following procedures:

- Remove save sets that comprise the oldest backup cycle from the client file index. For more information, see [“To Remove the Oldest Save Set Cycles” on page 98](#).
- Remove volume-based entries from the client file index. For more information, see [“Removing Volume-Based Entries from the Online Indexes” on page 98](#).
- 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 only works 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 [“Browse and Retention Save Set Policy Modification” on page 133](#).

If the size of the client’s online file index is still too large, consider moving the location of the index. For more information, see [“To Move a Client File Index” on page 95](#).

Reducing Media Database Size

Reduce the size of Sun StorEdge EBS server’s media database by using one or more of the following procedures:

- 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 98](#).

When a volume is removed from the media database, the entries associated with that volume are removed from the media database and the client’s online file index. If you select this option, you will still be able to recover the volume’s data using the scanner program.

Note – You gain very little disk space by removing a media database entry. By leaving a volume’s index entries in the media database, you are prevented from accidentally labeling another volume with the same name.

- Recycle volumes that contain recyclable save sets. For more information, see [“To Change the Mode of a Volume” on page 241](#).

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 contents 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 [“Browse and Retention Save Set Policy Modification” on page 133](#).

- Compress the media database. For more information, see [“To Compress the Media Database” on page 100](#).

▼ 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. In the Sun StorEdge EBS Administrator program, select Indexes from the Clients menu to open the Indexes dialog box.
2. Select the save sets whose oldest cycle you want to remove, and click Remove Oldest Cycle.

After the Remove Oldest Cycle operation has finished, the statistics in the Indexes dialog box are updated to reflect the current state of the 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:

```
# nsrmm -d -P -S ssid
```

where `ssid` is the save set ID for the save set.

Use `mminfo` to determine the save set ID:

```
# mminfo -v -c client_name
```

For more information, refer to the `nsrmm` and `mminfo` man pages.

When save sets are marked recoverable, users may no longer browse and recover these files using the `nwrecover` program. 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 [“Save Set Recovery” on page 303](#).

Removing Client File Index and Media Database Entries

You can remove both the client file index and media database entries. This action removes all knowledge of the volume from the Sun StorEdge EBS server. You should only remove a volume from the media database if you think the volume has been physically damaged and is unusable. However, even if you remove the database entries for a volume, as long as the volume is undamaged, the data remains recoverable by using the `scanner` program. For more information on the `scanner` program, see [“The scanner Program” on page 316](#).

Typically, you do not want to 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 volume’s entry 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.

▼ To Remove Online Index Entries

To remove volume-based entries from the online indexes:

1. In the Sun StorEdge EBS Administrator program, select Volumes from the Media menu to open the Volumes window.
2. Select the volume whose entry you want to remove from the online indexes.
3. From the Volume menu, select Remove.
4. Select OK to verify that the volume should be removed.
5. Click Apply.

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 command:

```
# nsrmm -d -S ssid
```

To remove all information related to a particular volume, enter the following command:

```
# nsrmm -d volume_name
```

▼ 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 `/nsr/mm/.cmprssd`.
2. Enter the following command:

```
# nsrim
```

Server Maintenance Tasks

This section describes tasks you might need to perform after you install and configure the Sun StorEdge EBS server.

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 `/nsr/logs` directory. You can use environment variables and/or startup scripts to automatically control the size of the log.

▼ To Set the 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 lines:

```
(echo 'starting Sun StorEdge EBS daemons:') > /dev/console
```

Enter the environment variables 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 1,024 KB.
- 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=512  
export NSR_MAXLOGSIZE
```

Note – You must 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:

- The `daemon.log` file is renamed `daemon.001` and a new empty `daemon.log` is created.
- 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.
- This process is repeated until the value in `NSR_MAXLOGVERS` is reached, at which point the highest numbered log is removed.



Caution – The trimming mechanism only functions when you start `nsrd`. The `nsrd` daemon does not check periodically to determine if the log file has exceeded `NSR_MAXLOGSIZE`. If the `nsrd` daemon runs for a long time, the log file can still grow very large. To activate the trimming mechanism, enter `nsr_shutdown` to stop the Sun StorEdge EBS daemons, and then use the Sun StorEdge EBS startup script to restart the daemons.

Client Operations

The Sun StorEdge Enterprise Backup Software product provides preconfigured settings for backing up clients, including the server, which is considered a client of itself. You can use these preconfigured settings, or create new settings. This chapter provides information that you might need when configuring Sun StorEdge EBS clients. The information describes Sun StorEdge EBS clients, the associated characteristics of a client and/or a save set, and how to best customize client configurations to suit the needs of your environment.

The main sections in this chapter are:

- [“Sun StorEdge EBS Client” on page 104](#)
- [“Client Configuration” on page 104](#)
- [“Multiple Clients from the Same Computer” on page 116](#)
- [“Save Set Status Values” on page 118](#)
- [“Browse and Retention Policies” on page 120](#)
- [“Browse and Retention Save Set Policy Modification” on page 133](#)
- [“Backup Command Customization” on page 136](#)
- [“Directives” on page 144](#)
- [“Adding or Changing the Sun StorEdge EBS Servers for a Client” on page 154](#)
- [“Client Priority” on page 155](#)
- [“Logical Volume Backup” on page 156](#)
- [“Sun StorEdge EBS Client Licensing” on page 157](#)
- [“Legato NetWorker ClientPak for Mac OS X” on page 159](#)

Sun StorEdge EBS Client

A Sun StorEdge EBS 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 that are 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 [“Multiple Clients from the Same Computer” on page 116](#).

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 Linux computer.

Client Configuration

This section provides an overview of how Sun StorEdge EBS clients are created and integrated into the Sun StorEdge EBS storage management environment. The following information is included:

- [“Client Name” on page 105](#)
- [“Save Sets” on page 105](#)
- [“To Create a Client” on page 110](#)
- [“To Edit a Client” on page 112](#)
- [“To Delete a Client” on page 113](#)
- [“To Restore a Deleted Client” on page 113](#)
- [“The Client ID” on page 113](#)
- [“To Change a Client Name” on page 114](#)
- [“Editing a Client’s nsrla.res Database” on page 115](#)

Sun StorEdge EBS client software must be installed on computers that are clients to the Sun StorEdge EBS server. Because each Sun StorEdge EBS server also backs itself up, each Sun StorEdge EBS server also has the client software installed. Client resources are created on the Sun StorEdge EBS server that specify the configuration choices for each Sun StorEdge EBS client.

Client resources determine the following:

- The data on the client filesystem to be backed up (referred to as the [save set](#)).
- The schedule that is used to back up that data.
- The length of time that the client's backed up data is available for browsing and recovery.
- Any additional features, such as archiving, that are enabled.

Client Name

The name that you give the Client resource must be a valid hostname for a computer that has the Sun StorEdge EBS client running on it.

If you are creating multiple Client resources for the same computer, you must use the same name for each Client resource. This allows the Sun StorEdge EBS server to associate different collections of backed up data with the same client. Unique settings for the Client resource attributes, such as the Save Set attribute, distinguish one Client resource from another.

There are a number of situations where you may want to create multiple Client resources for the same client. For example, if you use a Legato NetWorker Module to back up database files, you may want to create one Client resource to back up the database files and another to back up the client computer's operating system files.

Save Sets

The primary function of a Client resource is to identify the client data to be backed up by the Sun StorEdge EBS server. The collection of all data items backed up during a backup session between the server and client is called the *save set*. A save set can consist of a group of files or entire filesystems, as well as application-generated data such as a database or operating system information.

The client's save set is defined in the Save Set attribute of the Client resource. Multiple entries can be included in this attribute. For example, you can enter multiple files or filesystems, or a combination of both:

```
/usr  
/home  
/etc/hosts
```

Note – If a save set's name includes a symbolic link, recovery by save set ID is not supported.

If you are backing up large client filesystems, you can use the Save Set attribute to balance load for the client. For example, you could create multiple Client resources, dividing the large filesystems among each resource. You can then create different schedules for each resource to better distribute the client's backup load.

Predefined Save Sets

In addition to entering files and/or filesystems in the Save Set attribute of the Client resource, you can also use predefined save sets when configuring the Sun StorEdge EBS client. UNIX clients have one predefined save set, the All save set. For more information, see [“The All Save Set” on page 107](#).

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

- All
- SYSTEM STATE
- REPAIRDISK

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

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

For Microsoft 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 detailed information about VSS, refer to the *Legato NetWorker Administrator's Guide, Microsoft Windows Version*.

The All Save Set

The default setting used for the Save Set attribute when a client is first created is All, which refers to the All save set.

- For UNIX operating systems, the All save set generally backs up all locally mounted filesystems.
- For Microsoft Windows operating systems, the All save set *also* backs up operating system files that determine the client system's state.

TABLE 4-1 provides a list of the save sets included in the All save set.

TABLE 4-1 Components in the All Save Set

Operating System	Files/Save Sets
UNIX	Backs up all filesystems listed in the following locations: • Solaris and IBM DYNIX/ptx: /etc/vfstab • HP-UX, HP Tru64, Linux, and SGI: /etc/fstab • AIX: /etc/filesystems
Microsoft Windows XP Professional, Microsoft Windows Server 2003 with no VSS client license or with VSS disabled	SYSTEM STATE SYSTEM DB SYSTEM FILES SHAREPOINT ASR All local, physical drives
Microsoft 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
Microsoft Windows 2000	SYSTEM STATE SYSTEM DB SYSTEM FILES SHAREPOINT ASR All local, physical drives
Microsoft Windows NT 4.0	SYSTEM STATE (includes the registry) REPAIRDISK All fixed disks

SYSTEM and VSS SYSTEM Save Sets

Microsoft Windows operating systems contain files and databases that define the state of the computer at any given time. A simple backup of the computer's filesystem does not capture system state information. Sun StorEdge EBS software provides five SYSTEM save sets and six VSS SYSTEM save sets to back up the files that define the system state.

[TABLE 4-2](#) provides a list of the components in each of the SYSTEM save sets used with Microsoft Windows Server 2003 with no VSS client license or with VSS disabled, Microsoft Windows XP Professional, or Windows 2000 operating systems.

TABLE 4-2 Components in the SYSTEM Save Sets

Save Set	Components
SYSTEM STATE	ASR System Files; most elements of the Windows System State
SYSTEM FILES	Windows System File Protection (SFP) components; IA-64 Extensible Firmware Interface (EFI) FAT partition
SYSTEM DB	Windows system databases
SHAREPOINT	Web Storage System files; Microsoft Search resources; 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; and any shortcuts or content sources that reference the local filesystem
ASR	Information necessary to perform an automated reinstallation of Windows, restoration of the system configuration, and recovery of one or more disk volumes to return the failed computer to its condition at the time of the last ASR backup

[TABLE 4-3](#) provides a list of the components in each of the VSS SYSTEM save sets used with Microsoft Windows Server 2003 with VSS licensed and enabled.

TABLE 4-3 Components in the VSS SYSTEM Save Sets

Save Set	Components
VSS SYSTEM BOOT	All elements of the Windows system state except the VSS System Writer writer; ASR System Files legacy component
VSS SYSTEM FILESET	VSS System Writer writer; IA-64 EFI FAT partition legacy component

TABLE 4-3 Components in the VSS SYSTEM Save Sets

Save Set	Components
VSS SYSTEM SERVICES	All elements of the Windows System Services; Disk Quota database legacy component
VSS USER DATA	Application data that is not system-critical
VSS OTHER	Application data that is not system-critical
VSS ASR DISK	ASR legacy component

These save sets can be specified in the Save Set attribute of the Client resource, along with file-based save set specifications. For example, the following entries in the Save Set attribute would back up the following five SYSTEM save sets, as well as the C: drive, on a computer that is running Microsoft Windows Server 2003 with no VSS client license or with VSS disabled, Microsoft Windows XP Professional, or Microsoft Windows 2000:

SYSTEM STATE:
SYSTEM FILES:
SYSTEM DB:
SHAREPOINT:
ASR:
C:

The following entries in the Save Set attribute would back up the following six VSS SYSTEM save sets, as well as the C: drive, on a computer running Microsoft 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:
C:

For detailed information about the contents of the Microsoft Windows SYSTEM and VSS SYSTEM save sets, refer to the *Legato NetWorker Administrator's Guide, Microsoft Windows Version*.

▼ To Create a Client

Before a client can be backed up by a Sun StorEdge EBS server, the client computer must have the appropriate Sun StorEdge EBS client software installed. For more information, refer to the *Installation Guide*.

The steps presented here are a general guideline for client configuration. However, additional attributes may be present, particularly if hidden attributes are displayed (see [“Hidden Attributes” on page 42](#)). For information about these additional attributes, refer to Online Help.

To add and configure a new Sun StorEdge EBS client:

1. In the Sun StorEdge EBS Administrator program, select Client Setup from the Clients menu.
2. Click Create.
3. In the Name attribute, enter the client’s hostname. For more information, see [“Client Name” on page 105](#).

If hidden attributes are displayed, the Client ID attribute appears after the name. This attribute should be left blank; the Sun StorEdge EBS server assigns an ID to the client. For more information, see [“The Client ID” on page 113](#).
4. (Optional) In the Comment attribute, enter a description of the client.
5. In the Archive Services attribute, select disabled if you do not want this client to be able to use the Archive feature. For more information, see [“Enabling Archive Services for a Sun StorEdge EBS Client” on page 328](#).
6. In the Schedule attribute, select a schedule. For more information, see [“Schedules” on page 263](#).
7. In the Browse Policy attribute, select a browse policy. See [“Browse and Retention Policies” on page 120](#) for more information.
8. In the Retention Policy attribute, select a retention policy. For more information, see [“Browse and Retention Policies” on page 120](#).
9. In the Directives attribute, select a directive. For more information, see [“Directives” on page 144](#).
10. In the Groups attribute, select a group to which your client will be added. For more information, see [“Backup Groups” on page 250](#).
11. In the Save Set attribute, enter the files, filesystems, or predefined save set.

Enter All to back up all of the client’s filesystems. Multiple entries can be included in this attribute by entering the information and clicking Add. For more information about save sets, see [“Save Sets” on page 105](#).

12. In the Remote Access attribute, enter the user and user's hostname of other clients that are allowed to back up or recover this client's files, by using the formats specified in [“Administrator Privileges” on page 59](#).

If you enter a hostname or `host=hostname`, any user on that host is allowed to recover this client's files. To enter a username without specifying the host, you must enter `user=name`.

Leave this attribute blank if you do not want remote users to browse and recover this client's data. Users who are already defined as administrators of the Sun StorEdge EBS server already have remote access to the client and do not need to be included in this attribute. This attribute is also used to enable directed recovery; for more information, see [“Directed Recovery” on page 294](#).

Note – Only members of the Administrators group, or members of a group with the Change Security Settings privilege, can change the Remote Access attribute. For more information, see [“Sun StorEdge EBS User Groups” on page 61](#).

13. In the Remote User attribute, enter the user ID of a client to allow remote access and operations on this client.

Note: The Password attribute, immediately following the Remote User attribute, is used only for clients running NetWare. If you enter a password in this attribute, it is used by the Sun StorEdge EBS server to perform remote commands on this client.

14. Optional: In the Backup Command attribute, enter the name of a program file to instruct the Sun StorEdge EBS server to execute a customized backup command to back up the client data. For more information, see [“Backup Command Customization” on page 136](#).

15. Leave the Server Network Interface (SNI) attribute blank to instruct the Sun StorEdge EBS software to use the default network interface card (NIC) associated with the server's hostname. This is a hidden attribute.

In some cases, such as private backup networks, the server's hostname is not the interface that the client should use for backups and recoveries. In these cases, use this attribute to specify the unique hostname of the preferred interface.

When using an SNI, be sure to include this hostname in the list of servers allowed to backup the client. For more information, see [“Adding or Changing the Sun StorEdge EBS Servers for a Client” on page 154](#).

16. In the Aliases attribute, enter any nicknames for this client computer that queries can match. Otherwise, queries match only the client name.

Note – Each Client resource should include both the DNS short name and long name in the Aliases attribute.

17. In the Archive Users attribute, enter each user who is authorized to use the archive services.

If no users are listed, then only administrators are allowed to use the archive services on the client. The backslash (\) and the at symbol (@) are not allowed in the user names for this attribute.

Note – Only members of the Administrators group, or members of a group with the Change Security Settings privilege, can change the Archive Users attribute. For more information, see [“Sun StorEdge EBS User Groups” on page 61.](#)

18. In the Storage Nodes attribute, enter the name of each storage node to which the client can back up data.

The client backs up to the first active enabled storage node in the order they appear in the attribute. The default storage node *nsrserverhost* represents the server.

19. In the Clone Storage Nodes attribute, enter the storage node names for cloned data in the same manner as [Step 18.](#)

This attribute applies to clients that are also serving as Sun StorEdge EBS storage nodes; no information should be entered for clients that are not storage nodes.

20. Complete other attributes as necessary, and click Apply.

Note – If you try to add more clients beyond the number of connections you purchased for your Sun StorEdge EBS server, you receive the error message:

Too many clients - maximum is *number*.

▼ 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 [“To Change a Client Name” on page 114.](#)

To edit a client:

1. In the Sun StorEdge EBS Administrator program, select Client Setup from the Clients menu.
2. In the Clients list, select the appropriate client.
3. Edit the attributes of the client, and click Apply.

▼ To Delete a Client

When you delete a client, 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 explicitly removed. The browse and retention policies for the remaining entries are replaced with the policies assigned to the Sun StorEdge EBS server, defined as a client of itself.

Note – If you delete a client from the Sun StorEdge EBS server, the client’s previously backed-up data is still accessible and recoverable directly from the volume containing the data by using the **scanner** program. For more information, see [“The scanner Program” on page 316](#).

To delete a client:

1. In the Sun StorEdge EBS Administrator program, select Client Setup from the Clients menu.
2. In the Clients list, select the client to be deleted.
3. Click Delete.
4. Click OK to confirm the deletion.
5. Click Apply.

▼ To Restore a Deleted Client

After you delete a client, the Sun StorEdge EBS server can restore the client.

To restore a deleted client:

1. In the Sun StorEdge EBS Administrator program, create a new client.
2. In the Name attribute, enter the name of the deleted client.

The Sun StorEdge EBS server recalls the client ID for this name and inserts it into the Client ID attribute.

The Client ID

The Client ID attribute of the Client resource allows you to change the name of a Client resource while still maintaining the uniqueness and integrity of that resource. Normally, leave this attribute blank when creating a Client resource, which causes

the server to automatically assign a value. The only time you need to access the Client ID value of an existing client, or enter a value in the Client ID attribute, is when you need to change the name of a client or restore a deleted client.

The Sun StorEdge EBS server determines a value for this attribute based on the following criteria:

- When creating the first instance of a Client resource, the server creates a new unique identifier.
- When creating subsequent instances of a Client resource, the new instance inherits the Client ID value from the existing instance of the resource (that is, all Client resources of the same name have the same identification number).

If you create a Client resource and later create a new Client resource using the same client name, the Sun StorEdge EBS server recalls the Client ID used previously for this name and reuses it for the new resource.

▼ To Access the Client ID

To access the client ID:

1. In the Sun StorEdge EBS Administrator program, select Client Setup from the Clients menu.
2. From the View menu, select Details to display hidden attributes.
3. In the Clients list, select the appropriate client.

The client ID is displayed in the Client ID attribute.

▼ To Change a Client Name

Note – This procedure is not valid for renaming the client residing on a Sun StorEdge EBS server.

To change a client name:

1. Write down the Client ID of the original client.
For more information on understanding and accessing the Client ID attribute of the Client resource, see [“The Client ID” on page 113](#).
2. In the Sun StorEdge EBS Administrator program, delete all instances of the original client.
3. Stop all Sun StorEdge EBS daemons.

4. On the Sun StorEdge EBS server that backs up this client, rename the directory containing the client file index for this client from `old_client_name.domain.com` to `new_client_name.domain.com`. The default location for the client file index is:
 - For UNIX:
`/nsr/index/client_name.domain.com`
 - For Microsoft Windows:
`<NetWorker_install_path>\index\client_name.domain.com`
5. Restart the Sun StorEdge EBS daemons.
6. Create a new Client resource. For more information, see [“To Create a Client” on page 110](#).
7. In the Client ID attribute, enter the Client ID value of the original client.
You must have hidden attributes displayed to access the Client ID; for more information, see [“Hidden Attributes” on page 42](#).
8. Complete the other attributes as necessary, and click OK.

Editing a Client’s nsrla.res Database

The `nsrla.res` database is controlled by the client daemon, `nsrexecd`. This 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 must be edited using the character-based `nsradmin` program.

Note – For information about the Disable Directed Recover attribute in the NSRLA resource, see [“The Disable Directed Recover Attribute” on page 299](#).

▼ To Edit the nsrla.res Database

To edit the `nsrla.res` database:

1. Log in as root on the Sun StorEdge EBS client.
2. Enter the following at the command prompt:

```
# nsradmin -s server -p nsrexecd -c type:NSRLA
```

3. Select Edit from the Command menu.

4. Make changes to any appropriate attributes by arrowing to the attribute and typing the new information.
5. To save changes, press the Esc key and select Yes to confirm.

▼ To Display Hidden Attributes in the nsradmin Program

Some attributes, such as the Disable Directed Recovery attribute, are hidden and must be displayed to be edited.

To display hidden attributes:

1. Open the nsradmin program as described in [“To Edit the nsr_la.res Database” on page 115](#).
2. Enter the following at the nsradmin prompt:

```
nsradmin> option hidden
```

For information about the attributes in the NSRLA resource, refer to the nsr_la.5 man page or the *Legato Command Reference Guide*.

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's save sets are 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 filesystem 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.

Using Unique Combinations of Clients and Save Sets

Save sets are collections of data to be backed up by the Sun StorEdge EBS server. 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 Edit Client dialog boxes. Save sets can be any of the following:

- All of the data on a client (this is the default condition indicated by the value All)
- Partition on a disk
- Single directory
- Single file
- Raw partition (such as from a database)

If a client has a large volume of data, you may want to schedule the client computer for several, separate client and save set backups. By redefining a large filesystem into multiple client and save set instances, you are able to:

- Automatically back up a large client filesystem.
- Balance the load by avoiding a full backup of the entire filesystem at one time.

You can redefine a single filesystem into a series of separate client and save set instances for the filesystem. 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 runs its full backup on a different day of the week. For more information on schedules, see [“How the Sun StorEdge EBS Server Uses Schedules” on page 264](#).

Client and Save Set Combination

The same save set can appear in the Client resource's Save Set attribute for multiple client instances. This characteristic permits you to associate the same save set with more than one group or schedule for backup.

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 listed in the Client resource.

- To configure the client so that a specific filesystem is backed up, change the value in the Save Set attribute to the filesystem pathname.
- To define more than one save set on a client, enter each save set (partition, filesystem, or file) on a separate line in the Save Set attribute.

Save Set Status Values

The Sun StorEdge EBS server assigns to each backed-up save set a status based on the success of the backup or the age of the save set data. The status of a save set is displayed in the Volumes window.

[TABLE 4-4](#) provides a list of all the possible values for save set status.

TABLE 4-4 Save Set Status Values *(1 of 2)*

Status Value	Meaning	Description
abort	aborted	You aborted the backup for this save set manually, or a crash occurred during the operation. This save set is considered immediately eligible for recycling.
brows	browsable	The files in this save set retain entries in the client file index. You can restore all the files using an index-based recover.
inpro	in progress	This save set is currently being backed up.
recov	recoverable	The files in this save set do not have browsable entries in the client file index, and have not passed the retention policy.
recyc	recyclable	The save set, and all save sets dependent on this save set for recovery, have exceeded their retention policies.

TABLE 4-4 Save Set Status Values (2 of 2)

Status Value	Meaning	Description
scann	scanned-in	The entry for this save set was restored by using the scanner program. The policies were reset to the values that were applied when the save set was created (or, if the save set was created on Sun StorEdge EBS servers from releases prior to 6.0, when the save sets were upgraded to the 6.x format). For more information, see “The scanner Program” on page 316 .
susp	suspect	An attempt to recover this save set failed. The <code>recover</code> command could not read all the blocks of the save set. This can occur, for example, when there is a bad spot on the tape.
p	purged	The save set has been purged from the media database.

▼ To View the Status of a Save Set

To view the status of a save set:

1. In the Sun StorEdge EBS Administrator program, select Volumes from the Media menu to open the Volumes window.

Note – You can also click the Volumes button to open the Volumes window.

2. Select the appropriate volume.

The status of save sets on this volume are listed in the Status column of the Save Sets window. The save set status might change in any of the following situations:

- The save set has exceeded its browse policy. When this condition occurs, the save set status changes from browsable to recoverable.
- The save set has exceeded its retention policy and all save sets dependent on this save set have also exceeded their retention policies. When this condition occurs, the save set status changes from recoverable to recyclable.
- An administrator manually changes the save set status.

▼ To Change the Status of a Save Set

The Sun StorEdge EBS server marks a volume suspect if an error occurred while reading the media during backup. You may want to change the status of the save set the server marked as suspect to normal if you know that the save set data is not

really suspect, but that the server had difficulty reading the data for other reasons. For example, perhaps the media drive heads were dirty, causing the server to mark the save set suspect.

In some cases, you may want to mark a functioning save set as suspect in order to omit it from recoveries. This is necessary if you want to recover data from a cloned volume rather than from the original, because the Sun StorEdge EBS server will always look for the original save set during recoveries. Only when the original is suspect or nonrecoverable will the server then request a cloned save set and its corresponding clone volume. This process requires that a cloned volume containing the cloned save set be available. If a cloned volume containing the cloned save set is not available, then the suspect volume is used.

To change the status of a save set:

1. Display the save set status. For instructions, see [“To View the Status of a Save Set” on page 119](#).
2. If the status is *suspect*, select Normal from the Save Set>Change Status menu to change the status to normal. If the status is normal, select Save Set>Change Status>Suspect to change the status to suspect.

For more information about these status values, see [TABLE 4-4 on page 118](#).

Browse and Retention Policies

This section covers the following topics:

- [“How the Sun StorEdge EBS Server Uses Browse and Retention Policies” on page 121](#)
- [“Managing the Data Life Cycle” on page 127](#)
- [“Assigning Multiple Policies to a Single Client” on page 129](#)
- [“Preconfigured Policies” on page 129](#)
- [“To Create a Policy” on page 130](#)
- [“To Edit a Policy” on page 131](#)
- [“To Delete a Policy” on page 131](#)
- [“Browse and Retention Policies for Manual Backups” on page 132](#)

How the Sun StorEdge EBS Server Uses Browse and Retention Policies

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. Each time a backup is completed, the server creates entries for the backed-up files in the client file indexes. The *media database* stores one entry for each save set and storage volume during each backup operation.

Each client file index is a browsable structure of data from a single client computer. Users can specify anything from a single file to a complete filesystem. The Sun StorEdge EBS server can then be directed to reconstruct the data during a recover session to look exactly as it did at a specific time. The information that the client index contains and coordinates enables the server to automatically handle situations. Such situations include assembling data from backups based on levels, and renaming or deleting files or directories. The server uses browse policies to manage the life cycle of data and to automatically control the size of the client file index.

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 in the Sun StorEdge EBS recover program (nwrecover), and select individual files or entire filesystems for recovery. After the browse policy for a file is exceeded, the Sun StorEdge EBS server automatically deletes the entry for that file. The server deletes these entries to manage the size of the client index, which can grow rapidly: one entry for each file backed up during each scheduled backup of the client.

The *retention policy* determines how long save sets are maintained in the Sun StorEdge EBS server's media database. For at least the period of the retention policy, you can recover a client's backed-up save sets from media:

- No save set is considered recyclable until, at a minimum, it has exceeded its retention policy.
- No storage volume can be relabeled and written over until, at a minimum, all save sets on the storage volume (including save sets that depend on them) have exceeded their retention policies.

Theoretically, entries for a save set or a storage volume can remain in the media database forever, long after the retention policy has been exceeded. Entries are removed from the media database under these conditions:

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

The media database is the structure that tracks the location of save sets on storage volumes. The Sun StorEdge EBS server uses the retention policy to manage the longevity of Sun StorEdge EBS managed data. Data is recoverable as long as entries

exist in the media database; there is nothing to be gained by rushing to delete media database entries. For all these reasons, the media database retention policy does not trigger the automatic removal of media database entries. Instead, the retention policy determines how long an entry for a save set remains protected from being accidentally written over.

How the Browse Policy Works

You can recover a file that has an entry in the client file index through the Sun StorEdge EBS recover program (`nwrecover`). This enables users to browse and mark files, and initiate data recovery. Client file index entries are not necessarily deleted the same day that the browse policy is exceeded. 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 263](#), and for more information about backup levels, see [“Backup Levels” on page 273](#).

Example: One-Week Browse Policy

In [FIGURE 4-1 on page 123](#), 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 4-1 One-Week Browse Policy

To further illustrate, suppose that on October 12, you decide that you want to 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 combination 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 the example shown in [FIGURE 4-1 on page 123](#), entries from the Week 1 backup cycle are removed from the client file index on October 15.

Example: Two-Week Browse Policy

In [FIGURE 4-2 on page 124](#), 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 4-2 Two-Week Browse Policy

How the Retention Policy Works

The Sun StorEdge EBS media retention policy specifies a period during which backed-up data is protected from accidental overwrite. After the retention period is exceeded, the save set is eligible to change its status from *recoverable* to *recyclable*. The term *recyclable* is best understood as “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.

The Sun StorEdge EBS server changes the mode of that storage volume to 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* is a passive reminder that you can overwrite the volume if conditions are right. The volume is available for relabel and use if you do the following:

- Place the volume in an autochanger or mount the volume in a stand-alone device.
- Enable the Auto Media Management attribute in the Device resource.

The existing data is nonrecoverable after the volume is relabeled, so the entries for the overwritten save sets are removed from the media database. For more details about this feature of auto media management, see [“How the Sun StorEdge EBS Server Selects a Volume for Relabeling” on page 233](#).

Save set entries are also removed from the media database when you manually delete a volume from the Sun StorEdge EBS volume inventory. 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 program.
- If you re-create the save set's entries in the media database, the root user can recover data by using save set recovery.

For more information on how to use the `scanner` program, see [“The scanner Program” on page 316](#).

Example: Three-Week Retention Policy

FIGURE 4-3 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.

Managing the Data Life Cycle

The browse and retention policies that you associate with a client save set control both the growth of the client file index and the media database, and how long data remains available for recovery.

[FIGURE 4-4 on page 128](#) traces the data life cycle 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 `nwrecover` program. 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 itself changes from either full or appendable to recyclable.

While the status of a save set is either recoverable or recyclable, you can recover any save set from the storage volume 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 `nwrecover` program. For more information about using save set recover and the `scanner` program, see [“Save Set Recovery” on page 303](#) and [“The scanner Program” on page 316](#).




FIGURE 4-4 Data Life Cycle 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 226](#).



Caution – 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. Neither the save set recover feature nor the `scanner` program can be used to recover the data. At this point, the volume is ready for new data; 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 from the original client. The Sun StorEdge EBS server employs the Browse Policy and Retention Policy attribute values that correspond to the unique combination of values from the Client resource's Name, Save Set, and Group attributes.

Example: Assigning Different Policies for an Identical Client

You already have client *saturn* with a save set value of All and assigned to backup group *General*. The browse policy is Weekly and the retention policy is Monthly. You decide to create an otherwise identical copy of *saturn*, but assign it to backup group *Special*. For this version of *saturn* you designate the browse policy as Weekly and the retention policy as 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 you can 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 [“To Create a Policy” on page 130](#).

▼ To Create a Policy

Resource attributes vary depending on the server. Use the steps here as a general guideline. However, additional attributes may be present (particularly if Hidden Attributes is enabled); similarly, attributes mentioned here may be absent. For guidance in completing these additional attributes, refer to Online Help.

To create a policy:

1. In the Sun StorEdge EBS Administrator program, select Policies from the Customize menu to open the Policies window.
2. Click Create.
3. In the Name attribute, enter a name for the policy.

Choose a name that describes how long the policy lasts. For example, the name Quarter indicates the policy lasts a quarter of a year. The name you choose must be unique for the Sun StorEdge EBS server.

4. From the Period list, select a period.

This attribute works in conjunction with the Number of Periods attribute. You can choose days, weeks, months, or years as the period. For example, a quarterly policy is based on the period of a month (a quarter is three months). The Sun StorEdge EBS software defines periods in the following ways:

- Week: Seven days beginning on Sunday.
- Month: The number of days in the month during which a backup occurs; for example, for backups performed in February, the month period is 28 days.
- Year: 366 days.

5. In the Number of Periods attribute, enter a number.

This value indicates the number of periods applied to the policy. For example, if you choose month for the Period attribute, and three for the Number of Periods attribute, then the policy lasts for three months (one quarter).

6. Complete the other attributes as necessary, and click Apply.

Note – The Sun StorEdge EBS server does not allow a client's browse policy to exceed its retention policy. This means a file index entry must be removed before the volume on which the file is stored can be marked *recyclable* or removed from the media database.

▼ To Edit a Policy

Note – You cannot change the name of a policy.

To edit a policy:

1. In the Sun StorEdge EBS Administrator program, select Policies from the Customize menu.
2. Select the appropriate policy.
3. Edit the attributes and click Apply.

▼ To Delete a Policy

Note – You cannot delete a preconfigured policy, or any policy currently assigned to a client.

To delete a policy:

1. In the Sun StorEdge EBS Administrator program, select Policies from the Customize menu.
2. Select the appropriate policy.
3. Click Delete.
4. Click OK 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.

Browse and Retention Policies for Manual Backups

Browse and retention policies for manual backups adopt the same policies as the instance of a particular client name that has the longest retention time. For example, if you have three different instances of client *mars*, each with the one of the following retention policy periods:

- One week
- One month
- One year

A manual backup of *mars* without additional modifications automatically adopts a retention policy of one year, and a corresponding browse policy assigned to the client instance (one with a one-year retention policy).

You can modify the browse and retention policies of manual backups, but only when performing the backup using the `save` command.

- For more information about modifying browse and retention policies, see [“To Modify Manual Backup Policies” on page 134](#).
- For more information on manual backups in general, see [“Manual Data Backups” on page 285](#).

Browse and Retention Save Set Policy Modification

You can modify browse and retention policies prior to a save set backup by editing the Browse Policy and Retention Policy attributes in the Client resource. Or, use the save program at the command prompt to specify policy changes of manual backups.

You can use the `nsrmm` program to modify browse and retention policies inherited by the save set records once the backup has occurred.

The `nsrmm` program updates browse and retention policies by using these options:

- `-e retention_time` – updates retention time
- `-w browse_time` – updates browse time

These options (`-e` and `-w`) should be used together with the `nsrmm` option `-S` (to specify save set ID).

Note – The browse time cannot be after retention time, the retention time cannot be before the browse time, and the browse time must be after the *insertion time*. The insertion time is the time this save set record was most recently introduced into the save set database. When the `-e` and `-w` options are used with `nsrmm`, a validity check on these options fails if any of the following are false:

- `browse-time > insert-time`, if `-w` option was specified
- `retention-time > insert-time`, if `-e` option was specified
- `retention-time ≥ browse-time`

Both the browse time (specified with `-w` option) and the retention time (specified with the `-e` option) must be stated in time and date formats acceptable to the `nsr_getdate` program, which converts them to formats acceptable to the Sun StorEdge EBS software. For details about changing browse and retention time, or about time formats acceptable for use, refer to `nsrmm` or `nsr_getdate` in the *Legato Command Reference Guide*.

Example: Changing Browse and Retention Policies with nsrmm

These are examples of how one might use nsrmm to change browse and retention policies:

1. Change the retention time to midnight, January 1, 2007; change the browse time to midnight, January 1, 2004.

```
nsrmm -S 3315861249 -e "01/01/07 23:59:59"  
-w "01/01/04 23:59:59"
```

2. Change the browse time to six months from the current date and time.

```
nsrmm -S 5315861249 -w "6 months"
```

3. Change the retention time to two years from the current date and time.

```
nsrmm -S 3315861249 -e "2 years"
```

▼ To Modify Manual Backup Policies

You can modify the browse and retention policies of a manual backup only when performing the backup using the **save** program at the command prompt. You cannot modify manual backups using the nwbackup graphical interface program.

Use the **save** program with the following options in the same manner as modifying scheduled backup policies by using the nsrmm program:

- **-y** *retention_time* to update retention time
- **-w** *browse_time* to update browse time

For example, to perform a manual backup of *file_name* and provide a retention time of one month, you would enter the following command:

```
save -y "1 month" file_name
```

The nsr_getdate man page provides the terms and syntax to use when specifying browse and retention times. For more information, refer to the nsr_getdate and save man pages or the *Legato Command Reference Guide*.

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. Each line of the report displays the following items:

- Save set creation date
- Stored browse and retention dates (the term *undef* is displayed when connecting to a server using a release of Sun StorEdge EBS software earlier than release 6.0)
- Save set ID
- Client name
- Save set name

Granularity

Prior to the 6.0 release of Sun StorEdge EBS software, browse and retention policies provided only one form of binding, and that was through the Client resource. This was limiting, given that one policy affected all backups for a Client resource definition for past and future backups. This was also true about modifying policies; it affected all past and future backups.

The Sun StorEdge EBS server has the ability to modify browse and retention policies for save sets before and after backups. For example, a user may extend the save sets browse or retention value before the original time limit expires. The recycle functionality of `nsr` no longer needs to consult the Client resource. Instead, it uses save set browse and retention policies from the save set itself.

The Client resource policy settings provide the initial values for the save set record during backup. Later, these settings can be changed and another backup for the same Client resource can result in save sets having the new policy values. Different save sets for the same Client resource can have different browse and retention policies depending on the values they inherited from the Client resource during backup.

Overriding Browse and Retention Policies for Save Sets

Save sets are retained on volumes and in the media database until the save sets expire. Ordinarily, a save set expires and is recyclable when the save set, and all save sets that depend on it for recovery, exceed their browse and retention policies.

However, you can explicitly specify an expiration date for a save set that overrides the browse and retention policies, using the `-e` option with the `save` command. Dependency rules still apply; this means that a save set is not marked as recyclable until all save sets that depend on it are also marked as recyclable.

Backup Command Customization

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

- Shuts down a mail server or database *before* the Sun StorEdge EBS server performs a backup operation.
- Restarts the mail server or database *after* the backup has completed.
- Prints a message (such as, "Backup started at 3:33 a.m.") before the backup operation begins.
- Executes the backup.
- Prints a message (such as "Backup completed at 6:30 a.m.") when the backup has completed.

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, and then enter the name of that script in the Backup command text box of the client computer's Client dialog box. When the client is backed up, the customized program is invoked instead of the standard `save` program. When a client is backed up using `save` in a customized script, the instructions in the backup program are run separately for each save set backed up for the client. See [“Using the save Command with a Customized Backup Program” on page 137](#).
- By entering `savenpc` in the Backup command text box of the client computer's Client dialog box. This way, the client backup invokes the `savenpc` program instead of the `save` program. The first time the client is backed up, `savenpc` creates a default backup program file, which you can then customize for future backups of the client. When a client is backed up using `savenpc`, the instructions in the backup program are invoked just once during the backup session. See [“Using the savenpc Command with a Customized Backup Program” on page 140](#).

Consider the following questions to determine what level of customization will work best:

- Is there client data that does *not* need to be backed up every time (for example, company e-mail)?
- Do you want to have 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

You can enter the name of a customized script in the Backup Command attribute in the Client resource that includes additional processing instructions. If an entry exists in the Backup Command attribute, the script associated with it is executed, instead of the default **save** program, when scheduled backups are initiated.



Caution – Keep in mind that, 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 filesystem 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 you create a customized backup program for a client, immediately try backing up the client. Any configuration or network problems that might prevent a successful backup should show up during this test.

The syntax you use to create the backup program or batch file must adhere to the criteria described in the following list. The list is detailed and includes programming information. Do not attempt to write your own backup command unless you can follow these recommendations:

- 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.

▼ 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.

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

The following is an example script that backs up a ClearCase VOB (version object base). This customized 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). You 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.



Caution – You must include the **save** command within your script in order for the backup to occur. If the script file is not in the same directory as the **save** program, 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 which save reports on 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 the backup status in the 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>&

# make sure to gracefully exit out of this shell script
exit 0

```

Using the savenpc Command with a Customized Backup Program

By using the `savenpc` program instead of the `save` program, you can run preprocessing and postprocessing commands on a client that 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 `savenpc` command are identical to those for the `save` command. For further details about `savenpc` command options, refer to the `savenpc` man page or the *Legato Command Reference Guide*.

To execute the program:

1. In the Sun StorEdge EBS Administrator program, select Client Setup from the Clients menu.
2. Select the appropriate client.
3. Enter `savenpc` in the Backup Command attribute.
4. Run a backup of the client.

The first time a backup group with a client that uses `savepnpc` runs, a standardized `<group-name>.res` file is created in the `/nsr/res` (UNIX) or `<Networker_install_path>\res` (Microsoft Windows) directory on the client (where `group-name` is the same as the name in the group resource selected for that client). If the client belongs to multiple backup groups, a separate `<group-name>.res` file is created for each group to which the client belongs.

The initial `<group-name>.res` file contains 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` option in the `<group-name>.res` file determines what will happen to the preprocessing command when the `savegroup` process aborts prematurely. By default, the preprocessing command will execute as its own process and will therefore not be killed if the `savegroup` process aborts prematurely. If you want the preprocessing command to be killed 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 operation using the options you 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`. You must build the environment 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.
- If you do not want to include the environment variables in the `<group-name>.res` file, then you must include full pathnames for all commands and files.

- Resident commands, for which there is no executable file present, like `time` and `dir`, will not work as commands in your `<group-name>.res` file. The log reports that the executable file could not be found.
- On a Microsoft Windows client, do not use "@ECHO OFF" in the `<group-name>.res` file.
- To add more than one command sequence to the `precmd` and `postcmd` attributes, insert a comma (,) to separate the commands.
- A complete command line for an attribute must end with a semicolon (;).
- You must escape any backslash (\) characters in the `<group-name>.res` file. Therefore, a pathname like `c:\mydir\myprogram.exe` should 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";
postcmd: "V:\\usr\\sap\\PDB\\SYS\\exe\\run\\PDB-start.cmd > C:\\
WINNT\\system32\\PDBStart.log 2>&1";
timeout: "12:00pm";
```

It is not necessary to escape any backslash characters in scripts called in the `<group-name>.res`. One way to simplify the pathname issue is to include all of your 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 line in the `<group-name>.res` file. Note that in the preceding example, the comment "#end" has been added to the end of the file, ensuring that a line break follows the last semicolon.
- Text written to standard output by the preprocessing appears in the Sun StorEdge EBS completion notices. If this output is verbose, you can direct the output from the commands in your script to a log file for later examination. Output from the postprocessing is discarded because the connection to the server has already been dropped when these commands are run. Consider redirecting their output to a log file so you can examine the result if something failed to restart.

Timeout Attribute

The `Timeout` attribute indicates a point in time 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 a format that `nsr_getdate` can understand and must be enclosed in double quotes. For details, read the description of `nsr_getdate` in the *Legato Command Reference Guide*.

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. If you do not plan to use it, you can make the line on which it appears into a comment with a # character and that line of instructions will not be executed, for example:

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

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

▼ To Customize the `savenpc` Command for Multiple Groups

To customize pre- and postcommand processing for multiple groups:

1. Copy existing `<group-name>.res` files.
 - Microsoft Windows clients:
`<NetWorker_install_path>\tmp\<group-name>.res` to
`<NetWorker_install_path>\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, then you can activate the group for `savenpc` without the presence of this file. A default template will be created for you at `/nsr/res/<your_new_group>.res` or `<NetWorker_install_path>\tmp\<your_new_group>.res`, which can then be customized.

Message Logging by the `savenpc` Command

Messages generated by `savenpc` are written to the `savenpc.log` file, located in the following location on the client's filesystem:

- UNIX: `/nsr/logs`
- Microsoft Windows: `<NetWorker_install_path>\logs`

The format of the `savenpc.log` file is similar to the following:

```
09/03/99 13:56:43 preclntsave: All command(s) ran successfully.

09/03/99 13:57:43 preclntsave: All savesets on the worklist are
done.

09/03/99 13:57:51 pstclntsave: All command(s) ran successfully.
```

Directives

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 process to your specific needs, maximize the efficiency of backups, and apply special handling to individual files or directories.

Local and Global Directives

There are two 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.
- *Local directive files* — Users can create local directive files named `.nsr` anywhere on a client filesystem that they have permission to create files. The directives these files contain apply only to the immediate data within the path where the directive file is located.

Note – In the event of conflicting directives, global directives are enforced over local directives.

All Sun StorEdge EBS directives use essentially the same format, however, global directives contain an additional `<<directory>>` line, which specifies the path of the data the directive acts upon. For example, to force scheduled backups to skip all `*.tmp` files in the `/mydir` directory on a particular client, you could create a directive by using either of the following methods:

- Using the Sun StorEdge EBS Administrator program, create a global directive in the following format and apply it to the appropriate client:

```
<</mydir>>
skip: *.tmp
```

For instructions on creating global directives, see [“To Create a Directive Resource” on page 153](#).

- Using a text editor, on the Sun StorEdge EBS client create a file named `/mydir/.nsr` that contains the following directive:

```
skip: *.tmp
```

For more information about directive formats, see [“Directive Format” on page 149](#).

Application Specific Modules

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 included in local directive files or server-side Directive resources, and are then executed during backup of client data. Directive files and Directive resources can contain one or more ASM. [TABLE 4-5](#) describes the Sun StorEdge EBS ASMs.

TABLE 4-5 Sun StorEdge EBS ASMs (1 of 3)

ASM Name	Description	Example
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 4-5 Sun StorEdge EBS ASMs (2 of 3)

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
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.

TABLE 4-5 Sun StorEdge EBS ASMs (3 of 3)

ASM Name	Description	Example
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 using an encrypted password. For more information, see “Password ASMs” on page 148 . Note: Available only on Microsoft Windows and NetWare versions of Sun StorEdge EBS software.	pw1: *.* password "-o{koo{o"
pw2	The same as the pw1 ASM, except that the file itself is also encrypted. For more information, see “Password ASMs” on page 148 . Note: Available only on Microsoft Windows and NetWare versions of Sun StorEdge EBS software.	pw2: *.* password "-o{koo{{"
rawasm	Used to back up /dev entries (that is, block and character-special files) and their associated raw disk partition data. On some systems, /dev entries are actually symbolic links to device specific names. Unlike other ASMs, this ASM follows symlinks, allowing the shorter /dev name to be configured. For more information, including important precautions when using this ASM, see “Important Precautions When Using rawasm” on page 148 .	rawasm: /dev/oracle1
skip	Omits the specified files and directories from the backup, and does not place the filename in the online index.	+skip: *.tmp
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

The password ASMs, available for Microsoft Windows and NetWare clients only, are used to password-protect (pw1) or encrypt (pw2) data. Directives that use these ASMs require a local file named `Networkr.cfg`, which the Sun StorEdge EBS User program creates when you set a password. For more information about entering a password in the Sun StorEdge EBS User program, refer to the *Legato NetWorker Administrator's Guide, Microsoft Windows Version*.

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"
```

When writing a directive that uses pw1 or pw2, include the encrypted password by copying it from the `Networkr.cfg` file.

Important Precautions When Using rawasm

The `rawasm` directive is used to back up raw disk partitions. Raw partitions are often used to store active database management system (DBMS) data. If your raw partition contains data managed and updated by an active DBMS product, `rawasm` alone will not give a consistent backup. The database must not be updating the data in an uncontrolled fashion while `rawasm` saves or recovers data on the partition. The partition must be offline, the database manager shutdown, or the partition placed in an appropriate state for backup. As a result, it is better to use a Legato NetWorker Module application to back up databases.

Similarly, if `rawasm` is used to save a partition containing a UNIX filesystem, the filesystem must be unmounted or mounted read-only to obtain a consistent backup.

Using rawasm with Recovery

When recovering, `rawasm` requires that the filesystem 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 which 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, only the amount of data originally saved will be recovered, and since rawasm does not discover the size completed, the estimated size reported on recovery is not accurate.

Directive Format

Both global and local directives use a format that contains three primary types of directive 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(5)` and `nsr_directive(5)` man pages.

Directory Specifications

Directory specifications specify the highest-level directory to which the directives following this entry apply. The format for a directory specification is:

```
<<directory>>
```

For example, to apply ASMs to the `/etc` directory, use the following format:

```
<</etc>>
    directive
```

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

```
<</etc>>
    directive
<</tmp>>
    directive
```

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.

Note – For directives applied to clients on Microsoft 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:

```
<<"C:\MYDIR">>
```

Application Specific Module Specifications

ASM specification directives specify an ASM to run, as well as parameters associated with the ASM. The syntax for ASM specifications is:

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

Where the optional plus sign (+) indicates that the directive applies to both the current directory and all subdirectories, *asm* is the name of the ASM to be executed, and *argument* is a list of filenames that are acted upon by the ASM. The *argument* can include multiple filenames, separated by spaces, and can use wildcards as well. Only filenames and wildcards are allowed; no directory information can be included. For a description of available ASMs, see [“Application Specific Modules” on page 145](#).

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 filematching patterns are supported.

On UNIX systems, when applying a directive to all files, including hidden files (. * files), use * .?* (note that there must be a space between the first asterisk and the dot).

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 4-6 on page 151](#) lists the three save environment keywords.

TABLE 4-6 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 G:\SRC, except those in the G:\SRC\SYS directory, enter: <<G:\SRC>> +skip: *.o <<G:\SRC\SYS>> forget This uses the skip ASM to instruct the Sun StorEdge EBS server to skip all files named *.o in the G:\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: <<G:\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 G:\HOME\DOC directory to be applied, enter: <<G:\HOME>> ignore <<G:\HOME\DOC>> allow

Preconfigured Directive Resources

The Sun StorEdge EBS software comes with a number of preconfigured Directive resources. Some of these Directive resources contain preset global directives; others do not, but can be modified as needed.

All preconfigured Directive resources can be modified, but they cannot be deleted.

TABLE 4-7 lists the preconfigured directives and their descriptions.

TABLE 4-7 Preconfigured Directives

Directive Resource	Description
DOS standard	Legacy resource used to back up Microsoft Windows 95 and Windows 98 clients. By default, this resource has no directives.
DOS with compression	Legacy resource used to back up and compress Microsoft Windows 95 and Windows 98 clients. Applies the compressasm ASM to all files.
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 filesystem. • 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 e-mail 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 resource, along with applying the compressasm ASM to all files.

▼ To Create a Directive Resource

For instructions on the syntax to use when entering directives, see [“Directive Format” on page 149](#).

To create a Directive resource:

1. In the Sun StorEdge EBS Administrator program, select Directives from the Customize menu.
2. Click Create.
3. In the Name attribute, enter a name for the new directive.
4. In the Directive attribute, enter one or more directives.
5. Click Apply.

The Directive resource is now available to be selected in the Directive attribute for Client resources.

▼ To Edit a Directive Resource

Note – You cannot change the name of a directive.

To edit a Directive resource:

1. In the Sun StorEdge EBS Administrator program, select Directives from the Customize menu.
2. From the Directives list box, select the appropriate directive.
3. Modify the directives as necessary, and click Apply.

▼ To Delete a Directive Resource

Note – You cannot delete preconfigured directives or any directives currently applied to a Client resource.

To delete a Directive resource:

1. In the Sun StorEdge EBS Administrator program, select Directives from the Customize menu.

2. Select the appropriate directive in the Directives list box and click Delete.
3. Click OK to confirm the deletion.

Adding or Changing the Sun StorEdge EBS Servers for a Client

The Sun StorEdge EBS software uses the contents of the `/nsr/res/servers` (UNIX), or the `<NetWorker_install_path>\res\servers` (Microsoft Windows) file 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
- Migration
- Another client requesting a directed recover

During Sun StorEdge EBS installation, you have can add the names of Sun StorEdge EBS servers to this file. To add additional servers at a later date, a text editor must be used to add the server 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.



Caution – If the `servers` file is empty, then any Sun StorEdge EBS server has tasking rights This is a potential security concern.

▼ To Edit the servers File

To give clients tasking rights by adding their names to the `servers` file:

1. Shut down the `nsrexecd` daemon.
2. Open the `/nsr/res/servers` file in a text editor.

The default installation location for this file is:

- `nsr/res/servers` (UNIX)
- `<NetWorker_install_path>\res\servers` (Microsoft Windows)

3. Enter one server name per line.
4. Save the changes and exit the text editor.
5. Restart the nsrexecd daemon.

Client Priority

The Client Priority attribute in the Client resource specifies the order in which participating clients are probed for the information needed to complete the save set worklist for that client. The Client Priority attribute is part of the Client resource, and is only visible when Hidden Attributes is enabled. This 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 Client 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 contact order is random.

While the Client 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 list of clients to be contacted.

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 [“To Create a Group” on page 254](#).

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.

Logical Volume Backup

A *logical volume* is a type of primary (disk) storage on a client computer that can span several physical disk volumes. The logical volume has its own device address, and it is treated similarly to a disk partition by the filesystem. When the Sun StorEdge EBS server backs up data from clients, it has to determine how many save sessions to allocate to each client for best performance. To avoid contention, there should not be more than one backup operation running per physical disk. Different sessions are allocated across different physical disks for this reason.

To determine how many save sessions to allocate, the Sun StorEdge EBS server probes (queries) the clients in a backup group using the `savefs -p` command. This command determines:

- What data to back up.
- Where the data is physically located.
- Whether there are logical volumes.

The Sun StorEdge EBS server stores this information in two variables, *disk-number* and *maximum-sessions*, according to the following rules:

- When the group of volumes or disks that contain logical volumes is not part of the device path, all logical volumes on the client computer are assigned to the same *disk-number*, and *maximum-sessions* is set to the number of logical volumes on the client computer.
- When the group of volumes or disks that contain logical volumes is part of the device path, all logical volumes within the volume group are assigned to the same *disk-number*, and *maximum-sessions* is set to the number of logical volumes within the volume group.

The server uses the output from the `savefs` probe to allocate its save sessions (up to the maximum server parallelism) across the clients in the backup group.

It allocates the save sessions as follows:

1. The server allocates one save session per client in the backup group.
2. If there are still save sessions available, it allocates one save session per physical disk on each client.
3. If there are still save sessions available, it allocates save sessions to each *disk-number* value, up to the limits in *maximum-sessions* for each client and client parallelism.

Sun StorEdge EBS Client Licensing

When a backup begins, a Sun StorEdge EBS client passes its attributes to the Sun StorEdge EBS server. The Sun StorEdge EBS server uses this information to verify that the client is allowed to back up to it. The following read-only attributes in the Client resource are used for enforcing client licensing:

- Client OS Type
- Sun StorEdge EBS Version
- CPUs
- Enabler in Use

When you enable a Sun StorEdge EBS server, the enabler also allows you to use a specific number of clients of the same platform type. For example, a Sun StorEdge EBS server on Solaris is licensed for a specific number of Solaris or SunOS clients. If you want to connect clients of other platforms to that Sun StorEdge EBS server, you must purchase a Legato NetWorker ClientPak[®] enabler for each of those platforms.

You can define more clients than the number of clients allowed by all enabler codes stored in the server. However, the Sun StorEdge EBS server enforces the number and type of licensed clients at backup time. It does this by verifying:

- The number of connected clients is less than or equal to the number of clients allowed by its enabler codes.
- The clients are of the types allowed by the enabler codes.

Client connection licenses are not tied to a particular operating system. Except for NDMP and cluster clients, a connection license can be used for any client operating system that is either the same as the server or has its compatible ClientPak enabler installed. NDMP and cluster clients require dedicated connection licenses.

The nsrlic Program

The `nsrlic` program generates reports about all licenses currently active on a Sun StorEdge EBS server. The program shows a breakdown of the types of workstations and servers that are in use. In addition, `nsrlic` has a `-v` option, which provides a list of all connected clients, as well as additional license information. For information about the `nsrlic` command, refer to the *Legato Command Reference Guide*.

The reports generated by the `nsrlic` program still separate client licenses into universal and workstation licenses:

- All computers with operating systems that could be a server (for example, Solaris, AIX or Microsoft Windows NT) are listed under universal.

- All others (for example, Microsoft Windows 2000 or Microsoft Windows XP Professional) are listed as workstations.

Workstation licenses will be “borrowed” from the universal licenses, unless the server has old workstation client connection licenses.

Client Licensing Error Messages

In the following cases, Sun StorEdge EBS backups will be rejected because of licensing enforcement, and an error message will be generated:

- Backups will be rejected if the client type is not allowed by the enabler code, or if a Legato NetWorker ClientPak enabler is not present. For example, the following error message might be sent for a Solaris server with an unlicensed Microsoft Windows NT client:

RAP error, ClientPak for NT does not exist in the system.

To remedy, you must purchase and enter a Legato NetWorker ClientPak for NT enabler.

- Backups will be rejected if no more licenses are available for that type of client. For example, you might receive an error message similar to this:

RAP error, Too many clients. Maximum is 2

To remedy, you must purchase and enter an enabler for additional clients.

For more information about purchasing enablers, refer to your Installation Guide.

Legato NetWorker ClientPak for Mac OS X

This section describes Legato NetWorker client support for Mac OS X computers.

Mac Legacy Files Supported

The Legato NetWorker client on Mac OS X supports backup and recovery of Hierarchical File System (HFS) and HFS+ resource fork and Finder metadata. This compatibility ensures the proper backup and restore of legacy files used with Mac OS X.

Supported Filesystems

The Sun StorEdge EBS client for Mac OS X software supports the following filesystems:

- HFS+ (including journaled)
- HFS
- UFS

Known Limitations

The following are known limitations for this release of the Legato NetWorker ClientPak:

- There is no graphical user interface in this release, however there is full support of the command line interface.
- Recovery of Mac OS X save sets to non-Mac OS X clients is not supported.
- Non-English locales are not supported.

Backing Up a Sun StorEdge EBS Client on Mac OS X

This section provides information on configuring and performing backups of data on the Legato NetWorker client on Mac OS X.

Performing Backups of Data on Mac OS X Clients

For both scheduled and manual Sun StorEdge EBS backups, each Mac OS X computer to be backed up must have a Client resource on the Sun StorEdge EBS server.

Scheduled Backups

Scheduled backups of Legato NetWorker clients on Mac OS X are performed by the Sun StorEdge EBS server. For information about scheduling backups, see [Chapter 6](#).

Manual Backups

To manually back up files or directories of Sun StorEdge EBS Legato NetWorker clients on Mac OS X, use the **save** command, 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 networker_server` option.

MAC OS X Required Directives

To ensure a consistent state after recovery, certain files and directories should *not* be backed up on Mac OS X systems. [TABLE 4-8](#) lists files and directories that should *not* be backed up on Mac OS X.

TABLE 4-8 Mac OS X Files or Directories to Exclude from Backup

Files or Directories	Description
Desktop DB Desktop DF	Mac OS 9 desktop database files
.DS_Store	Finder settings
TheVolumeSettingsFolder/	Desktop state details
VM Storage	Mac OS 9 virtual memory
private/var/db/netinfo	The NetInfo database

To ensure that these files and directories are not backed up, define them on the Sun StorEdge EBS server in a Directives resource. For an example of a Directive that skips these files, see [“Example: Standard Mac OS X Directives” on page 161](#). For more information on Directives, see [“Directives” on page 144](#).

Example: Standard Mac OS X Directives

```
<< "/" >>
    skip: Desktop\ DB
    skip: Desktop\ DF
    +skip: .DS_Store
    skip: cores
    skip: VM_Storage
    skip: TheVolumeSettingsFolder
<< "/private/var/db/netinfo" >>
    skip: .
<< "/nsr/logs" >>
    logasm: .
<< "/private/var" >>
    logasm: .
<< "/private/var/vm" >>
    swapasm: .
```

Backing Up the NetInfo Database for Disaster Recovery

The following sections describe issues that are specific to the backup and recovery of the Mac OS X NetInfo database. The NetInfo database contains system configuration information that is essential for disaster recovery.

Backing Up the NetInfo Database Before a Save

The `save` program does not back up the resident NetInfo database. To ensure consistency of disaster recovery, save the NetInfo database in the `/var/backups/networker.nidump` file before each Sun StorEdge EBS backup. This can be done manually with the following command:

```
# nidump -r / -t localhost/local /var/backups/networker.nidump
```

To automatically generate a backup file of the NetInfo database before each Sun StorEdge EBS save, use the Sun StorEdge EBS `savenpc` command, configured as follows:

1. Enter `savenpc` as the Backup command when configuring the Mac OS X client as a Legato NetWorker Client resource in the Sun StorEdge EBS Administrator or `nsradmin` program.
2. Create a custom `savenpc` script in the `/nsr/res` directory with the name `<group_name>.res`, where `<group_name>` is the Group resource selected for that client.

Example: Custom `savenpc` Script for Mac OS X

A Mac OS X client that belongs to the Default group will have a `/nsr/res/Default.res` script with the following content:

```
type: savenpc;  
  
precmd: "/usr/bin/nidump -r / -t localhost/local  
         /var/backups/networker.nidump"
```

In this script, the `savenpc` command backs up the NetInfo database to the `/var/backups/networker.nidump` file before each scheduled save.

For additional information on the `savenpc` script, see [“Using the `savenpc` Command with a Customized Backup Program” on page 140](#). For additional information on recovering from a disaster, refer to the *Legato NetWorker Disaster Recovery Guide*.

Restarting the Sun StorEdge EBS Client on Mac OS X

This section describes how to shut down and restart the Legato NetWorker client daemon on Mac OS X.

▼ To Shut Down the Sun StorEdge EBS Client Daemon on Mac OS X

To shut down the Legato NetWorker client daemon on Mac OS X:

1. Open the Mac OS X Terminal application utility.
2. Enter the following command to shut down the Sun StorEdge EBS daemons:

```
# SystemStarter stop NetWorker
```

3. Verify that the Legato NetWorker client daemon, `nsrexecd`, is not running by entering the following command at the system prompt:

```
# ps -ax | grep nsrexecd
```

▼ To Start the Sun StorEdge EBS Client Daemon on Mac OS X

To start the Legato NetWorker client daemon on Mac OS X:

1. Open the Mac OS X Terminal application utility.
2. Start the Legato NetWorker daemon by entering the following command:

```
# SystemStarter start NetWorker
```

3. Verify that the Legato NetWorker client daemon, `nsrexecd`, is running by entering the following command at the system prompt:

```
# ps -ax | grep nsrexecd
```

Recovering Individual Files and Directories

To browse and recover files with Sun StorEdge EBS software, use the `recover` program. For information about the `recover` program, see [“Recovering Data from the Command Prompt” on page 290](#).

Device and Media Management

This chapter describes device and media operations you can perform through the Sun StorEdge Enterprise Backup Software server.

The main sections in this chapter are:

- [“Pools” on page 166](#)
- [“Volume Labeling” on page 180](#)
- [“Devices” on page 187](#)
- [“Dynamic Drive Sharing” on page 198](#)
- [“Sun StorEdge EBS DiskBackup” on page 206](#)
- [“Media Handling Error Resolution” on page 225](#)
- [“Auto Media Management” on page 226](#)
- [“Storage Nodes” on page 228](#)
- [“Storage Management Operations” on page 231](#)
- [“Volume Status” on page 234](#)
- [“Sun StorEdge EBS Volume Operations” on page 236](#)
- [“Save Set Staging” on page 244](#)

Note – The term “autochanger” refers to a variety of backup devices: autoloader, carousel, datawheel, jukebox, library, and near-line storage.

Pools

A *pool* is a specific collection of media to which the Sun StorEdge EBS server writes data. The Sun StorEdge EBS server uses pools to sort and store data. The configuration settings for each pool act as filters that tell the server which volumes should receive specific data. Pools are used in conjunction with label templates to keep track of which data is on which specific volume. For detailed information about label templates, see [“Volume Labeling” on page 180](#).

This section covers the following topics:

- [“How the Sun StorEdge EBS Server Uses Pools” on page 166](#)
- [“Sun StorEdge EBS Pool Types” on page 167](#)
- [“How the Sun StorEdge EBS Server Uses Pool Criteria to Sort Data” on page 169](#)
- [“Directing Client Indexes and Bootstrap to a Separate Pool” on page 170](#)
- [“Meeting the Criteria for More Than One Pool Configuration” on page 171](#)
- [“Not Meeting the Criteria for Any Pool” on page 172](#)
- [“Configuring Pools” on page 173](#)
- [“Using Storage Devices and Pool Configuration to Sort Data” on page 175](#)
- [“To Create a Pool” on page 176](#)
- [“To Edit a Pool” on page 179](#)
- [“To Delete a Pool” on page 179](#)
- [“To Create an Archive Pool” on page 180](#)

How the Sun StorEdge EBS Server Uses Pools

The way you configure pools determines which volumes receive data. Each pool configuration contains criteria that the data must meet for the data to be written to associated volumes.

When a scheduled backup occurs, the Sun StorEdge EBS server tries to match the save set to a pool configuration. If the save set matches the criteria of a pool configuration, it is directed to a labeled volume from that pool. The server then checks if a correctly labeled volume for that pool is mounted on a storage device.

- If it is, the Sun StorEdge EBS server writes data to the volume.
- If it is not, the Sun StorEdge EBS server requests that such a volume be mounted and waits until an operator mounts the appropriate volume.

Sun StorEdge EBS Pool Types

The Sun StorEdge EBS software provides preconfigured pool types to keep different types of data separate. Sun StorEdge EBS software does not mix the following types of data on volumes within a pool:

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

Unless you specify other pools, all backup data is routed to the Default pool and all archive data is routed to the Indexed Archive pool. Cloned backup data is routed to the Default Clone pool, and cloned archive data is routed to the Indexed Archive Clone pool.

Preconfigured Pools

[TABLE 5-1](#) lists the preconfigured pools that come with the Sun StorEdge EBS product.

TABLE 5-1 Preconfigured Pools (1 of 2)

Type of Pool	Type of Data
Default	All backed up data when another pool is not specified.
Default Clone	Data from clones of the Default pool's volume.
Archive	Data from Archive Pool clients. This pool does not store data in the client file index.
Archive Clone	Data from clones of the Archive pool's volume.
Full	Data from full backups.
Indexed Archive	By default, this pool is used for data being archived. Unlike the Archive Pool, this pool stores entries in the client file index and will increase the size of the index.
Indexed Archive Clone	Data from clones of the Indexed Archive pool's volume.
Migration	Migrated data.
Migration Clone	Data from clones of the migration pool's volume.
NonFull	Data from incremental and other non-full backups.
Offsite	Data to be stored offsite.

TABLE 5-1 Preconfigured Pools (2 of 2)

Type of Pool	Type of Data
PC Archive	Data from PC Archive clients. Like the Indexed Archive pool, this pool puts entries in the client file index.
PC Archive Clone	Data from cloned PC Archive clients.

Default Pool

If the save stream does *not* match the criteria for any custom (user-created) pools, the Sun StorEdge EBS server directs the save stream to the *Default* pool. Even when you use customized pools or other preconfigured pools, always ensure that appropriate storage devices and media are available for the Default pool for those cases when the Sun StorEdge EBS server directs save streams there. If you do not have the Default pools properly configured for backup, you may receive an error message like the following in the Pending window:

```
media waiting (or critical): backup to pool 'Default' waiting for
1 writable backup tape or disk
```

Example: Using Pool Configurations to Sort Data

[FIGURE 5-1 on page 169](#) illustrates how the Sun StorEdge EBS server uses pool configurations to sort data. The save stream contains data from a full backup that was performed on client save sets in a group called *Accounting*. The Sun StorEdge EBS server looks for a pool configuration that matches the group *Accounting* and the level *Full*. When the server finds the matching pool configuration, it writes the data to a volume with a label from the corresponding *Accounting Full* pool of volumes that is mounted on one of the storage devices.




FIGURE 5-1 Using Pool Configurations to Sort Data

How the Sun StorEdge EBS Server Uses Pool Criteria to Sort Data

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

- Group (backup group)
- Sun StorEdge EBS client
- Save sets (file or filesystems)
- Backup levels (full, levels [1–9], incremental, manual)

If you begin by entering a group name in the Group attribute, the pool is immediately restricted to accept only data associated with the named group. If you add a second group name to the Group attribute, the pool accepts data associated with either group, but no others. Entries for a single attribute function as “OR” clauses; that is, the pool accepts data from clients in either group.

Each of the four configuration criteria, however, functions with the others as an “AND” clause. That is, 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 pool.

You cannot create pools that share identical settings for *all* of the following:

- Pool type
- Group
- Client
- Save set
- Level

If the settings for a new pool match the settings for an existing pool, you receive a warning message. Change the appropriate settings and reapply to save the Pool resource.

Directing Client Indexes and Bootstrap to a Separate Pool

You can use regular expression matching to direct the client indexes and bootstrap to a pool different 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 filesystem to the same pool, create a pool (in the Pool resource) with the attributes listed in

[TABLE 5-2](#):

TABLE 5-2 Example Attributes for Sending Bootstrap to a Separate 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 pools, while the Sun StorEdge EBS server’s bootstrap and index save sets are written to a separate volume labeled for the *Index* pool.

Meeting the Criteria for More Than One Pool Configuration

Depending on the pool configurations you create, you might have data that matches the criteria for more than one pool configuration. For example, if you configure one pool to accept data from a group called *Accounting* and you configure another pool to accept data from all full backups, the Sun StorEdge EBS server must determine the pool to which a full backup for the Accounting group is written. The server uses the following pool selection criteria:

- 1. Group (highest precedence)
- 2. Client
- 3. Save set
- 4. Level (lowest precedence)

When data matches the attributes for two pools, for example, Group and Level, the pool data is written to the pool specified in the Group attribute. For example, in the case where the data from the group matched the criteria for two different pools (one configured to accept data from the Accounting group and the other to accept data from all full backups) the data is routed to the pool that accepts data from the Accounting group.

TABLE 5-3 details the hierarchy that the Sun StorEdge EBS server uses to determine pool selection when a conflict arises. For example, the pool criteria for Group takes precedence over the pool criteria for Client, Save set, and Level. Data that meets the criteria for both pools is written to the pool associated with the Group. If data does not meet the criteria for any customized group, it is written to the Default group.

TABLE 5-3 Sun StorEdge EBS Hierarchy for Resolving Pool Conflicts (1 of 2)

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

TABLE 5-3 Sun StorEdge EBS Hierarchy for Resolving Pool Conflicts (2 of 2)

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

Not Meeting the Criteria for Any Pool

When you use customized pool configurations to sort data, you might inadvertently omit a client or save set. During a scheduled backup, if data does not meet the criteria for any customized pool configuration, the Sun StorEdge EBS server automatically sends the data to the Default pool. The server uses the Default pool to ensure that all data for clients in a backup group is backed up to a volume.

When the Sun StorEdge EBS server sends data to the Default pool, it looks for a labeled volume from the Default pool mounted on a storage device. If no Default 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 pool volume in the middle of a scheduled backup, but an operator is not present to mount it, the backup pauses until it has been mounted. If you have an operator available to monitor the backups, it is a good idea to keep a volume labeled for the Default pool close at hand in case this situation arises unexpectedly.

If you plan to use Sun StorEdge EBS software for unattended backups, run a test of the backup after making any configuration changes to ensure that all data is written to the appropriate volumes and to avoid an unexpected Sun StorEdge EBS request for a Default pool volume. For the procedure to test a scheduled backup, see [“To Start a Group Immediately” on page 258](#).

Configuring Pools

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

- Incremental backups
- Manual backups
- Clone data
- Archive data
- Migration data

Note – With Sun StorEdge EBS release 6.1 and later, you can create and edit pools while a backup session is running.

Incremental Backups

To create a separate pool for incremental backups, be aware that the Sun StorEdge EBS hierarchy of precedence (See [TABLE 5-3 on page 171](#)) If the Level attribute value is *incremental*, incremental data is routed to the associated pool but the corresponding changes to the client's file index are not. The Sun StorEdge EBS server saves all client file indexes at level 9 to speed the recovery operation, if one is needed.

If the client file indexes do not meet the criteria for the pool associated with the incremental backups, the Sun StorEdge EBS server:

- Matches the indexes to another pool (usually the Default pool).
- Looks for an appropriately labeled volume to write to.

If you need to recover all of your data, you need to use a large number of volumes. To store the client file indexes along with the incremental backup data and to speed the recovery operation, define the Level attribute in the Pool resource to accept both level 9 and incremental data.

You can use the Sun StorEdge EBS preconfigured NonFull pool settings to ensure that the client file indexes belong to the same pool as their incremental backups. When you keep the indexes in the same pool as their incremental backups, you reduce the number of volumes you need for a recovery.

Manual Backups

You can create a customized pool to receive data from a manual backup by specifying *manual* in the Level attribute. The Sun StorEdge EBS server, however, sorts data from a manual backup differently than data from a regularly scheduled

backup. Because a manual backup is not performed as part of a scheduled backup group, the data is not associated with any group name. Thus, when you perform a manual backup in which only a single client's save set data is saved, the group normally associated with that client's save set is not included as a criterion for pool assignment. As a consequence, data from a manual backup may be sent to a different pool than where it would be sent during a scheduled backup.

If you do not create a customized pool to receive data from manual backups, the Sun StorEdge EBS server uses the Default pool and looks for a mounted volume from the Default pool on which to write manually backed-up data. Because the Sun StorEdge EBS server tracks the volume location of all backup data, you do not need to worry about tracking which volume contains the manually backed-up data. If you need to recover the data, the server requests the correct volume.

When you perform a manual backup, the media database and resource files contained within the server bootstrap are not included in the backup. If you never perform regularly scheduled backups of the clients and server computers, the information vital to data recovery in the event of a disaster is not available. For further information on how the bootstrap is used during disaster recovery, refer to the Legato NetWorker *Disaster Recovery Guide*.

Clone Data

If you want to clone data, the Sun StorEdge EBS server requires a specific pool to receive the clone data and a minimum of two devices:

- One to read the source volume
- One to write the clone.

If you do not associate data to be cloned with a customized Clone pool, the Sun StorEdge EBS server automatically uses the Default Clone pool. You must mount an appropriately labeled volume on a separate storage device for the cloning process to proceed smoothly. For more information on the Sun StorEdge EBS cloning feature, see [Chapter 9](#).

Archive Data

If you want to use Sun StorEdge EBS Archive to archive data, the Sun StorEdge EBS software requires a specific pool to receive the archived data. You can then store these volumes off site. If you do not associate data to be archived with a customized Archive pool, the Sun StorEdge EBS server automatically uses the preconfigured Indexed Archive pool. You must mount an appropriately labeled volume on a storage device for the archive process to proceed smoothly. For more information on the Sun StorEdge EBS archive feature, see [Chapter 8](#).

The Indexed Archive pool includes file information in the client file index, and will cause the index to grow in size. In contrast, the Archive pool will not store information in the client file index and does not cause the index to grow. As a result, you cannot browse or recover individual files when the Archive pool is used.

Migration Data

Archive and migration data are both written in a different format than regular Sun StorEdge EBS save set data. Therefore, archive and migration data must each be written to different volumes. The client file indexes created during an indexed archive, premigration, or migration operation are kept online and will not be written to the media until the next savegroup. By default, they are written to a volume from the Default pool every time the savegroup involving those clients is performed.

Using Storage Devices and Pool Configuration to Sort Data

You can sort data by configuring pools of volumes in conjunction with your storage devices. You can either use specific media to receive data, or designate a specific storage device to receive data from a designated 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 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 118](#).

Using a Specific Device for Backup Data

You can associate a pool with a specific storage device. For example, you may want your full backups 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 pool mounted on the specific storage device.
- Associate a specific pool with the storage device in the Device resource. Then, in the Pool resource, associate the pool with the device in the Pool Configuration attribute list. All data will be written only to that device.

▼ To Create a Pool

To create a pool:

1. Create a label template for your pool. You also have the option of using a preconfigured label template. For a complete discussion of label template, see [“How the Sun StorEdge EBS Server Uses Label Templates” on page 180](#).

Note – If you do not select a label template when creating a pool, the Sun StorEdge EBS server notifies you that it will create the label template. To have the Sun StorEdge EBS server create the label template, click OK.

2. From the Media menu, select Pools.
3. Click Create.
4. Enter a name for the pool in the Name attribute.

Because a pool is associated with a label template, use a name that matches the label template created in [Step 1](#). If you are using a preconfigured label template, name the pool in a way that clearly associates the pool with the corresponding label template.
5. From the Enabled attribute, select Yes.
6. From the Pool Type attribute, select the appropriate pool type.
7. In the Label Template attribute, select the matching label template.
8. Select one or more of the following: group, client, save set, or level. For information on the syntax of save set paths, see [“Expression Matching of Save Sets to Pools” on page 177](#).
9. Select the devices that are permitted to mount volumes associated with this pool (optional).
10. In the Store Index Entries attribute, indicate whether client file index entries are to be generated for this pool.
11. In the Auto Media Verify attribute, select whether automated media verification will be performed while data is written to a volume associated with this pool.

With Auto Media Verification enabled, the Sun StorEdge EBS software verifies data written to media. It repositions the volume to read a portion of the data and compares the written data to the original data that is stored in memory. Media is verified whenever a volume becomes full during a backup, or whenever the volume is idle because all save sets being written to the volume are complete.
12. In the Recycle to Other Pools attribute, select whether volumes marked as recyclable can be used by other pools.

13. In the Recycle from Other Pools attribute, select whether this pool can recycle volumes from other pools.

14. (Optional) Select a preference for volume type in the Volume Type Preference attribute.

The value you state is used as a selection factor when a request is made for a writable volume.

15. Complete other attributes as necessary, and click Apply.

You cannot create pools that share identical settings for *all* of the following:

- Pool type
- Group
- Client
- Save set
- Level

If your new pool contains the same values for these attributes, the following warning message appears:

Pool *pool_name* has overlapping selection criteria.

Review the pool configuration and change overlapping criteria, and click Apply.

Expression Matching of Save Sets to Pools

When you specify save sets to include in a pool, you can specify exact save set names, or you can use regular expression matching to send a group of save sets to a specific pool.

Regular expression matching affects how the save set path is written. Backslashes (common in Microsoft Windows systems) and periods must be escaped with the backslash character to avoid being interpreted as part of the regular expression. For example, an entry reading:

/space/a.log

would match /space/axlog, /space/ablog, /space/a.log, and so forth. To match only the file a.log, this would have to be written:

/space/a\.log

Similarly, the Microsoft Windows directory `C:\SomeDir\MyFiles` would have to be written:

```
C:\\SomeDir\\MyFiles
```

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
/space/[jl]unk
c:\\Program Files\\bin
```

For more information on these commands and expression matching, refer to the *Legato Command Reference Guide* and the `nsr_regexp` and `nsr_pool` man pages.

Auto Media Verification

If the Auto Media Verify attribute in the Pool resource is enabled, the Sun StorEdge EBS server verifies data written to volumes from this pool by the following methods:

- Data is verified by repositioning the volume to read a portion of the data previously written to the media. If the data read matches the data written, verification succeeds; otherwise it fails.
- 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 Sun StorEdge EBS 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 Sun StorEdge EBS server is attempting to switch volumes, all save sets writing to the volume are terminated.

▼ To Edit a Pool

You cannot change the name of a pool.

Note – The following preconfigured pools cannot be modified: Archive, Archive Clone, Default, Default Clone, Migration, Migration Clone, Indexed Archive, Indexed Archive Clone, PC Archive, and PC Archive Clone.

To edit a pool:

1. In the Sun StorEdge EBS Administrator program, select Pools from the Media Menu.
2. In the Pools list, select the pool to edit.
3. Edit the pool's attributes, and click Apply.

▼ To Delete a Pool

Note – Preconfigured pools cannot be deleted. You can delete other pools as long as there are no active volumes assigned to the pool.

To delete a pool:

1. In the Sun StorEdge EBS Administrator program, select Pools from the Media Menu.
2. In the Pools list, select the pool to be deleted.
3. Click Delete.
4. When prompted, click OK to confirm the deletion.
5. Click Apply.

▼ To Create an Archive Pool

Note – An appropriately labeled volume must be mounted on a storage device for the archive process to proceed. For more information, see [Chapter](#) .

To create an Archive pool:

1. Create a new pool. For instructions, see [“To Create a Pool” on page 176](#).
2. Set the Pool Type to Archive.

As a result, when you schedule a new archive request, you may use the new Archive pool you created. If you choose to clone archived data, you can also create a new archive clone pool. The Sun StorEdge EBS server will write archive data only to an archive volume, and archive clone data only to an archive clone volume. For more information on pools, see [“Pools” on page 166](#).

Volume Labeling

The Sun StorEdge EBS software labels (initializes) each storage volume with a unique internal label that corresponds to a pool. During backup and other operations, Sun StorEdge EBS software identifies the pool to which a volume belongs by its label. 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. If you need to recover data, the server prompts you for the specific volume that contains the required data, identifying the required volume by the name given to it when it was labeled.

How the Sun StorEdge EBS Server Uses Label Templates

The Sun StorEdge EBS server writes a given set of data to a specific pool. For the server to recognize that a particular volume correctly belongs to a particular pool, the volume must have an internal identification label that associates it with the correct pool. The contents of the volume label follow rules defined in a specific label template that was created in the Label Template resource. You then associate a label

template with a specific pool in the Pool resource. If you do not associate data with a specific pool, the Sun StorEdge EBS server uses the preconfigured Default pool and corresponding Default label template.

FIGURE 5-2 illustrates how a pool configuration uses its associated label template to label a volume. You must configure a label template before you configure the associated pool for the custom template to be available in the Pool resource.




FIGURE 5-2 Labeling a Volume Using a Label Template

Preconfigured Label Templates

The following is a list of preconfigured label templates contained with the Sun StorEdge EBS server. These label templates correspond to the preconfigured pools.

- Default
- Default Clone
- Archive
- Archive Clone
- Full
- Indexed Archive

- Indexed Archive Clone
- Migration
- Migration Clone
- NonFull
- Offsite
- PC Archive
- PC Archive Clone
- Two Sided

▼ To Complete Label Template Attributes

There are certain guidelines to keep in mind when completing the attributes for a Label Template resource. [TABLE 5-4 on page 182](#) describes how to complete the key attributes for this resource. For more information on creating a label template, see [“To Create a Label Template” on page 186](#).

TABLE 5-4 Key Label Template Attributes (1 of 2)

Attribute	Guidelines
Name	Keep the label name consistent with the pool name, so that you can easily see how the data is organized. For example, you can create a label template called “AcctFull” for volumes that belong to a pool called “Accounting Full.” The Sun StorEdge EBS software does not allow the following characters in label template names: / \ * [] () \$! ^ ' ; ' ~ < > & { } In addition, you cannot use the following four characters, because they are used as separators in label templates: • colon (:) • dash (-) • period (.) • underscore (_)

TABLE 5-4 Key Label Template Attributes (2 of 2)

Attribute	Guidelines
Field	A label template is made up of one or more fields. Each field, or component, provides a layer of specificity to the organizational structure. There can be any number of components, but it is best to keep the template simple with as few as needed. The label cannot exceed 63 characters (60 characters for advanced file type devices). Space characters are not allowed. 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, <i>not</i> 1-99, or aaa-zzz, <i>not</i> aa-zzz. (This rule does not apply to a list of character strings or words; character strings should be separated by a blank space.) 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 5-5 on page 184 illustrates how label templates use components to create a number sequence for volume labels.
Separator	Choose which symbol you want to appear between component entries. You can use the period, dash, colon, or underscore to separate each component of the label template. If you do not select a separator, the label components do not have separators (for example, AA00aa), which makes the labels 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). • If you want 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. • If you want the Sun StorEdge EBS server to generate the first label for you, 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 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 5-5 lists examples of number sequences for volume labels.

TABLE 5-5 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 lower case 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)

Your 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 Sun StorEdge EBS server reaches the end of the template numbering sequence, it wraps around 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.

Label Template Naming Strategies

The Sun StorEdge EBS server is packaged with preconfigured label templates that correspond to the preconfigured pools. If you choose to create the templates, include any number of components in the Fields attribute. However, it is best to create it with as few components as necessary for the particular 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 5-6](#) illustrates several ways you can use components to organize the labels.

TABLE 5-6 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 suggestion is to name the volumes with the name of the Sun StorEdge EBS server followed by a three-digit number, for example:

```
jupiter.001
```

No matter what naming scheme you choose to use, keep in mind that the simpler a convention is, the easier it will 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.

When you label volumes, 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 the Sun StorEdge EBS software. For example, if you physically label a volume *mars.1*, its Sun StorEdge EBS name should also be *mars.1*.

▼ To Create a Label Template

For information on the template attributes, see [TABLE 5-4 on page 182](#).

To create a label template:

1. In the Sun StorEdge EBS Administrator program, select Label Templates from the Customize menu.
2. Click Create.
3. In the Name attribute, enter the name of the label template.
4. (Optional) In the Comment attribute, enter a description of the label template.
5. Enter the components for the label template in the Fields attribute. Press [Return] after entering each component.
6. In the Separator attribute, select a separator symbol.

The separator symbol is inserted between the components in the Fields attribute. For example, a label template with components of host, archive, and 001-999 and a dot (.) for a separator symbol will label tapes using the following format host.archive.[001-999].

If no symbol is selected, the components will have no separators, for example, hostarchive[001-999].

7. (Optional) For the Next attribute, enter the next label that the Sun StorEdge EBS server is to use to label the volume according to the template.
8. Click Apply.

▼ To Edit a Label Template

You cannot change the name of a label template. However, if you want 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. From the Customize menu, select Label Templates to open the Label Templates window.
2. In the Label Templates attribute, select the label template to be edited.
3. Edit the attribute values of the label template, and click Apply.

▼ To Delete a Label Template

You cannot delete the preconfigured label templates or a label template currently in use.

To delete a label template:

1. From the Customize menu, select Label Templates to open the Label Templates window.
2. In the Label Templates attribute, from the list displayed, select the label template to be deleted.
3. Click Delete.
4. When prompted, click OK to confirm the deletion.
5. Click Apply.

Devices

A device is a drive that reads and writes data to storage volumes or a disk file during backup, recovery, and other operations. You may add additional devices to the Sun StorEdge EBS server since the software supports multiple devices in a series. That is, it will use a tape in one device until it is full and then use the tape in the second device. If you add or remove a device from the Sun StorEdge EBS server, the server will automatically recognize and display the change.

You must configure each device before the Sun StorEdge EBS server can recognize and use it. The Device resource contains the attributes for each device that you must set. The instructions for configuring devices differ depending on whether the device is stand-alone or is contained in an autochanger or silo.

Device Support by Sun StorEdge EBS Edition

Each edition of Sun StorEdge EBS software has a different maximum number of devices a server can administer. Each enabled storage node connected to the Sun StorEdge EBS server can increase the device maximum. The maximum number of devices for any combination of Sun StorEdge EBS server and storage nodes is 256. Device levels for other Sun StorEdge EBS versions and releases can vary, so refer to the documentation of that particular server.

[TABLE 5-7](#) displays the maximum number of devices allowed for this Sun StorEdge EBS software release, including the maximum for each enabled storage node.

TABLE 5-7 Maximum Devices

Sun StorEdge EBS Product	Server	Each Storage Node
Workgroup and Business Editions	4	Not Applicable
Network Edition	16	16
Power Edition	32	32

If your server is configured to support more than one device, it will use the volume in a device until it is full, then use the volume in another device, and so on.

Use concurrent devices to spread backups evenly across your storage devices, and use the optional Sun StorEdge EBS Autochanger Software Module for unattended backups to an autochanger.

For a list of the currently supported storage devices, refer to the *Legato Compatibility Guides* on the Legato web site at www.legato.com.

▼ To View Supported Devices and Media Types

The Sun StorEdge EBS server supports a wide range a storage devices and media types, including tape, optical, and file type devices. The storage devices and supported media types can sometimes vary.

To view the media types available to a particular Sun StorEdge EBS server:

1. From the Media menu, select Devices to open the Devices window.
2. In the Media Type attribute, display the list of supported devices. This list displays abbreviations for all the device types supported in that Sun StorEdge EBS version.
3. From the Help menu, select On Window to view more detailed descriptions of each device type.

Filesystem devices are configured differently from traditional devices and are covered separately in this guide. For information on file device types, see [“Sun StorEdge EBS DiskBackup” on page 206](#).

Configuration of the Sun StorEdge Network Foundation HBA/Driver

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. However, before you run the script, you must ensure that all devices are powered on and connected to the HBAs.

Note – The `/usr/sbin/nsr/lus_add_fp_devs` script should only be used 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 you add a new device, or remove an old device, you must re-create and reinstall the new `lus` configuration. For information explaining how to add or delete a device, see [“To Add a Device” on page 192](#).



Caution – When you run `/usr/sbin/nsr/lus_add_fp_devs` script, check that the World Wide Port Names appear in the script output. If none appear, the fiber channel may be incorrectly configured. 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 `lus` is not currently loaded.

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

Note – If you select no, 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 you select no to the previous prompt, the program exits.

```
scsidev@0.0.0:FUJITSU MAJ3182M SUN18G 0804|Disk,  
/dev/rdisk/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
```

Nonrewinding Tape Device Usage

If you use tape drives as storage devices, you must use nonrewinding devices. The Sun StorEdge EBS server writes a filemark on the volume at the end of each backup and then appends data onto the volume based on the position of the filemark. If the device rewinds the media, the filemark position will be lost and previously written data is overwritten by the next backup. The pathnames for these devices must follow the Berkeley Software Distribution (BSD) semantic rules, for example, `/dev/rmt/0mbn`. The “b” in the pathname satisfies the BSD semantics requirement.

Block Size Mode

Ensure that the block size mode for tape devices used with Sun StorEdge EBS software is set to *variable*. Otherwise, data recovery may fail. The procedure for setting the device block size varies depending on your operating system.

For information about setting the tape device block size on your operating system, refer to the documentation that comes with your operating system.

▼ To Add a Device

You need to create a new Device resource for each stand-alone device or autochanger device. When you elect to create a Device resource, some of its attributes have default values. You can use these values if they apply to your storage device, or change them to a different value.

To add a device from the Sun StorEdge EBS Administrator program:

1. From the Media menu, select Devices to open the Devices window.
2. Click Create.
3. In the Name attribute, enter the device pathname.
If you use a file device, you must enter it as a directory path rather than only a filename.
4. From the Media Type attribute, select the appropriate media type.
If you want to back up to a file-type device configured on your disk, see [“Sun StorEdge EBS DiskBackup” on page 206](#) for specific instructions and restrictions. For more information on the logical selection, see [“Logical Volume Backup” on page 156](#).
5. (Optional) For the Read-Only attribute, select Yes to have the device read volumes only (instead of writing the data to volumes).
6. In the Target Sessions attribute, enter or select a value. For more information on how this attribute works, see [“Target Sessions” on page 58](#).
7. In the Cleaning Interval attribute, enter a date to set the period of time between cleanings. Note that the Cleaning Required attribute must be also enabled. Specify the time in days, weeks, or months. For example:
 - Four days
 - Two weeks
 - Month (same as one month)This attribute is only valid for autochanger tape devices.
8. (Optional) In the Auto Media Management attribute, select Yes to enable this feature. For information on this feature, see [“Auto Media Management” on page 226](#).
9. Complete other attributes as necessary, and click Apply.

▼ To Edit a Device

To edit a device from the Sun StorEdge EBS Administrator program:

1. From the Media menu, select Devices to open the Devices window.
2. In the Devices attribute, select the device to be modified from the list displayed.
3. Edit the attributes, and click Apply.

▼ To Delete a Device

Note – You cannot delete devices that have volumes mounted.

To delete a device from the Sun StorEdge EBS Administrator program:

1. From the Media menu, select Devices to open the Devices window.
2. In the Devices attribute, select the device to be deleted.
3. Click Delete.
4. When prompted, click OK to confirm the deletion.
5. Click Apply.

Device Parameter Environment Variables

A number of environment variables can be used to adjust several Sun StorEdge EBS device parameters:

- **NSR_DEV_BLOCK_SIZE_MEDIA_TYPE**
- **NSR_DEV_TAPE_FILE_SIZE_MEDIA_TYPE**
- **NSR_DEV_DEFAULT_CAPACITY_MEDIA_TYPE**
- **NSR_DEV_LOAD_TIME_MEDIA_TYPE**
- **NSR_DEV_LOAD_POLL_INTERVAL_MEDIA_TYPE**
- **NSR_DEV_LOAD_TRY_LIMIT_MEDIA_TYPE**

Where *MEDIA_TYPE* is the backup device type available to the Sun StorEdge EBS server (also found in the Media Type attribute of the Device resource in the Sun StorEdge EBS Administrator program). Note that the syntax for the device name must be all uppercase, with underscores (_) replacing blank spaces and hyphens. Therefore, a device displayed in the Sun StorEdge EBS software as "8mm Mammoth-2" would be listed as:

8MM_MAMMOTH_2

Environment variables should be set in the Sun StorEdge EBS startup script (/etc/init.d/networker) before the lines:

```
(echo 'starting Sun StorEdge EBS daemons:') > /dev/console
```

Enter the environment variables using the format:

```
ENV_VAR_NAME = value
export ENV_VAR_NAME
```

Note – You must restart the Sun StorEdge EBS server for the environment variables to take effect.

NSR_DEV_BLOCK_SIZE_MEDIA_TYPE

NSR_DEV_BLOCK_SIZE_MEDIA_TYPE is organized in units of kilobytes. 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 [“Setting Compatible Block Size for UNIX and Microsoft Windows” on page 390](#).

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 will generally lead to faster positioning. For example:

```
NSR_DEV_TAPE_FILE_SIZE_TZ89=512
```

The Sun StorEdge EBS software writes a filemark by closing and reopening the tape device; requiring a second or two. If this value is too small, throughput may be slowed.

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
```

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 may be unnecessary load failures. If it is too long, then labeling new tapes will take longer than necessary. The minimum allowable value is 10 seconds, the maximum value is 600 seconds. For example:

```
NSR_DEV_LOAD_TIME_DTL8000=300
```

NSR_DEV_LOAD_POLL_INTERVAL_MEDIA_TYPE

NSR_DEV_LOAD_POLL_INTERVAL_MEDIA_TYPE is the number of seconds that **nsrmmd** waits between 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`

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 one of two options in the CDI attribute of the Device resource:

- Not Used—this means the CDI feature is not enabled.
- SCSI Command— when this option is enabled, it means the Sun StorEdge EBS software is using the new method to send explicit SCSI commands to tape devices.

When the CDI feature is enabled, it does the following:

- Provides clearer tape status messages.
- Informs you 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, this method can be time-consuming if there are 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 contents –it simply needs to exist. This will disable the use of CDI for that server and all storage nodes controlled by that server.

The CDI feature is not supported in these cases:

- Within an NDMP environment.
- On an SGI or Sequent server or storage node.

Note – The CDI feature is not supported for tape devices connected directly on SGI or Sequent systems. However, it can be used by CDI- supported storage nodes (for example, Solaris) controlled by a Sun StorEdge EBS server running on SGI or Sequent.

Note – Using 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.

Essentially you should only set or disable the CDI feature on the advice of a Legato Technical Support representative. In a situation where you have enabled the CDI feature and are experiencing any tape or SCSI issues, contact Legato Technical Support. The CDI feature enables Sun StorEdge EBS to collect better diagnostic information and facilitates when Sun Technical Support needs to address any issues you may encounter.

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.

This section contains the following topics:

- [“Introduction to DDS” on page 198](#)
- [“DDS Configuration” on page 199](#)
- [“Block Size Compatibility” on page 200](#)
- [“Enabling DDS within NDMP” on page 201](#)
- [“DDS Attributes in the Devices Dialog Box” on page 203](#)
- [“To Create a Device and Enable DDS” on page 204](#)
- [“High Availability and DDS” on page 205](#)
- [“DDS Licensing” on page 206](#)

For information explaining how to enable DDS when configuring a device with the `jbconfig` program, see [“To Enable DDS When Configuring an Autochanger” on page 378](#).

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 enable DDS when creating a device through the Sun StorEdge EBS Administrator program or when configuring a device with the **jbconfig** program.

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 – Currently, DDS is only supported in a storage area network (SAN) Fibre Channel environment and not a direct-connect SCSI environment.

Benefits of DDS

Enabling DDS on your system provides you with the following benefits:

- Reduces storage costs — A single tape drive can be shared out among several storage nodes. In fact, since Sun StorEdge EBS software uses the same open tape format for UNIX, Windows NT, NetWare and Linux, you can actually share the same tape between different platforms (assuming that respective save sets belong to the same pool).
- Reduces LAN traffic — You can reconfigure clients as SAN storage nodes which can send save sets over the SAN to shared drives.
- Provides fault tolerance — Within a SAN environment you can configure your hardware to eliminate a single point of failure.
- Provides configuration over a greater distance — Enables you to configure a system over a greater distance than with SCSI connections.

DDS Configuration



Caution – DDS licenses must be installed before the **jbconfig** command is used to configure shared devices. See the Sun StorEdge EBS Roadmap for information about licensing DDS.

FIGURE 5-3 on page 200 illustrates the DDS process and potential configurations for sharing drives. This basic configuration consists of a server, two storage nodes, and an autochanger with two tape drives.

In this figure, storage nodes sn_1 and sn_2 are attached to the autochanger. Each node, on its own, has access to drive_1 and drive_2. With DDS enabled, both nodes have access to both drives and—more importantly—can recognize when a shared drive is in use.

For example, sn_1 sends a backup to drive_1 and subsequently, sn_2 starts a backup to drive_1. If the backup from sn_2 can use the same tape in drive_1, it will. If not, the backup started by sn_2 will be sent to the next available drive, which is drive_2.

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




FIGURE 5-3 Dynamic Drive Sharing

Block Size Compatibility

With DDS enabled, you can share a drive 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, you must 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 information on how to set block sizes for individual drives or tapes on different platforms, see [“Setting Compatible Block Size for UNIX and Microsoft Windows” on page 390](#).

Block Size Incompatibility between UNIX and Microsoft Windows

Incompatible block size settings between UNIX and Microsoft Windows storage nodes may result in any one of the following error scenarios:

- A backup taken on a UNIX node may not be recoverable on a Microsoft Windows node. 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 labelled 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 NT 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 NT in this situation fails because of positioning errors on the tape.

SCSI Reserve/Release Setting

If you are using DDS with either of following operating systems, you need to disable the scsi reserve/release setting for DDS to function correctly:

- HP Tru64 UNIX
To disable the scsi reserve/release setting on HP Tru64 UNIX, use the SCU command. For further information about using SCU, refer to the SCU man page.
- HP-UX 11.0 (with patch PHKL_20063)
To disable the SCSI reserve/release setting on HP-UX, set the kernel parameter `st_ats_enabled` to 0 to globally disable the SCSI reserve/release setting.

Enabling DDS within NDMP

The following sections explain the requirements for successfully enabling DDS with NDMP.

DDS on NDMP Nodes in a SAN Environment

You can share drives between NDMP nodes in a SAN environment. However, you must ensure the NDMP nodes are homogeneous.

For example, you can enable DDS 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 you to share drives between non-homogeneous NDMP nodes. However, there is no inherent limitation within DDS preventing this.

FIGURE 5-4 on page 202 illustrates a basic DDS configuration with NDMP.




FIGURE 5-4 DDS with NDMP

DDS Attributes in the Devices Dialog Box

The Devices dialog box includes two attributes used in configuring DDS:

- 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. You can update the Hardware ID during the device configuration process. Any character string can be used as a value.

The value for the Hardware ID attribute is a string which may be too long for display in the Sun StorEdge EBS Administrator graphical interface. Consequently, when hardware IDs are enabled, a unique positive integer is assigned to all the shared devices.

This number is shorter and is clearly displayed in the Sun StorEdge EBS Administrator interface to denote when two paths actually refer to the same device. The hardware ID is arbitrarily generated and assigned a run time. If you reboot the servers or enable and disable the DDS feature, the hardware ID will be regenerated and reassigned a run time.

Changing the Hardware ID attribute under these circumstances does not indicate any limitation in the DDS functionality.

Shared Devices Attribute

The Shared Devices attribute features values which 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 5-8](#) lists this attribute's allowed values, and their descriptions.

TABLE 5-8 Shared Devices Attributes

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

▼ To Create a Device and Enable DDS

You can enable the DDS feature when you create a device. If you assign a nonunique hardware ID to the device in [Step 4](#), you also enable DDS—when another device has the same hardware ID assigned to it.

Device instances sharing the same physical drive across multiple hosts have the same hardware ID. This indicates to Sun StorEdge EBS software that these multiple devices are the same physical drive.

To create a device and enable DDS:

1. From the Media menu, select Devices to open the Devices window.
2. Click Create.
3. Enter the device pathname into the Name attribute and complete the remaining attributes as required. For information about how to create a new device, see [“To Add a Device” on page 192](#).
4. Assign a nonunique hardware ID to the device. The hardware ID accepts any character string as a value.
5. Complete any other attributes as necessary and click Apply.

▼ To Stop Sharing a Device through the Devices Window

To stop sharing a drive that has been assigned a hardware ID, remove the value from the hardware ID attribute in the Devices window.



Caution – When you stop sharing a DDS-configured drive, this does not mean that the drive is disabled. This means that Sun StorEdge EBS software is no longer aware that a specified drive is shared.

Consequently, so as not to impede Sun StorEdge EBS operations when you stop sharing a device, remove the other device instances.

▼ To Change the Idle Device Timeout Attribute

You might notice that a tape remains mounted in its drive for 10 to 15 minutes after a backup has completed. Other requests for the drive from another device path will wait during this period.

You can modify this timeout value by adjusting the Idle Device Timeout attribute. To change this attribute:

1. In the Sun StorEdge EBS Administrator program, select Jukeboxes from the Media menu.

2. In the Jukeboxes dialog box, select Details from the View menu.
3. Scroll down to Idle Device Timeout.
4. Select a time that is appropriate for your system configuration.

The default Idle Device Timeout value is 10 minutes.

Note – Selecting zero means that the device never times out. Consequently, you would have to eject the tape 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, 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.

DDS Licensing

To enable DDS, you need one license per 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 autochanger and storage node licensing.

Sun StorEdge EBS DiskBackup

One Sun StorEdge EBS data storage option is to save data to a computer's local or network-attached disk, rather than use a tape device to write the data to storage media. Storing data using the Sun StorEdge EBS DiskBackup solution greatly reduces the time it takes to both save and recover data compared to using tape.

Sun StorEdge EBS software offers two variations of this method:

- File Type Device
- Advanced File Type Device

This section contains the following topics about the Sun StorEdge EBS DiskBackup Option:

- [“File Type Device” on page 206](#)
- [“Advanced File Type Device” on page 207](#)
- [“Differences between File Type and Advanced File Type Devices” on page 209](#)
- [“Differences in the Cloning Process” on page 212](#)
- [“Changes in How Sun StorEdge EBS Administrator Displays” on page 213](#)
- [“Configuring Sun StorEdge EBS DiskBackup” on page 218](#)
- [“Sun StorEdge EBS DiskBackup Licensing” on page 224](#)

File Type Device

You configure the Sun StorEdge EBS server to back up data to a file type device by creating a new Device resource for your Sun StorEdge EBS server the same way as for any other storage device.

When you configure a file type device, the following conditions apply:

- Multiple file device; however, each device must have a different name.
- Dynamic Drive Sharing and NDMP are *not* supported.

- If you plan to use multiple file type devices on the same disk, it is recommended that you partition the disk and create only one file type device per partition.

Note – This release of the Sun StorEdge EBS software also offers the advanced file type device. To take advantage of the advanced file type device, you can clone the file type save sets to an advanced file type device.



Caution – The upper limits of save set size—with file type device as a backup solution—depend on the upper limits supported by your operating system or the file size specified by your disk device vendor.



Caution – When you create a file type device on a network filesystem (NFS), the following restrictions apply:

- The filesystem 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 filesystem.
 - If the file device was used prior to setting the Volume Default Capacity attribute, the data for that file device must be staged or cloned to another device.
 - If you change the device capacity of a volume, you must relabel the volume for that change to take effect.
 - Subsequently, the directory contents must be deleted and the file device must be re-created. If not, the change in volume default capacity does not take effect.
-

Advanced File Type Device

The *advanced file type* (`adv_file`) is designed for very large disk devices. It differs from the file type device because the volume for the advanced file type device is not marked full unless the maximum number of save sets is reached.

When an advanced file type device runs out of disk space, the current backup is suspended. The following message appears:

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



Immediately after this message appears, the action that is associated 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 until more space is made available. If enough space has been cleared, the backup continues.

Caution – When using the advanced file type device, the upper limits of save set size depend on the upper limits supported by your operating system or the file size specified by your disk device vendor.

The advanced file type device deletes expired save sets depending on how your retention policy is set.

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 logs a message in the `/nsr/logs/media.log` file.

At this point, the backup ceases until you make more storage space available. The notifications listed here are new default notifications pertaining specifically to the advanced file system type device. For further information explaining how to create or edit a notification, see [“Notifications” on page 65](#).

For further information regarding backup issues, see [“How the Sun StorEdge EBS Software Backs Up Data” on page 11](#) or the `savegrp` man page.

Note – Dynamic Drive Sharing and NDMP are not supported with the advanced file type device.

Maximum Number of Save Sets for Advanced File Types

The maximum number of save sets the media database allows in a volume is 270,000. To avoid this limit when using a large disk array, do one of the following:

- Configure multiple `adv_file` type devices into different backup pools.
- Limit the number of save sets to an `adv_file` volume.

For more information about the save set limit, see [“Maximum Number of Save Sets Per Volume” on page 243](#).

Differences between File Type and Advanced File Type Devices

TABLE 5-9 lists the functionality differences between the file type device and the advanced file type device.

TABLE 5-9 Advanced File Type and File Type Device Differences
(1 of 2)

Function or Operation	Advanced File Type Device	File Type Device
Creating a device in Devices window	Select media type: <code>adv_file</code> 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: <code>file</code> UNIX/Linux: local or NFS. For further information, see the Important Note on page 208 . Windows storage node: local path only
Labelling the device	• Volume name cannot exceed 60 characters. • Automatically creates an <code>_AF_readonly</code> device, then labels and automounts a volume with a <code>.RO</code> suffix	• Volume name cannot exceed 63 characters
Concurrent Operations: For further information about concurrent operations, see TABLE 5-11 on page 212 .	Yes	No
Save set ID / clone ID: For further information about concurrent operations, see TABLE 5-11 on page 212 .	• One save set ID with two clone IDs created for each save set • A second (virtual) clone created on <code>_AF_readonly</code> device • User can specify virtual clone ID for concurrent operations.	One save set ID created with one clone ID for each save set.

TABLE 5-9 Advanced File Type and File Type Device Differences
(2 of 2)

Function or Operation	Advanced File Type Device	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.	The nsrim program removes both aborted and expired save sets, once every 24 hours, after a savegroup is completed. Note: When you upgrade the Sun StorEdge EBS server and storage node 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.	Changing device capacity requires volume to be relabeled for capacity change to take effect. For further information, see the Important Note on page 207 .
What happens when Filesystem or Volume is full	- Message displayed stating filesystem require more space nrsim invoked to reclaim space for expired save set on advanced file type device - Notification sent by e-mail stating device is full - Device waits until space become available; volume is never marked 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	Savestream (uasm) format (uses less space)	Legato Open Tape Format (OTF)

Advanced File Type Device Concurrent Operations

When the primary 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 you to conduct concurrent operations, such as reading from it. Do not make any changes to the *_AF_readonly* device resource. However, if you make a change to the primary volume or save set, you may propagate these changes to the secondary volume or save set.

See [TABLE 5-10 on page 211](#) for information detailing which changes must be propagated from the primary device, volume or save set to the secondary device, volume or save set.

For example, when you create and label the primary advanced file type device */mercury/dev1*, the following information would be displayed in the Devices window.

TABLE 5-10 Devices Window Display

Device	Type	Volume
<i>/mercury/dev1</i>	<i>adv_file</i>	<i>af_1</i>
<i>/mercury/dev1/_AF_readonly</i>	<i>adv_file</i>	<i>af_1.RO</i>

In this example, when you label the advanced file type device, 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 on this device.

[TABLE 5-10](#) illustrates that when you are backing up data to the device:

/mercury/dev1

you can concurrently read from the device:

/mercury/dev1/_AF_readonly

Note – The terms *_AF_readonly* name and the *.RO* suffix are reserved for the advanced file type. Do not create any devices or volumes using these terms.

TABLE 5-11 displays which Sun StorEdge EBS operations can be performed concurrently when you are using the advanced file type device for data storage.

TABLE 5-11 Concurrent Operations with Advanced File Type Device

Backups	Recovers	Manual Cloning	Automatic and/or Manual Staging
X	X		
X		X (1 clone)	
X			X (1 stage)

With the Sun StorEdge EBS DiskBackup option, you can only perform the following operations:

- Multiple concurrent backups and multiple recovers
- Multiple concurrent backups and only one manual cloning.
- One automatic or manual staging operation while performing multiple concurrent backups

Differences in the Cloning Process

For both file type and advanced file type devices, save sets are cloned one at a time. However, there are important differences in the cloning process for the two types of devices:

- For file type devices, both automatic and manually cloning begins only after all the save sets in a savegroup have been backed up.
- For advanced file type devices, automatic cloning begins after the save sets 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 you have 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 immediately 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 immediately launch the cloning process for that respective save set. Essentially, this is a time-saving feature that enables you to optimize the Sun StorEdge EBS operations.

You can only clone one save set at a time.

Changes in How Sun StorEdge EBS Administrator Displays

The following sections illustrate the display changes that occur with the implementation of the advanced file type device.

Device Display

As shown in [FIGURE 5-5](#), both the primary and the secondary *read-only* device with the *_AF_readonly* suffix are displayed.

- Volumes corresponding to both the primary and secondary *read-only* devices appear.
- Volume corresponding to the *_AF_readonly* device appears with the *.RO* suffix.

FIGURE 5-5 Device Display with Advanced File Type Device

Volume Display

As shown in [FIGURE 5-6](#), with the advanced file type device, the following differences are apparent:

- Read-only volumes appear with the *.RO* suffix.
- Cloned, *read-only* volumes are displayed with the *_c* and the *.RO* suffix.

Volumes							
File Volume Save Set							Help
Volumes							
Volume	Bar code	Pool	Written	%Used	Mode	Expiration	Location
TestBackup.007		TestBackup	0 KB	0%	appen	-	
TestBackup.007.R0		TestBackup	0 KB	0%	appen	-	
TestBackup.009		TestBackup	60 KB	0%	appen	10/11/03	
TestBackup.009.R0		TestBackup	0 KB	0%	appen	10/11/03	
TestBackup_c.009.R0		TestBackup	0 KB	0%	appen	10/11/03	

Save Sets				
Client	Date	Level	s	Save Set

FIGURE 5-6 Volume Display with Advanced File Type Device

FIGURE 5-7 on page 215 shows how a save set is displayed when the advanced file type device is used as a backup device.




FIGURE 5-7 Save Set Display with Advanced File Type Device

Propagating Changes from Primary to Read-Only

When you create an advanced file type device, an *_AF_readonly* device is automatically created and added to the Sun StorEdge EBS device list. Whenever data is backed up, there are some additional entries added to the Sun StorEdge EBS media index for the data backed up to that device.

When changes are made to the user-created, advanced file type device (primary device) attributes, you need to immediately make the same changes to the attributes of the *_AF_readonly* (secondary device). If you make manual changes to volume or save set properties within the Sun StorEdge EBS indexes, you must also propagate those changes to the other read-only volumes and save sets.

For example, if you change the mode of a primary volume, you also need to change the mode of the read-only volume.

Because propagating the changes depends on individual system use and requirements, this process has not been automated. Consequently, you must propagate the required changes manually.

To propagate the change manually, access the primary device, volume or save set and change the appropriate attribute. Subsequently, access the read-only device, volume or save set to match the attribute change made to the primary device, volume or save set.

Note – You should always make any attribute changes to the primary (read / write) device, volume or save set first and then propagate these changes to secondary (read-only) device, volume or save set. Because the secondary devices, volumes and save sets are essentially clones of the primary devices, volumes and save sets, it is more efficient to propagate changes from primary to secondary.

[TABLE 5-12](#) lists the attributes you must propagate manually from the primary device, volume or save set to the read-only device, volume or save set.

TABLE 5-12 Attributes Requiring Manual Propagation

Primary Device, Volume or Save set	Attribute
Device	Enabled
Volume	Mode
	Remove
Save set	Mode

Note – There are some attributes that do not need to be propagated from primary to secondary volume, which include:

- Volume location
- Volume mount

Configuration Support for Sun StorEdge EBS DiskBackup

The following configurations of the Sun StorEdge EBS DiskBackup solution are supported:

- Local disk of UNIX, Linux, or Microsoft Windows versions of Sun StorEdge EBS storage node. This is supported on both file type and advanced file type devices.

- 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 and advanced file type devices.
- 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 advanced file type devices.

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 with CIFS implementations that must behave as CIFS.

Example: DiskBackup within a System Configuration

[FIGURE 5-8 on page 218](#) illustrates how to use Sun StorEdge EBS DiskBackup within your system configuration. In this example:

- The advanced file type devices 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.
- The UNIX server stores the indexes and the bootstrap file.
- Linux/UNIX storage node A writes its backups to one of the following:
 - The advanced file type device through an NFS connection to Disk Device 1.
 - The advanced file type device on Local Disk 1.
- The Microsoft Windows Storage Node B uses a CIFS connection to back up to advanced file type device 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.

You can configure Sun StorEdge EBS DiskBackup to the specific configuration of your system.




FIGURE 5-8 Sample Sun StorEdge EBS DiskBackup Configuration

Reclaiming Space for Sun StorEdge EBS DiskBackup

When you are backing up data to file type and advanced file type devices you need to create a staging policy to move the data off the disk as soon as possible to prevent the filesystem from becoming full. To make space for additional backups, you can do any one of the following:

- Set up a staging policy
- Review and, if required, modify the retention policy of the clients

For more information on creating staging policies and using staging, see [“Save Set Staging” on page 244](#).

Configuring Sun StorEdge EBS DiskBackup

The following sections explain how to configure a Sun StorEdge EBS DiskBackup Device:

- On local or NFS-mounted UNIX/Linux disks. See [“To Configure a File Type or an Advanced File Type Device on UNIX and Linux” on page 219](#)
- On local Microsoft Windows disks. See [“To Configure a File Type or Advanced File Type Device on a Local Disk of a Microsoft Windows Storage Node” on page 220](#)
- On CIFS-mounted Microsoft Windows storage nodes (advanced file type device only). See [“Configuring a CIFS-Mounted Advanced File Type Device on a Microsoft Windows Storage Node” on page 221](#).

▼ *To Configure a File Type or an Advanced File Type Device on UNIX and Linux*



Caution – When you create a file type device on an NFS filesystem, the following restrictions apply:

- The filesystem 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 filesystem.
- 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. If these actions are not completed, the change in volume default capacity does not take effect.

Note – Do not modify any directories or files created by the file type or advanced file type device. If you make any changes to the device path or its contents, the Sun StorEdge EBS software cannot locate the device.

Creating either an advanced file type device or a file type device requires creation of a directory in the disk filesystem that the Sun StorEdge EBS server will recognize as the device name (and the destination for the data), for example:

`/export/backup`

For an NFS-mounted advanced file type device, ensure the NFS mount point is available by either appropriately automounting it or by manually mounting as required.

Note – Do not use the `tmp` directory for Sun StorEdge EBS DiskBackup devices because the data might be overwritten.

Note – Since the advanced file type device is never marked full, there is no need to add another device to the pool.

To configure an advanced file type device or advanced file type device:

1. Create one directory for each disk (or partition) being used for advanced file type device or file type device.

2. In the Sun StorEdge EBS Administrator, select Devices from the Media menu.
3. Click Create.
4. In the Name attribute, enter the complete path for the directory.
5. For the Media Type attribute, select one of the following:
 - File to indicate the device is a file type device.
 - Adv_file to indicate the device is an advanced file type device.
6. For the Target Sessions attribute, enter or select a value.
7. Leave the options for cleaning at their default disabled settings so that automatic cleaning is not invoked.
8. Do *not* enable the Auto Media Management feature for advanced file type device or file type device.
9. Complete other attributes as necessary, and click Apply.

You label and mount the Sun StorEdge EBS DiskBackup Device in the same manner as for other offline media used for Sun StorEdge EBS backups.

Note – Because the advanced file type device automatically affixes the .RO suffix to a volume name, you must limit the volume name to 60 characters when you label the volume.

Save set staging is especially useful when you are using either device type of the Sun StorEdge EBS DiskBackup feature. For more information, see [“Save Set Staging” on page 244](#).

▼ *To Configure a File Type or Advanced File Type Device on a Local Disk of a Microsoft Windows Storage Node*

You configure Sun StorEdge EBS to back up data to a file type or advanced file type device by creating a new Device resource for the Sun StorEdge EBS server the same way as for any other storage device.

Do not use the *TEMP* directory for the advanced file type device because the data might be overwritten.

Note – Because advanced file type devices are never marked as full, there is *no* need to configure another device in the same pool.

To configure the Sun StorEdge EBS DiskBackup option:

1. Create one directory for each disk (or partition) being used for file type device storage. Creating a file type device requires the creation of a directory in the disk filesystem that the Sun StorEdge EBS server will recognize as the device name (and the destination for the data), for example:

C:\device01

2. In the Sun StorEdge EBS Administrator program, right-click Devices from the Configure tab and select Create.
3. In the Name attribute, enter the complete path for the directory.
4. For the Media Type attribute, do one of the following:
 - Select file to indicate the device is a file type device.
 - Select *adv_file* to indicate the device is an advanced file type device.
5. Complete the other Device attributes as necessary and click OK.
6. For the Target Sessions attribute, enter or select a value.
7. Leave the options for cleaning at their default disabled settings so that automatic cleaning is not invoked.
8. Do not enable the Auto Media Management feature for advanced file type device or file type device.

Label and mount the Sun StorEdge EBS DiskBackup Device in the same manner as for other offline media used for Sun StorEdge EBS backups.

Note – Because the advanced file type device automatically affixes the *.RO* suffix to a volume name, you must limit the volume name to 60 characters when you label the volume.

Save set staging is especially useful when you are using either device type of the Sun StorEdge EBS DiskBackup feature. For more information, see [“Save Set Staging” on page 244](#).

Configuring a CIFS-Mounted Advanced File Type Device on a Microsoft Windows Storage Node

The following section explains how to configure the advanced file device on a CIFS connection on a Microsoft Windows storage node. Before you configure the advanced file type device, however, you must start the service as administrator.

▼ *To Start the Sun StorEdge EBS Service on a Microsoft Windows Storage Node*

You must start the Sun StorEdge EBS services as administrator to successfully access the network advanced file type devices on a Microsoft Windows storage node.

To start the Sun StorEdge EBS services as administrator, complete the following steps:

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.
4. Right-click Sun StorEdge EBS Remote Exec Service.
5. Select Properties.
6. Select the Log On tab.
7. Select This Account radio button.
8. Replace LocalSystem with Administrator.
9. Enter the administrator password in the appropriate text box.
10. Select OK.
11. Right-click the Sun StorEdge EBS Remote Exec service for Windows Storage Node.
12. Start the Sun StorEdge EBS Remote Exec service for Windows Storage Node.

▼ **To Configure a CIFS-Mounted Advanced File Type Device on a Microsoft Windows Storage Node**

The following section explains how to configure an advanced file type device on a Window storage node defined on a UNIX/Linux server.




FIGURE 5-9 CIFS-Mounted Advanced File Type Device

Creating an advanced file type device requires the creation of a directory in the disk filesystem that the Sun StorEdge EBS server will recognize as the device name (and the destination for the data), for example:

`\export\backup`

Note – Do not use the `temp` directory for the advanced file type device because the data might be overwritten.

Note – Because advanced file type devices are never marked as full, there is no need to configure another device in the same pool.

To configure the Sun StorEdge EBS DiskBackup option:

1. Create one directory for each disk (or partition) being used for advanced file type device.
2. In the Sun StorEdge EBS Administrator program, select Devices from the Media menu.
3. Click Create.
4. In the Name attribute, enter the complete path for the directory. For the Microsoft Windows storage node, you must use the UNC format:

```
rd=<sn_host>:\\<CIFS_host>\<share_point_name>\path
```

5. For the Media Type attribute, select `adv_file` to indicate the device is an advanced file type device.
6. For the Target Sessions attribute, enter or select a value.
7. Leave the options for cleaning at their default disabled settings so that automatic cleaning is *not* invoked.
8. Do *not* enable the Auto Media Management feature for advanced file type device or file type device.

Note – For the following steps, you must enter the administrator and password as the CIFS filesystem user name and password.

9. Enter administrator in the Remote User attribute. For example, in the configuration illustrated in [FIGURE 5-9 on page 223](#), you would enter the administrator of the Disk Device.
10. Enter the administrator’s password in the Password attribute. In the configuration illustrated in [FIGURE 5-9 on page 223](#), you would enter the remote password of the Disk Device.
11. Complete other attributes as necessary, and click Apply.



Caution – After you provide a new password for CIFS advanced file type device, you must restart the Sun StorEdge EBS daemons on both the server and storage node for the new password to take effect and to reestablish the CIFS connection correctly.

Note – You label and mount the advanced file device in the same manner as for other offline media used for Sun StorEdge EBS backups. Because the advanced file type device automatically affixes the `.RO` suffix to a volume name, you must limit the volume name to 60 characters when you label the volume.

Save set staging is especially useful when you are using either device type of the Sun StorEdge EBS DiskBackup feature. For more information about save set staging, see [“Save Set Staging” on page 244](#).

Sun StorEdge EBS DiskBackup Licensing

Both Sun StorEdge EBS DiskBackup options— file type and advanced file type device —are available in all editions of the Sun StorEdge EBS 7.x software. 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, you must obtain the permanent license. If the enabler for your 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 will no longer be available for saves.

Without a valid license, you will be able to recover the backed up data. You will *not*, however, be able to create new Sun StorEdge EBS DiskBackup devices or mount any Sun StorEdge EBS DiskBackup device volumes, except as read-only. When the license has expired, you will not be able to send backups to already-mounted writable volumes on advanced file type devices.

Sun StorEdge EBS DiskBackup Licensing Error Messages

This section lists the error messages that might appear if your system is not licensed appropriately.

- If the server is *not* licensed for the advanced file device type, the device creation process will fail and a message similar to the following appears:

```
Server is not licensed for adv_file support. Install adv_file
enabler
```

- If the enabler for your Sun StorEdge EBS server or Sun StorEdge EBS DiskBackup feature expires (for example, the evaluation period is terminated), any advanced file type devices that exist will no longer be available for saves. If you attempt to write to the device, the following error message is written to the log file:

```
Server is not licensed for advanced file support. Install adv_file
enabler
```

Media Handling Error Resolution

A damaged tape can produce media handling errors. Trying to mount and read a damaged tape may result in a continuous loop of failed actions. The device can continuously try to mount the tape, replace it in the slot, and then retry the action again with the same result.

▼ To Limit Repeated Device Errors

To minimize the number of instances that this type of loop can occur:

1. From the Media menu, select Devices to open the Devices window.
2. From the View menu, select Details to display the hidden attributes.
3. In the Max Consecutive Errors attribute, lower the value (set to 20 by default), and click Apply.

When the number of errors meets the value set in this attribute, the device stops attempting the activity and becomes disabled.
4. Once the host running the Sun StorEdge EBS server is idle, remove any media from the disabled device and ensure that the device is in good working order.
5. With the Devices window still open, re-enable the device.
6. From the list in the Devices attribute, select the disabled device.
7. In the Enabled attribute, select Yes.
8. In the Write Enabled attribute, select Yes, unless you want the volume to be read-only.
9. Click Apply.

If the disabled device is part of an autochanger, you may need to reset the device by doing the following:

1. Open a terminal window.
2. Change the path to the directory containing the Sun StorEdge EBS executables.
3. Enter the following command:

```
# nsrjb -HE
```

Auto Media Management

The auto media management feature gives the Sun StorEdge EBS server automatic control over media loaded in the storage device. If you enable the auto media management feature in the Device resource, the server automatically:

- Labels the volume.

- Mounts the volume.
- Overwrites volumes it considers unlabeled.
- Recycles volumes eligible for reuse that are loaded into the device.

The auto media management feature is only enabled for stand-alone devices in the Device resource. To enable auto media management for devices in an autochanger, see [“Auto Media Management with Autochangers” on page 399](#).

The Sun StorEdge EBS server considers a volume unlabeled with the following conditions:

- Has no internal label.
- Is labeled with information other than a recognizable Sun StorEdge EBS label.
- Is labeled with a Sun StorEdge EBS label, but the density indicated on the internal label differs from the density of the device where the volume is mounted.

Because the auto media management feature can relabel a volume with a different density, it is possible to inadvertently overwrite data that still has value. For this reason, be careful if Sun StorEdge EBS volumes are shared between devices with different densities.

If you do not enable the auto media management feature, the Sun StorEdge EBS server ignores unlabeled media and does not consider it for backup.

If you enable the auto media management feature for a stand-alone device, the following processes occur when a volume becomes full during a backup:

1. 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.
2. The device is monitored and the software waits for another volume to be inserted into the device.
3. After a volume is detected, a check is performed to determine whether the volume is labeled. If it is already labeled, the volume is mounted into the device. The Sun StorEdge EBS server checks to see whether the newly mounted volume is a candidate to write data to.
 - If so, the write operation continues.
 - If not, the Sun StorEdge EBS server continues to wait for a writable volume to continue the backup.
4. If the volume is recyclable and is a member of the required pool, it is recycled the next time a writable volume is needed.
5. If the volume is unlabeled, it is labeled when the *next* writable volume is needed for a save.

If a volume that is not full is unmounted from a stand-alone drive and you enabled auto media management, the Sun StorEdge EBS server waits for 60 minutes before it automatically remounts the volume in the drive. This hour delay provides time to unload the volume after unmounting.



Caution – 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.

Storage Nodes

Storage nodes are host computers with attached remote 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. By using Sun StorEdge EBS storage node software, you can route client data directly to a storage node's storage devices without the data first going to the Sun StorEdge EBS server. For a complete description of storage nodes, see [“Storage Nodes” on page 83](#).

▼ To Configure a Storage Node

You must configure each storage node device. This procedure is similar to the tasks required for adding and configuring devices in [“Devices” on page 187](#).

Here is an overview of how to configure a storage node:

1. Install a client (with the required storage node software).
2. Install the driver for each remote device.
3. On the storage node, run `jbconfig -s server` from the command line, where *server* is the host on which the Sun StorEdge EBS server resides, to configure the jukebox.
4. Verify that the storage node device name follows this convention:

`rd=storage_node_hostname:device_name`

Administrator Privileges

The administrator on the storage node only needs to be on the Sun StorEdge EBS server's Administrator attribute to run `jbconfig` and `scanner` on the storage node. For normal operation, such as after `jbconfig` is run on the storage node, the administrator on the storage node can be removed from the server's Administrator list.



Caution – It is recommended that you do not attempt to convert an existing Sun StorEdge EBS server to a Sun StorEdge EBS storage node, as there is a potential for conflict when the resource database, media data base, and client file indexes from the retired server are merged into databases and client file indexes of the controlling Sun StorEdge EBS server.

Installing the Software on a New Storage Node

When you install storage node software to create a new storage node, you complete following:

1. First update the server to the same Sun StorEdge EBS release.
2. Run the `jbconfig` program on the storage node computer to add a Jukebox resource.

The Jukebox resource on the server displays values that the `jbconfig` program automatically selects based on the type of autochanger you configure.

Installing an Update Only on the Sun StorEdge EBS Server

If you choose to update only the Sun StorEdge EBS server with the new storage node software, any existing storage nodes will continue to function as usual, but they will not be able to use the features of the updated `nsrjb` program.

Installing an Update on an Existing Storage Node

If you have existing storage nodes and plan to update both the server and storage node software to this release, first install the updated software packages on the server. Then install the software on the storage node.

▼ To Configure Autochangers and Silos as Storage Node Remote Devices

To configure a remote autochanger or silo device as a storage node remote device:

1. Verify that administrator at the storage node is listed in the Administrator attribute in the controlling Sun StorEdge EBS server's Server resource.
 - For UNIX storage nodes, use the *root@hostname* format, where *hostname* is the hostname of the storage node.
 - For Microsoft Windows storage nodes, use the *administrator@hostname* format, where *hostname* is the hostname of the storage node.
2. Run the `jbconfig` program on the storage node to define each device in the autochanger or silo:

```
jbconfig -s servername
```

▼ To Set Timeouts for Storage Node Remote Devices

To set a timeout for a mount request on a storage node remote device and redirect the save to another storage node:

1. From the Media menu, select Devices to open the Devices window.
2. From the View menu, select Details to display the hidden attributes.
3. 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.

Dedicated Storage Node

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 only back up its own local data.

Set this attribute at the time a device is created on a remote storage node. If the Dedicated Storage Node attribute is set to Yes, you will need a Dedicated Storage Node License for the storage node.

However, if the Dedicated Storage Node 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 will behave as if the Dedicated Storage Node attribute is set to the default No.

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

Storage Management Operations

The following sections detail the various tasks and components involved in the operation of storage volumes and their relationship with pools.

How the Sun StorEdge EBS Server Uses Volume Labeling

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. 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:

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.
 - An override volume name you entered.

If you relabel a recyclable volume from the same pool, 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.

How the Sun StorEdge EBS Server Uses Pools

Every volume belongs to a pool. Each pool has a matching label template associated with it. Volumes are labeled according to the rules of these label templates. Label templates provide a method for consistently naming and labeling volumes so there is no need to keep track of how many volumes you have used. You can either take advantage of the preconfigured pools and associated label templates that come with the Sun StorEdge EBS product, or create your own pools, label templates, and pool-template associations. By customizing label templates, you can better control over how the data storage is organized.



Caution – Placing a new label on a storable volume, or recycling a volume by relabeling it, makes any existing stored data unavailable for recovery.

How the Sun StorEdge EBS Server Selects and Mounts a Volume during a Backup

When a backup takes place, the Sun StorEdge EBS server searches for a volume from the appropriate pool to accept the data that needs to be written. The available volumes are those:

- Mounted on stand-alone devices.
- Accessible to the Sun StorEdge EBS server through auto media management or via the autochanger.

The results of backup and recovery procedures will vary depending on the availability and condition of the volume:

- If you try to back up files when an appropriate volume is not mounted, the Sun StorEdge EBS server requests a writable volume by displaying the following, or similar, message:

media waiting (or critical): backup to pool 'Default' waiting for 1 writable backup tape or disk

- If you are recovering data, the Sun StorEdge EBS server requests the volume name you need to mount in the device by displaying the following, or similar, message:

media waiting: recover waiting for 8mm 5GB *tape_volume_name*

- If you need more than one volume to recover your data, the Sun StorEdge EBS server displays all the volumes in the order they are needed. During the recovery process, the server requests the volumes it needs, one at a time.
- If more than one volume is mounted on the Sun StorEdge EBS server's devices, the server uses the following hierarchy to select a volume on which to write data:
 1. A mounted, appendable volume from the appropriate pool.
 2. A mounted, recyclable volume from the appropriate pool, but not currently in use.
 3. An appendable volume that is not currently mounted in the device, but is from the appropriate pool.
 4. A recyclable volume that is not currently mounted in the device, but is from the appropriate pool.

How the Sun StorEdge EBS Server Selects a Volume for Relabeling

Backup data is destined for a specific pool. When the data is ready to be written, the Sun StorEdge EBS server monitors the active devices to locate a volume from the appropriate pool as follows:

- If only one volume from the pool is mounted and appendable, the data is directed to that volume.

- If two volumes from the same pool are mounted on devices, the Sun StorEdge EBS server factors in the following to guide its volume selection:
 - Volume mode

The Sun StorEdge EBS server first writes to a mounted appendable volume from the appropriate pool if one is available. However, if the appendable volume is also recyclable — that is, if all the save sets on the appendable volume have the status recyclable — the volume is treated as recyclable. A mounted recyclable volume from the appropriate pool is the second choice. A mounted recyclable volume that belongs to a different pool will not be recycled. The choice is a new, unlabeled volume.
 - Volume label time (the time when the volume was labeled).
 - Volumes with the oldest label time are selected before volumes that were labeled more recently.
 - Number of sessions currently being written to the device.
 - If the Sun StorEdge EBS server cannot find a mounted volume from the appropriate pool, a mount request is initiated. If auto media management is not enabled or if the server only has stand-alone devices available, mount requests will continue to be generated until a volume is mounted and writing begins.
 - If an appropriate volume has been found in the autochanger and more than one empty drive is available, the Sun StorEdge EBS server mounts the volume on the lowest-numbered drive. If drives are full, the server will select a loaded idle drive and unload it to make room for the needed volume.
-

Volume Status

The Sun StorEdge EBS server reports on the status of volumes using values such as *Written*, *%Used*, *Location*, and *Mode*. This section defines some of the most common terms used in reports about volumes.

In displays that appear within the Sun StorEdge EBS Administrator program, the volume name is the same as the name that appears on the volume label. At the end of the name the following designations may appear:

- (A), which indicates an archive volume.
- (R), which indicates a volume that is considered *read-only*.

The value of *Written* always indicates the exact number of bytes written to the volume.

The value of *%Used* is an estimate of percentage used based on the total capacity of the volume, based entirely on the specified value of the Media Type of the Device resource. 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 media has occurred. This is in contrast to a Used value of 100%, which means that the value of *Written* is equal to or exceeds the estimate for this volume.

The volume location refers to an administrator-defined character attribute that describes a physical location meaningful within the user's environment, such as "2nd shelf, Cabinet 2, Room 42."

For information on changing volume modes, see ["To Change the Mode of a Volume" on page 241](#). For more information about save set status, see ["Client and Save Set Combination" on page 118](#).

TABLE 5-13 lists the Sun StorEdge EBS volume modes and their definitions.

TABLE 5-13 Volume Modes

Mode Value	Meaning	Description
appen	appendable	This volume contains empty space. Data that meets the acceptance criteria for the pool to which this volume belongs can be appended.
man	manual recycle	This volume is exempt from automatic recycling. The mode can only be changed manually.
(R)	read-only	The save sets on this volume are considered read-only. The mode can only be changed manually. See the following section for more information on the read-only mode.
recyc	recyclable	The save sets on this volume have exceeded both the browsing and retention policies. The mode can only be changed manually.
full	full	This 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 browse and retention policies.

Read-Only Mode

Changing the mode of a volume to read-only prevents new data from being 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 change 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 ["To Change the Mode of a Volume" on page 241](#).

Sun StorEdge EBS Volume Operations

If you try to back up files when a volume is not mounted, you receive one of three messages suggesting that you perform one of the following tasks:

- Mount a volume.
- Relabel a volume.
- Label a new volume.

If you are recovering files, the Sun StorEdge EBS server requests the volume name you need to mount. If you need more than one volume to recover the files, the server lists all the volumes in the order they are needed. During the recovery process, the server requests each volume it needs, one at a time. If you back up to an autochanger, the server automatically mounts volumes stored in the autochanger.

▼ To View Volume Information

To view volume information in the Sun StorEdge EBS Administrator program:

1. From the Media menu, select Volumes to open the Volumes window.

Note – You can also click the Volumes button.

2. Select the volume name whose save sets you want to view.

The Volumes window displays the information for the server's volumes, including:

- Volume name
- Amount of data written to each volume
- Percentage currently used
- Other topics

Check the Volumes window to view the contents of the media database and as a starting point for removing volumes, changing volume modes, and setting locations for your collection of volumes.

[TABLE 5-14](#) lists the volume-related categories displayed in the Volumes window.

TABLE 5-14 Volumes Window

Category	Description
Name	Volume name.
Barcode	Barcode label, if one exists.
Pool	Name of the pool to which the volume belongs.
Written	Amount of data written on the volume.
% Used	Percentage of the volume used compared to its estimated total capacity. When 100% appears in the Used column, the volume has met or exceeded its expected capacity. However, space may still be available because the end of the tape mark has not been reached.
Mode	Status or kind of volume. For more details on the type of volume modes, see “To Change the Mode of a Volume” on page 241 .
Location	Location of the volume (entered by the user as a reminder).

▼ To View Save Set Information

To view save set information in the Sun StorEdge EBS Administrator program, select Volumes from the Media menu.

Note – You can also click the Volumes button.

The Volumes window lists the save sets stored on a volume. This window displays information, such as:

- Client name
- Save set name
- Save set size
- Other categories

You can also change the status of a save set from here.

[TABLE 5-15](#) lists the categories related to save sets that appear in the Volumes window.

TABLE 5-15 Volumes Window with Save Set Information

Category	Description
Client	Name of the Sun StorEdge EBS client computer that created the save set.
Date	Date when the save set was created.
Level	Level of backup that generated the save set. This refers only to scheduled backups. For manual backups, the level is blank.
Status	Type of save set. For a listing of save set values and descriptions, see “Save Set Status Values” on page 118 .
Save Set	Pathname of the filesystem containing the save set. This column also includes clone information. If the save set has a clone, the pathname is marked <i>has clones</i> and the cloned save set is marked <i>clone save set</i> .

▼ To Mount a Volume

To mount a volume:

1. Insert a volume in the Sun StorEdge EBS server storage device or ensure that volumes are already present for the server to access.
2. Select the appropriate device in the Devices window of the Sun StorEdge EBS Administrator program.
3. From the Media menu, select Mount to mount the volume.

Note – You can also click the Mount button.

After you label and mount a volume, the volume name appears in the Devices list of the Sun StorEdge EBS Administrator program beside the pathname of the device.

To perform an unattended backup using a stand-alone device, you must mount labeled volumes in the device before leaving it unattended.



Caution – You can only use nonrewinding devices with the Sun StorEdge EBS server. If you use a rewinding device, the read/write head is repositioned at the beginning of the volume, and the previously backed-up data is overwritten. The pathnames for these devices must follow the BSD semantic rules, for example,

/dev/rmt/0mbn. The "b" in the pathname satisfies the BSD semantics requirement. If you use a file device, you must enter it as a directory path (the same as for other device types) rather than as just a filename.

▼ To Label and Mount a Volume in One Operation

If you have more than one storage device connected to the Sun StorEdge EBS server, you must first select the device you want to use for labeling from the Device list in the Sun StorEdge EBS Administrator program. Relabeling 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:

1. Insert an unlabeled or recyclable volume in the Sun StorEdge EBS server storage device, or ensure that these types of volumes are already present for the Sun StorEdge EBS server to access.
2. Select the appropriate device in the Devices list of the Sun StorEdge EBS Administrator program.
3. From the Media menu, select Label to open the Label window.

Note – You can also click the Label button.

4. Select a pool in the Pool attribute.

This action also generates a name for the Volume Name attribute. The following rules apply:

- If the volume is unlabeled, the Sun StorEdge EBS server assigns the next sequential label from the label template associated with the chosen pool. If you are relabeling a recyclable volume from the same pool, the volume label name and sequence number remain the same. However, access to the original data on the volume is destroyed, and the volume becomes available.
- If you do not select a pool for the volume that you are about to label, the Sun StorEdge EBS server automatically applies the label template associated with the Default pool.
- To create individual label names not associated with a template, edit the Volume Name attribute in the Label resource and enter a unique label name.

5. Select the Manual Recycle attribute if you want to manually recycle the volume.

If you enable the Manual Recycle attribute when you label a volume, 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.

6. Select the Mount After Labeling attribute and click OK. The Sun StorEdge EBS server automatically labels the volume, and then mounts the volume into the device.

After a volume is labeled and mounted in a device, the volume is available to receive data. Because the Sun StorEdge EBS label is internal and machine-readable, you should put an adhesive label on each volume that matches the internal volume label.

For information on using barcode labels, see [“Labeling Volumes with Barcodes” on page 404](#).

▼ To Label without Mounting

You may prelabel volumes without mounting them. To label a volume without mounting, follow the same procedures as for labeling and mounting in one operation, but deselect the Mount After Labeling attribute in the Label window.

▼ To Unmount a Volume

To unmount a volume:

1. In the Devices list of the Sun StorEdge EBS Administrator program, select the appropriate device.
2. Select Unmount from the Media menu.

Note – You can also click the Unmount button.

Once the volume is unmounted from the device, the Devices list in the Sun StorEdge EBS Administrator program displays that information.

▼ To Find a Volume Name

If the adhesive label on the volume is missing or illegible, you can determine its name by using one of these methods:

- Mount the volume in the storage device. The volume name appears next to the device name in Sun StorEdge EBS Administrator program. For more information, see [“To Mount a Volume” on page 238](#).
- Mount the volume in the storage device and enter the `nsrmm` command at the prompt to determine the volume’s name:

```
# nsrmm -pv -f devicename
```

Note – When you use the `nsrmm` command with the `-f devicename` attribute and the device name contains backslashes, you must use the backslash (escape) character with the backslash. Therefore, a device with the name `rd=sn1:\\.\\Tape0` would be written at the command line as: `rd=sn1:\\\\.\\Tape0`.

- Load the volume in the storage device and select Label from the Media menu. The name of the label appears in the Volume Name attribute. Click Cancel to close the dialog without relabeling the volume.

▼ To Change the Mode of a Volume

The mode of a volume determines whether the Sun StorEdge EBS server can write to it during a backup. A volume is automatically recyclable if all the save sets on the volume are recyclable. Recyclable save sets have passed the time period designated by the Retention policy for the client.

You can manually change the mode of a full or appendable volume to recyclable if you do not need the data on the volume. The Sun StorEdge EBS server then overwrites the data on the volume the next time it needs a volume for a scheduled or manual backup.

To change the mode of a volume:

1. If the volume is mounted, unmount the device:
 - a. In the Device list of the Sun StorEdge EBS Administrator window, select the appropriate device.
 - b. From the Media menu, select Unmount.
2. From the Media menu, select Volumes to open the Volumes window.

Note – You can also click the Volumes button.

3. From the Volume menu, select Change Mode and one of the following values:

- Appendable
- Read Only
- Recyclable

For a listing of the volume modes, see [TABLE 5-13 on page 235](#).

4. Select a Mode button and click OK.

The changed volume mode is now displayed in the Devices list of the Sun StorEdge EBS Administrator program.

Marking a Volume as Full for Off-Site Storage

When removing a volume from your autochanger 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 do this, enter the following at the command prompt:

```
# nsrmm -V valid -o full
```

Where *valid* is the volume identifier of the volume. When a volume is marked as full, it is also marked as read-only. The `mminfo` command can be used to determine the volume identifier. "For more information about these commands, refer to the man pages for `nsrmm` and `mminfo` or the *Legato Command Reference Guide* for more information about these commands.

Removing Bad Volumes from the Online Indexes

Remove a volume entry from the media database only if the volume has become physically damaged or unusable. When you select a volume for removal, the Sun StorEdge EBS server checks to see if the volume has a clone. If a clone is available, the volume entry is not removed from the media database. The Sun StorEdge EBS server accesses the clone volume instead of the original volume during recoveries or other actions requiring the volume. If a clone of the volume is not available, then the volume entry will be removed from the media database, and data located on that volume will no longer be browsable.

Removing volume-based entries in the online indexes has the added advantage of reducing the size of the online indexes. However, be cautious when deleting index entries, as it can limit or eliminate the ability to recover data from affected volumes.

The steps to remove volume-based entries are the same whether you are removing a bad volume, reducing the size of your indexes, or both. For instructions on how to perform these tasks, see [“Removing Volume-Based Entries from the Online Indexes” on page 98](#).

Maximum Number of Save Sets Per Volume

The maximum number of save sets the media database allows per volume is 270,000. When a volume reaches this number of save sets, it is marked as full. Prior to reaching this number, warning messages are logged to the *daemon.log* file and to the messages window to indicate that the volume will soon reach its maximum number of save sets.

In order to avoid reaching this limit, do *one* of the following:

- If the device is either a file type or advanced file type device, use staging to move save sets from the volume to alternate storage media. You can automatically stage save sets based on attributes such as the File System Check Interval and Max Storage Period, or you can manually stage save sets as needed. For more information about staging, see ["Save Set Staging" on page 244](#).
- Set the client's retention policy to a period that will prevent the 270,000 limit from being reached. For example, if you reach this limit in six months, setting the retention policy to a period of less than six months will remove the earliest database entries and prevent the maximum number of save sets from being reached. For information about setting retention policies, see ["Browse and Retention Policies" on page 120](#).

If the save set limit is reached and the volume is marked as full, you can remove some save sets from the volume using the **nsrmm** program:

```
# nsrmm -d -S ssid
```

where *ssid* is the save set ID for the save set to be removed.

Note – Using this command will remove all information about the save set from the media database, and will result in the save set being overwritten during subsequent backups.

After removing save sets from the volume, do one of the following:

- If the device is either a file type or advanced file type device, use the **nsrim** command line program to recover additional space on the disk and change the volume mode to *append*. Save set data stored in file type or advanced file type devices is automatically removed and space recovered after the next scheduled backup (**savegrp**) in which the **nsrim** program is run, or during staging when the Recover Space operation is performed.
- For all device types, manually change the volume mode to *append*. [For instructions, see ["To Change the Mode of a Volume" on page 241](#).

Save Set Staging

Save set staging is a process of transferring data from one storage medium to another. Staging reduces the time it takes to complete a backup by directing the initial backup to a high-performance file type or `adv_file` device. The data can then be staged to a storage medium, freeing up the disk space. Any volume type, such as Default, Index Archive, or Default Clone, can be staged.

You can create, edit, and delete staging policies as you can for other Sun StorEdge EBS resources. As part of the client setup, the use of a staging device can be selected for each pool (or set of pools) for backup, archive, and migration. The files are retained for the specified time in the disk staging pool before being moved to a tape device or optical disk. Any number of devices can be in the staging pool, and a save set can be staged as many times as required, for example to disk, to optical disk, to a local tape device, and to a remote tape device. Also, a volume can be staged to a second volume, and then that data on the second volume can be staged back to the first volume.

The staging process is driven by one of the following events:

- As part of an automatic process, such as keeping the save set for 30 days on the staging device before staging the data to the next device.
- As part of an event driven process, such as when available space in the staging pool drops below a set threshold. When this happens, the oldest save sets are moved first, until available space reaches the upper threshold that has been set.
- As part of an administrator initiated process, such as allowing the administrator to either reset the threshold and kick off staging or manually select save sets to stage.

When you enable a staging policy, the Sun StorEdge EBS server creates a clone of the save set you specify on a clone volume of the medium you specify. If the save set was stored on a filesystem device, the save set is deleted from the filesystem to free

the space. The Sun StorEdge EBS server tracks the location of the save set in the media database. The retention policy for the save set does not change when the data is staged.

Note – If the file type volume is on a storage node that is running Sun StorEdge EBS software 6.1 or earlier, the tape is not automatically marked appendable after the staging operation.

Staging Data from a File Type or `adv_file` Device

Staging is particularly well suited for data that has been backed up on file type or `adv_file` devices. Staging allows the occupied disk space on file type or `adv_file` devices to be reclaimed so that the disk space can be used for other purposes. Use staging to move the data to more permanent storage, such as an optical or tape volume, or even another, lower-priority device. Staging also allows data to be moved off the device outside the backup period, ensuring that sufficient disk space is available for the next backup session.

Staging and Cloning from the Command Line

Staging a save set from the command line works differently than staging a save set from the Sun StorEdge EBS Administrator program. When staging from the Sun StorEdge EBS Administrator program, the save sets belonging to a single device, as selected in the Staging resource, are staged. When staging a save set from the command line, specifying the save set IDs determines which save sets are staged.

When a save set is cloned, the cloned save sets are given the same save set ID as the original save set. Therefore, when staging a save set from the command line, the Sun StorEdge EBS software stages all the save sets with the specified 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 make sure that the original save set and all clones are not removed, specify a clone ID with the save set ID when staging from the command line. An example of the syntax is:

```
# nsrstage -m -S ssid/cloneid
```

▼ To Create a Staging Policy

Note – Before creating a staging policy, make sure you have configured all appropriate devices. Otherwise, no devices will be listed in the Devices attribute.

To create a staging policy from the Sun StorEdge EBS Administrator program:

1. From the Customize menu, select Staging to open the Stage window.
2. From the View menu of the Stage window, select Details to display hidden attributes.
3. Click Create.
4. In the Name attribute, enter the name of the staging policy.
5. In the optional Comment attribute, enter a description of the staging policy.
6. In the Enabled attribute, select Yes to have the staging policy start functioning right away.

You have the option of enabling or disabling staging at any time.

7. 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 filesystem partition that the device is on. Staging continues until the lower mark is reached (see [Step 8](#)).

Note: The High-Water Mark must be greater than the Low-Water Mark.

8. In the Low-Water Mark (%) attribute, enter or select a number.

This is the point at which the staging process should stop, measured as the percentage of available space used on the filesystem partition that the device is on.

9. From the Save Set Selection attribute, select a save set. This attribute allows you to determine which save sets are selected for staging.
10. From 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.

11. In the Devices attribute, specify the file type and adv_file devices controlled by this policy as the source device of 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.

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 filesystem 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. Once a save set is staged from the selected device, its media database entry is removed, and the space once occupied on the volume can be freed and recovered from disk.
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 filesystem check operations.
17. In the File System Check Interval Unit attribute, select Minutes or Hours.
18. In the Start Now attribute, select one of the choices to apply the selected operation on all devices associated with this policy.

The operations available are:

- Recover space: Recover space for save sets that have no entries in the media database and delete all recycled save sets.
- Check file system: Check filesystem and stage data if necessary.
- Stage all save sets: Stage all save sets to the destination pool.

Note – The choice you make takes effect immediately after clicking Apply. If you do not want any action to start now, leave this attribute alone. After the staging operation has completed, this attribute is returned to its default setting (blank).

19. Click Apply.

▼ To Edit a Staging Policy

Note – You cannot change the name of a policy.

To edit a staging policy:

1. In the Sun StorEdge EBS Administrator program, select Staging from the Customize menu.
2. In the Stage attribute, select the staging policy to be edited.
3. Edit the attribute values of the staging policy and click Apply.

▼ To Delete a Staging Policy

To delete a staging policy:

1. In the Sun StorEdge EBS Administrator program, select Staging from the Customize menu.
2. In the Stage attribute, select the staging policy to be deleted.
3. Click Delete.
4. When prompted, click OK to confirm the deletion.
5. Click Apply.

Backup Groups and Schedules

This chapter provides information about Sun StorEdge Enterprise Backup Software software scheduling features, including automated group backups and customizable backup schedules.

The major sections in this chapter are:

- [“Overview of Sun StorEdge EBS Scheduling” on page 249](#)
- [“Backup Groups” on page 250](#)
- [“Group Backup Management” on page 257](#)
- [“Bootstrap Generation and Printout” on page 261](#)
- [“Backing Up Open Files” on page 262](#)
- [“Schedules” on page 263](#)
- [“Backup Levels” on page 273](#)
- [“Save Set Consolidation” on page 278](#)

Overview of Sun StorEdge EBS Scheduling

The following two Sun StorEdge EBS resources work in tandem to perform scheduled backups of client data:

- Group
- Schedule

Together, the Group and Schedule resources provide maximum flexibility in scheduling large numbers of differing client backups each day, with the least amount of disruption and inconvenience to your enterprise.

Group Resource

The Group resource (backup groups) specifies a time of day when a backup occurs. Typically, backup times are scheduled after regular work hours, such as late evenings and on the weekends. These times are preferred so that the heavy load from network traffic associated with backups does not interfere with regular work activities. Clients assigned to a backup group will have their specified save sets backed up during their group's scheduled time.

Schedule Resource

The Schedule resource provides the ability to specify the day of the week and date during the month when the backup occurs. The type and nature of the backup are also specified in this resource. Because scheduled backups are configured on a weekly or monthly basis, the type of backup scheduled for a particular day of the week is in turn duplicated for:

- The same day in each subsequent week.
- The particular day of a particular week once each month.

Backup Groups

Sun StorEdge EBS backup groups allow you to designate a starting time for a client's scheduled backup. You can:

- 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 taking the client's backup schedule into account when you decide which clients to include in a specific group.
- Stagger groups at different times of the day to avoid excessive strain on the network or the Sun StorEdge EBS server.

For more information about creating schedules that stagger the days that different clients perform full backups, see [“Schedules” on page 263](#).

Each client backing up to a Sun StorEdge EBS server must be assigned to a backup group, either a preconfigured group or a group you create yourself. In many cases, the Sun StorEdge EBS default group is sufficient for the clients, and creating additional groups is not necessary.

When you create a group, the group appears as a choice in the Client resource. The Client resource is where you associate a client with a group. The group name also appears as a selectable option in the Pool dialog box, which is where you associate a group with a pool. You can assign a client's save sets to more than one group.

How the Sun StorEdge EBS Server Uses Groups

When a Sun StorEdge EBS Client resource is created, the client is assigned to a backup group. The group determines the time when the Sun StorEdge EBS server begins to backup the client, along any other clients also assigned to the same group.

The clients in each backup group begin their automatic scheduled backups according to the start time of the group. You can balance the backup loads by taking the client's backup schedule into account when you decide which clients to include in a specific group.

Example: Using Groups to Back Up Multiple Clients

[FIGURE 6-1 on page 252](#) illustrates how the Sun StorEdge EBS server uses two backup groups to back up multiple client save sets. In the example, three client computers—mars, jupiter, and saturn—are part of the group named Weekly Full, which 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. Group Accounting 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 back up 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; it uses groups to determine what time clients start their scheduled backups.

Group:	Weekly Full	Group:	Accounting
Daily Start Time:	24:00 (Midnight)	Daily Start Time:	19:00 (7:00 p.m.)
Client:	mars	Client:	mercury
Save Set:	All	Save Set:	All
Schedule:	Full on Monday Incremental, all other days	Schedule:	Full on 1st of month Level 5 on 15th of month Incremental on all other days
Client:	jupiter	Client:	venus
Save Set:	All	Save Set:	/usr/home
Schedule:	Full on Tuesday Incremental on all other days	Schedule:	Full on 1st and 15th of month Incremental on all other days
Client:	saturn		
Save Set:	All		
Schedule:	Full on Wednesday Incremental on all other days		




FIGURE 6-1 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 you need.

Key Group Attributes

[TABLE 6-1](#) displays some of the key attributes for the Group resource.

TABLE 6-1 Group Attributes

Attribute	Action
Autostart	This attribute controls 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.
Start Time	The Default group is preconfigured to start its daily backup at 3:33 a.m. This time can be changed.
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.
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 “To Set the Group Backup Time Interval” on page 255 . This is a hidden attribute.
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. This is a hidden attribute.
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. This is a hidden attribute.

Aborted Group Backups

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 NFS filesystem 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 comes back up, the backup continues and is completed.

▼ To Create a Group

To create a group:

1. In the Sun StorEdge EBS Administrator program, select Groups from the Customize menu to open the Groups window.
2. Click Create.
3. In the Name attribute, enter the name of the group.
4. (Optional) In the Comment attribute, enter a description of the group.
5. For the Autostart attribute, select Enabled.
6. In the Start Time attribute, enter a new time, unless you want to keep the default time of 3:33 a.m.

Note – When you select a start time for different groups, schedule them far enough apart so that one group has completed backing up before the next group starts.

7. If the NetWorker PowerSnap Module software is also installed, configure the PowerSnap attributes as follows:
 - a. For the Snapshot attribute, select True.
 - b. For the Snapshot Pool attribute, select the pool that was configured specifically for PowerSnap operations.
 - c. For the Snapshot Policy attribute, select a preconfigured or a custom policy.
8. In the Client Retries attribute, select the number of retries.

This value specifies the number of times the Sun StorEdge EBS server will attempt to back up a failed client.
9. In the Clones attribute, select Yes to automatically clone the group's backed-up data. You must have appendable volumes in the appropriate clone pool.

Note – If you select No, do not complete the clone-related step and skip to [Step 12](#).

10. In the Clone Pool attribute, select the appropriate clone pool for the cloned data.
11. If your environment supports it, select the appropriate migration clone pool in the Migration Clone Pool list.

12. In the Printer attribute, enter the name of the printer on which `savegrp` will print the bootstrap save set information.

13. Complete other attributes as necessary, and click Apply.

When you finish creating a backup group, the group appears as a choice in the Group attribute of the Client resource. This attribute assigns an individual client to a particular group.

▼ To Set the Group Backup 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 group backup 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 group backup time intervals:

1. In the Sun StorEdge EBS Administrator program, select Groups from the Customize menu to open the Groups window.
2. Select the appropriate Group.
3. Enter the Interval attribute value in the hh:mm format.

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

4. Click Apply.

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

▼ To Force an Incremental Backup

For groups that have more than one scheduled backup within a 24-hour period, you can use the Force Incremental attribute in the Group resource to instruct the Sun StorEdge EBS server to perform only 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 time of the first backup will be incremental, meaning that only changed files will be backed up regardless of the configured level. This may be advantageous when an environment requires frequent scheduled backups, but you want to minimize the impact on systems being backed up. 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. In the Sun StorEdge EBS Administrator program, select Groups from the Customize menu to open the Groups window.
2. From the View menu, select Details to display hidden attributes.
Because the Force Incremental attribute is a hidden attribute, you must set the Sun StorEdge EBS Administrator program to display hidden attributes before you can access and change the Force Incremental attribute.
3. Select the appropriate Group.
4. In the Force Incremental attribute, select Yes.

Note – To disable the force incremental function, select No for this attribute.

5. Click OK.

For more information about configuration group backups that occur more often than once every 24 hours, see [“To Set the Group Backup Time Interval” on page 255](#).

▼ To Edit a Group

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

To edit a group:

1. In the Sun StorEdge EBS Administrator program, select Groups from the Customize menu to open the Groups window.
2. Select the group to be edited.
3. Edit the attributes of the group and click Apply.

Note – To learn more about each attribute, see the Online Help.

▼ To Delete a Group

Note – You cannot delete the preconfigured Default group or any group currently applied to a Client resource.

To delete a group:

1. In the Sun StorEdge EBS Administrator program, select Groups from the Customize menu to open the Groups window.
2. Select the group to be deleted.
3. When prompted, click OK to confirm the deletion.
4. Click Apply.

Group Backup Management

This section describes how to accomplish the following tasks from the Sun StorEdge EBS Administrator program:

- [“To Start a Group Immediately” on page 258](#)
- [“To Stop a Group Immediately” on page 258](#)
- [“To Manually Restart a Group” on page 259](#)
- [“To View Group Control Details” on page 259](#)
- [“To Preview a Group Backup” on page 260](#)

▼ To Start a Group Immediately

When you start a scheduled backup group manually (on demand), each client in the group is backed up.

The level of a manual backup is based on the level that was defined in the clients's Schedule resource for that particular day.

To start a group immediately:

1. In the Sun StorEdge EBS Administrator program, select Group Control from the Server menu to open the Group Control window.

Note – You can also click the Group Control button.

2. Select the name of the group to be started immediately.
3. Click Start.
4. Click to start the group.

The Sun StorEdge EBS server overrides the group's scheduled start time and immediately backs up the clients in the group.

▼ To Stop a Group Immediately

To stop a group immediately:

1. In the Sun StorEdge EBS Administrator program, select Group Control from the Server menu to open the Group Control window.

Note – You can also click the Group Control button.

2. Select the name of the group to be stopped immediately.
3. Click Stop.
4. Click OK to stop the group.

The group stops immediately, any running save sessions are aborted, and the Sun StorEdge EBS Administrator program displays Not Finished in the Status attribute in the Group Control window.



Caution – If you manually stop a scheduled backup, or if the backup is interrupted for some reason (for example, due to a power failure), the Autorestart attribute is automatically set to Disabled. To ensure that the scheduled backup restarts automatically, set the Autorestart attribute to Enabled.

The Autorestart attribute is a hidden attribute in the Group resource. For more information about editing the Group resource, see [“To Edit a Group” on page 256](#). For more information about displaying hidden attributes, see [“Hidden Attributes” on page 42](#).

▼ To Manually Restart a Group

To manually restart a group:

1. In the Sun StorEdge EBS Administrator program, select Group Control from the Server menu.

Note – You can also click the Group Control button.

2. Select the name of the group to be restarted.
3. Click Restart.
4. Click OK to restart the group.

After you initiate the restart of a group backup, the Sun StorEdge EBS server resumes the scheduled backup for the group and displays Running in the Status attribute in the Group Control window.

▼ To View Group Control Details

The Group Control Details window displays save sets that:

- Are in the process of backing up.
- Have completed their backups.

Open this window to determine which client's save sets were backed up successfully and which save sets may have failed.

To view group control details:

1. In the Sun StorEdge EBS Administrator program, select Group Control from the Server menu.

Note – You can also click the Group Control button.

2. Select the name of the group with the details to be viewed.
3. Click Details.

The Group Control Details window displays the status of client save sets in the backup process in one of three message panes:

- Pending Save Sets — Displays the client save sets that have not yet been backed up.
- Completed Save Sets — Displays the client save sets that the Sun StorEdge EBS server has successfully backed up.
- Failed Save Sets — Displays the client save sets that the Sun StorEdge EBS server did not back up (typically because of a computer or network crash).

If the server was unable to back up save sets for a client in the group, the backup may have failed for one of the following reasons:

- The Sun StorEdge EBS server failed.
- The Sun StorEdge EBS client failed.
- Licensing problems.
- The network connection failed.
- Customized backup failed.

▼ To Preview a Group Backup

Sun StorEdge EBS software allows you to simulate a backup for a specific group using the Group Control Preview feature. This feature runs a probe step on each client and outputs the following information:

- Filesystem
- Level
- Pool

This feature may be useful to identify potential problems before the Sun StorEdge EBS server runs an upcoming group backup. The simulated results are then displayed in the Group Control Details window.

To preview a group backup:

1. In the Sun StorEdge EBS Administrator program, select Group Control from the Server menu.

Note – You can also click the Group Control button.

2. In the Groups list, select the appropriate group.
3. Click Preview to open the Group Control Details window.

The simulated results of the backup are displayed in the Completed Save Sets list of the Group Control Details window.

Savegroup Completion Report

When the backup is completed, the Sun StorEdge EBS server generates a Savegroup Completion Report about the success of the scheduled backup. The server sends the root user an automatic notification and displays the same information in the Sun StorEdge EBS Administrator program.

Bootstrap 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 *Legato Networker Disaster Recovery Guide*.

By default, 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.

▼ To Save the Bootstrap to a File

To save the bootstrap to a file:

1. In the Sun StorEdge EBS Administrator program, select Notifications from the Customize menu to open the Notifications window.
2. In the Notifications attribute, select Bootstrap.

The Action attribute displays the `lp` command with appropriate options. For example, on Solaris the following value is displayed:

```
/usr/bin/lp -s -c -t bootstrap -d printer_name
```

3. To save the bootstrap information to a file, rather than sending it to a printer, change the Action attribute to direct the bootstrap to a file. For example, on Solaris, enter:

```
/bin/cat >> /directory/filename
```

▼ To Send the Bootstrap File by E-mail

To send the bootstrap file to a user ID by e-mail, change the Action attribute to use the `mail` program. For example, on Solaris, enter:

```
/usr/ucb/mail -s nwserver_bootstrap user@hostname
```

System Console Log

The UNIX system log displays messages passed from the Sun StorEdge EBS server. When the Sun StorEdge EBS server is installed, it adds lines to the configuration log file (`syslog.conf`). These tell the system log facility what types of notices to direct to which file or user. For example:

```
daemon.notice /dev/console
daemon.notice /nsr/logs/messages
daemon.notice operator
local0.notice /nsr/logs/summary
local0.alert root, operator
```

Backing Up Open Files

You can back up most open files using the Sun StorEdge EBS product. If an open file changes during a scheduled backup, the Sun StorEdge EBS server recognizes that it is changing and backs up the file. The server also checks to see if the size of the file

changed. If the file size did change, the server reports the discrepancy. Files that change during backup are displayed in a warning message in the Group Control Details window similar to this:

warning: *file_name* changed during save

To make sure the changed file gets backed up, you can do one of the following:

- Restart the backup group.
- Back up the file manually.

Backing Up Open Files with VSS

In Sun StorEdge EBS release 7.1, 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. 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.

If a [writer](#) is available, enabled, and a VSS client license exists, Sun StorEdge EBS software defaults to backing up data using VSS technology. If a VSS client license does not exist or a writer does not exist for an application, Sun StorEdge EBS software backs up data using the [legacy method](#). During a given backup for an individual client, either the VSS method or the legacy method is used, but not both. For more information about VSS, see the *Legato NetWorker Administrator's Guide, Microsoft Windows Version*.

Schedules

The Sun StorEdge EBS server backs up each client computer across a network according to a backup schedule. A client's backup schedule tells the Sun StorEdge EBS server what level of backup operation (for example, full or incremental) to perform on a given day. For instance, on Fridays you might perform a full backup on a client and the rest of the week perform incremental backups.

Schedules can be very simple or very complex, depending on the needs of the environment. All clients can share the same schedule, or each client can have its own unique schedule. The scheduling feature allows you to create customized schedules that you can apply to client save sets. For more information about client configuration, see [“Client Configuration” on page 104](#).

This section describes the following:

- [“How the Sun StorEdge EBS Server Uses Schedules” on page 264](#)
- [“Preconfigured Sun StorEdge EBS Schedules” on page 265](#)
- [“Backup Cycles” on page 266](#)
- [“Schedule Planning Considerations” on page 268](#)
- [“Scheduling Large Client Filesystems” on page 269](#)
- [“Key Schedule Attributes” on page 269](#)
- [“Resource Configuration Order for Schedules” on page 271](#)
- [“To Create a Schedule” on page 271](#)
- [“To Edit a Schedule” on page 272](#)
- [“To Delete a Schedule” on page 272](#)
- [“Overriding a Client’s Regular Backup Schedule” on page 272](#)

How the Sun StorEdge EBS Server Uses Schedules

The Sun StorEdge EBS server uses a client’s backup schedule to determine what level of backup operation to perform on a given day for specified save sets. The time of day the backup operation begins is determined by the group with which the client save sets are associated.

The type and scope of the backup is determined by the [backup level](#) you specify. These levels can be set to back up a client’s entire filesystem, or only data that has changed since the last full backup. For more information on how backup levels work, see [“Backup Levels” on page 273](#).

The Schedule resource allows you to customize each client’s backup schedule. For example, some clients may have data to be backed up at the full level every three days, with incremental backups in between. Other clients may have less critical data that only needs a full backup once a month, with incremental backups or level 1–9 backups in between.

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, you can create new ones 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. However, a schedule’s name cannot be modified.

[TABLE 6-2](#) describes the preconfigured schedules.

TABLE 6-2 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, and 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. The following items are demonstrated:

- [“Example: Weekly Backup Cycle” on page 266](#)
- [“Example: Default Schedule with Multiple Clients” on page 266](#)
- [“Example: Staggered Weekly Schedules for Multiple Clients” on page 267](#)

Example: Weekly Backup Cycle

[FIGURE 6-2](#) 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 6-2 Weekly Backup Cycle

You can 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 6-3](#) illustrates how the Default schedule works for three clients.




FIGURE 6-3 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 278](#).

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 6-4 illustrates how to use a staggered backup schedule for multiple clients.




FIGURE 6-4 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.

Schedule Planning Considerations

Sun StorEdge EBS software makes it easy to set up backup schedules. Deciding which schedules best fit an environment, however, requires some 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, you must 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 on data life cycle management, see [“Browse and Retention Policies” on page 120](#).

Scheduling Large Client Filesystems

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 the backup takes.

You can 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 filesystems individually, add and configure the same client several times by addressing the different filesystems in the Client resource. For example, configure the first Client resource to back up one filesystem () with a single backup schedule in one group. Then, configure the second Client resource to back up another filesystem () with a second backup schedule in another group.



Caution – When you create separate backup schedules and explicitly list save sets, any files or filesystems not included in that list are omitted from backup. This includes any new disk volumes that are added to the system. You can avoid this risk of omission by entering the value All in the Save Set attribute. This way, the Sun StorEdge EBS software automatically adds the new disk volumes to the backups.

Key Schedule Attributes

To create a customized backup schedule, you must define values in the Schedule resource in the Sun StorEdge EBS Administrator program.

TABLE 6-3 describes the key attributes contained in the Schedule resource.

TABLE 6-3 Key Schedule Attributes

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 simple, descriptive name; for example, Monday Full. You cannot change the name of an existing schedule. For example, if you want to change the schedule <i>Full Every Friday</i> to <i>Full Every Monday</i>, you must delete the <i>Full Every Friday</i> schedule, and then create a <i>Full Every Monday</i> schedule.
Period	The period refers to how often you want a full backup to run. You can 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.
Prev Month	Click to display the previous month's backup schedule in the calendar.
Next Month	Click to display the next month's backup schedule in the calendar.
Calendar	Displays the days of the month and the backup level scheduled for each day. To set the backup level for a selected day,. In addition to full and incremental backups, you can set intermediate backup levels. These allow you to help maintain server efficiency and manage the number of volumes needed for a recovery. You can include one or more of the following levels in a backup schedule: • Full • Incremental • Level [1 – 9] • Consolidated • Skip The 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 273 and “Save Set Consolidation” on page 278.

Note – The Force Incremental attribute, located in the Group resource, 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.

Resource Configuration Order for Schedules

If a customized schedule is used, the schedule must be created before it can be included in the Schedule attribute of the Client resource. The start time for your scheduled backup is determined by the Group resource selected in the Group attribute of the Client resource. The length of time that the data is available for browsing or recovery is determined by the browse and retention policies selected for the Client resource, rather than by the schedule.

▼ To Create a Schedule

To create a schedule:

1. In the Sun StorEdge EBS Administrator program, select Schedules from the Customize menu.
2. Click Create.
3. In the Name attribute, enter the name of the schedule.
4. (Optional) In the Comment attribute, enter a description of the schedule.
5. For the Period attribute, select Week or Month. The selected value changes the way levels set for each day are duplicated:
 - Week: Levels selected for each day of the week are duplicated to the corresponding day of each subsequent week. For example, if you select a full backup for a Friday, every Friday is configured to have a full backup.
 - Month: Levels selected for each date in the month are duplicated to the corresponding date in each subsequent month. For example, if you select a full backup for the first of the month, every month has a full backup on the first of the month.
6. Select a backup level for each day of the period (week or month) displayed in the calendar by clicking on the day and selecting the backup level from the menu provided. Click Previous Month or Next Month to see the backup calendars for preceding or future months.

For more information on how backup levels work, see [“Backup Levels” on page 273](#).

7. Click Apply.

You can set up as many Sun StorEdge EBS backup schedules as you want, as long as you give each one a unique name.

To assign a backup schedule to a client, you must enter that value in the Schedule attribute of the Client resource configuration.

▼ To Edit a Schedule

To edit a schedule:

1. In the Sun StorEdge EBS Administrator program, select Schedules from the Customize menu.
2. Select the schedule to be edited.
3. Edit the attributes and click Apply.

▼ 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. In the Sun StorEdge EBS Administrator program, select Schedules from the Customize menu.
2. Select the schedule to be deleted.
3. Click Delete.
4. When prompted, click OK to confirm the deletion.
5. Click Apply.

Overriding a Client's Regular Backup Schedule

You can use the Level and Schedule attributes in the Group resource to override a client's regular backup schedule. For example, one evening you might want to run a full backup on all the clients in a group, regardless of the clients' regular backup schedules. The entry you make in the Level attribute overrides the backup level setting for every client in the group.

Alternatively, you might want a group of clients to 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 preconfigured setting), the clients follow their individual backup schedules.

The Level and Schedule attributes are hidden attributes. For information about displaying hidden attributes, see [“Hidden Attributes” on page 42](#).

Backup Levels

Because it may not be practical or efficient for you to run full backups every day, the Sun StorEdge EBS server allows you to specify the level of the backup operation performed during its automatic, scheduled backups. Limiting how often you perform a full backup can help maintain server efficiency, while still ensuring that the data is protected. Different backup levels allow 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 6-4](#) describes the five kinds of backup levels supported by the Sun StorEdge EBS software for filesystem data:

TABLE 6-4 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 subsequently merges these changes with the last full backup. For more information, see “Save Set Consolidation” on page 278 .
Skip	Skips the scheduled backup. For example, you may want to skip a backup on a holiday if you know that no one will be available to change or add more media volumes.



Caution – A level skip backup does not back up any data. However, the Sun StorEdge EBS server still contacts the client for the scheduled backup at the Start Time indicated for the backup group. The server's `savegrp` program generates a Savegroup Completion Report that shows that a level skip backup was performed and no data was backed up. Any problem that could cause the Sun StorEdge EBS server to wait indefinitely for the client to respond has the same effect during a level Skip backup as for other level backups, even though no data is actually backed up.

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 maintain control over the number of volumes you use. A carefully planned backup strategy enables you to recover everything to disk using a minimum number of volumes. The fewer volumes you need to recover from a disk crash, the less time you must spend restoring the disk.

You can also control the size and time it takes to back up data by using directives, which compress and eliminate unnecessary data from the backups. For example, you can use a directive that tells the Sun StorEdge EBS server to skip certain files or filesystems when performing a backup. For more information on directives, see [“Directives” on page 144](#).

The following three examples illustrate the concept of how Sun StorEdge EBS backup levels work, with the requirements for recovery in the event of data loss.

Example: Backup Levels (Part 1)

As shown in [FIGURE 6-5 on page 275](#), 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 6-5 Backups for October 1 through October 7

Example: Backup Levels (Part 2)

FIGURE 6-6 continues the example illustrated in **FIGURE 6-5** 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 6-6 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 6-7 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 6-7 Backups for October 1 through October 15

How the Sun StorEdge EBS Server Uses Backup Levels

A backup schedule defines what level backup the Sun StorEdge EBS server should perform on a given day during a backup cycle. You can apply one or more of these backup levels to customize a backup schedule. If you plan 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 essential difference, however, 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 6-5 lists advantages and disadvantages of each backup level.

TABLE 6-5 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 manner:

- 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 from regular files and other data. Knowing how the online indexes are backed up is particularly important during disaster recoveries. [TABLE 6-6](#) compares the level at which the Sun StorEdge EBS server backs up regular files and the online indexes.

TABLE 6-6 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

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. Essentially, save set consolidation is a process rather than an attribute; it describes the process of merging a level 1 backup with an existing full level 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.

As such, you cannot find *consolidation* as a backup level in most of the various resources where Sun StorEdge EBS backup levels appear. For example, you cannot find *consolidation* as a potential backup selection in the Pool resource. The one

exception is the Schedule attribute, where you can use the default consolidation schedules provided, or create a consolidation schedule by overriding a scheduled level with the special level `c`.

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 takes place entirely on the server. The burden of backups is shifted entirely to the server, reducing Client resource use and network traffic. This shift to the server also results in more frequent level saves, resulting 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 there are no existing full backups the first time a consolidation backup begins, the consolidation backup defaults to a level full. Once a full backup exists, subsequent consolidation backups are performed as a level 1 followed by consolidation with the existing full.

Requirements

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

Consolidation of save sets larger than 2 GB is not supported.

Save Set Consolidation Usage

You should 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 filesystems 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 – Do not consider save set consolidation as a performance improvement without considering these conditions. Under the appropriate conditions, save set consolidation can be used to 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.

In addition, 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 instance, save set consolidation will not produce a measurable benefit.
- The filesystems 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 filesystem.
- It is cost-prohibitive to allocate three (or the minimum of two) tape drives for the exclusive use of the server while it consolidates the full backup.



Caution – Even if a consolidated backup cannot be completed because of a system crash, tape drive errors, or other problem, data integrity is ensured. The consolidated backup is aborted. The transactional log assists save set consolidation in putting back changes made to the online index. In most instances, the Sun StorEdge EBS server only makes online index changes when data has been written to tape.

Invoking Save Set Consolidation

You can schedule save set consolidation through the Sun StorEdge EBS Administrator program, or invoke it from the command prompt.

▼ To Schedule Save Set Consolidation

To schedule save set consolidation through the Sun StorEdge EBS Administrator program:

1. In the Sun StorEdge EBS Administrator program, select Schedules from the Customize menu.
2. Select the schedule to which you want to apply save set consolidation.
3. In the calendar, click the day save set consolidation should occur.
4. Select Consolidate from the pop-up menu.

The letter *c* appears for that day, to signify the newly created consolidated backup. The Sun StorEdge EBS server will invoke the save set consolidation backup on this date.

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

▼ To Invoke Save Set Consolidation from the Command Prompt

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 *c* 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 further information, refer to the `savegrp` and `nsrssc` man pages refer to the *Legato Command Reference Guide*.

▼ 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 most recently backed up.

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

1. Configure a Group resource for consolidation backups. See [“To Create a Group” on page 254](#).
2. Configure a Pool resource for consolidation backups. See [“To Create a Pool” on page 176](#).
3. Select the name of the Group resource you created in [Step 1](#) as the value for the Pool resource’s Groups attribute.
4. Edit the Client resource for each client for which you want to perform consolidated backups, and assign the client to the group that you created for consolidation backups.

Example: Consolidation of Level 1 and Full Backups

[FIGURE 6-8 on page 283](#) illustrates how the Sun StorEdge EBS server merges a level 1 backup with an existing full level backup to “consolidate” both to create another full level backup.

On October 16, the level “c” automatically performs a level 1 backup, which backs up all the data that has changed since the last full backup. As part of its consolidation process, the Sun StorEdge EBS server merges this level 1 backup with the last full backup (created on October 1) and builds a new full level backup. To recover from a disk crash on October 17, you need one volume—the full volume that was created using the consolidation process on October 16.




FIGURE 6-8 Consolidated Backup

Backup and Recovery

This chapter describes the method for backing up data manually, rather than using the scheduled group backup method. It also details the various methods for recovering backed-up data, and how to direct recovered data to a specific location.

The main sections in this chapter are:

- [“Manual Data Backups” on page 285](#)
- [“Data Recovery” on page 287](#)
- [“Disaster Recovery” on page 294](#)
- [“Directed Recovery” on page 294](#)
- [“Save Set Recovery” on page 303](#)
- [“Comparison of Recovery Methods” on page 309](#)
- [“Recovering Online Indexes” on page 312](#)
- [“The scanner Program” on page 316](#)

Manual Data Backups

The Sun StorEdge Enterprise Backup Software server backs up client data regularly using scheduled backups. Scheduled backups are the preferred backup method because they occur automatically, and can be recovered more easily.

Manual backups are intended for users to make quick backups of a few files. They do not generate bootstrap files or backup indexes, so they cannot replace scheduled backups. You can manually start a scheduled backup group at any time using the Start Group option in the Sun StorEdge EBS Administrator program. For more information about starting a scheduled backup group at any time, see [“To Start a Group Immediately” on page 258](#). For a complete description of scheduled backups and groups, see [Chapter 6](#).

There may be times, however, when you want to back up one or more data items to the Sun StorEdge EBS server right away, instead of waiting for the scheduled nightly backup or performing an entire group backup. For example, if you have just finished working on a spreadsheet, you might like to back up the file immediately.

You can designate the entire filesystem, specific directories, or individual files for manual backup. If you select a directory for backup, all of its nested subdirectories and files are also selected for backup. To exclude certain files from manual backup, you must unselect them.

You can also search the filesystem for files you would like to back up by selecting Search from the View menu to open the Search window, then entering the name or a partial name of the item you are seeking.

▼ To Start a Manual Backup

To start a manual backup by using the `nwbackup` program:

1. Start the `nwbackup` program by entering the following command at the prompt:

```
# nwbackup
```

The `nwbackup` program provides a graphical interface through which you can perform manual backups. For general information about the `nwbackup` program, see [“The `nwbackup` Program” on page 45](#).

2. Select each directory and/or file you intend for manual backup by performing one of the following:

- Select the directory or file and click the Mark button.

To clear an item, click the Unmark button.

- Select the directory or file and select Mark from the Mark menu.

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

- Select the check box next to each directory or file listed.

To clear a check box, click the check mark next to the selected item.

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

3. Click the Start button to begin the manual backup.

When the Sun StorEdge EBS server has successfully finished the backup, the `nwbackup` program displays the following (or a similar) message:

```
Backup completion time: 2-15-03 3:27p
```



Caution – There are certain kinds of corrupt files or errors on computer disk volumes that Sun StorEdge EBS software cannot detect during a backup. It backs up the corrupt data. It is a good idea to run diagnostic programs regularly on your computers to correct disk volume errors.

Backing Up Data Manually from the Command Prompt

You can also perform a manual backup 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, refer to the `save` man page.

Data Recovery

Recovery of backed-up data is performed using the `nwrecover` program. This program is similar to the `nwbackup` program. For information about the `nwrecover` program, see [“The nwrecover Program” on page 46](#).

When the Sun StorEdge EBS server backs up a file to media (for example, magnetic tape), it creates an entry for the file in the client file index. The Sun StorEdge EBS server uses this index to map the backed-up files to volumes. If the computer has not been backed up yet, the client file index will be empty, and therefore no backups will appear in the `nwrecover` program’s main window.

By using `nwrecover`, you can browse backed-up files and filesystems from Sun StorEdge EBS clients, and recover files to the state they were in when they were backed up. The versions of files and filesystems available in `nwrecover` are determined by the time period specified in the browse policy. For more information about the browse policy, see [“Browse and Retention Policies” on page 120](#).

If remote access is enabled, users may also be able to recover files from other client computers. For more information on remote access, see [“Directed Recovery” on page 294](#).

When recovering data to a filesystem where other versions of the same data reside, there is the likelihood of a file conflict. You have three options in such instances:

- Determine what course of action to take every time a conflict occurs during a recovery.
- Instruct the Sun StorEdge EBS server in advance as to what action to take for all conflicts.
- Redirect the recovered data to a different directory where there is no danger of conflicts.

▼ To Recover Data

You can designate the entire filesystem, specific directories, or individual files for recovery. Note that if you select a directory for recovery, all of its nested subdirectories and files are also selected for recovery. To exclude certain files from recovery after selecting an entire directory, you must unselect them.

To recover data:

1. Start the `nwrecover` program by entering the following command at the prompt:

```
# 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 filesystem.

2. Mark each directory and/or file to be recovered by performing one of the following:

- Select the directory or file, and click the Mark button.
To clear an item, click the Unmark button.
- Select the directory or file and select Mark from the Mark menu.
To clear an item, select Unmark from the Mark menu.
- Select the check box next to each directory or file listed.
To clear a check box, click the check mark next to the selected item.

When you mark a directory or file for recovery, a check mark appears next to that item.

Note – To recover data from a scheduled backup other than the most recent one, you must change the browse time. See [“Recovering Data from a Previous Backup” on page 290](#).

3. 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 in the `nwrecover` program for each recovered file, so 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 with the file you are recovering, the Sun StorEdge EBS server prompts you to choose one of several workarounds. For more information about naming conflicts, see [“Resolving Recovered Data Conflicts” on page 291](#).

If the recovery is successful, the `nwrecover` program displays the following (or a similar) message:

```
Received 1 file(S) from NSR server server

Recover completion time: 5-30-03 2:03p
```

Recovering Data from the Command Prompt

You can also perform a recovery using the `recover` program, an interactive command line interface. 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 /dev/rmt/0cbn
Requesting 1 file(s), this may take a while...
./myfile
Received 1 file(s) from NSR server `jupiter'
Recover completion time: Tue Jan 21 08:33:04 2003
recover> quit
```

For more information, refer to the `recover` man page.

Recovering Data from a Previous Backup

From the `nwrecover` program, you can browse the entries for each backed-up file listed in the client file index. By changing the browse time, you can view the entries for the backed-up files for a particular period. You might want to do this if you need to:

- Look for data that has been deleted from the filesystem and may not be present in the most recent backup.
- Retrieve an earlier version of existing data.

The browse policy determines how long data is retained and whether you can browse and recover this data. For example, to recover a lost file that was backed up six months ago requires the appropriate browse policy. For more information, see [“Browse and Retention Policies” on page 120](#).

▼ To Change the Browse Time

Note – Before changing the browse time, 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.

To change the browse time:

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 time format.
4. Click OK.
5. Select the files or directories you want to recover as described in [“To Recover Data” on page 288](#).

Note – If there are no entries for files backed up at the time selected in the Browse Time field, the Sun StorEdge EBS software displays files that were backed up at the closest preceding time. However, if the time selected in the Browse Time field is earlier than the first entry in the client file index, the Sun StorEdge EBS software warns you that there are no files backed up at the specified time.

Resolving Recovered Data Conflicts

By default, the Sun StorEdge EBS server attempts to recover data to its original location, using the original file and directory names. However, if another file or directory with the same name is present in that same location, a file recovery conflict occurs. You can instruct the Sun StorEdge EBS software to respond to these recovery conflicts either automatically or on an file-by-file basis.

When you start a recovery using the `nwrecover` program, the Conflict Resolution window opens and asks if you want to be prompted when a conflict occurs.

- If you select Yes, the Sun StorEdge EBS server prompts you to select a resolution option each time there is a conflict.
- If you select No, then you must select a conflict resolution option that the server will follow each time there is a conflict during the recovery.

[TABLE 7-1](#) lists the naming conflict resolution options available during a recovery.

TABLE 7-1 Naming Conflict Resolution Options

Option	Action
Rename Recovered File	Renames the recovered files by appending each conflicting name with <code>.R</code> .
Discard Recover File	Discards the recovered file in order to prevent the Sun StorEdge EBS server from overwriting an existing file.
Overwrite Existing File	Replaces the file on the filesystem with the recovered version.

You can also relocate recovered data to a directory where you know there are no conflicts, or you can create a new directory using the relocate feature. However, even if you relocate recovered data to a new and/or empty directory, the Sun StorEdge EBS server still prompts you to enter conflict resolution information. For more information on this feature, see [“Relocating Recovered Data” on page 292](#).

Relocating Recovered Data

When you recover a file, by default the Sun StorEdge EBS server attempts to copy the file to its original location. Unless you relocate the recovered file, the server attempts to place it in the file’s original directory. If an existing file on the local disk has the same name as the one being recovered, the server prompts you for instructions on how to recover the file (rename, overwrite, or discard).

Alternatively, you can instruct the Sun StorEdge EBS server where to relocate the recovered file prior to the recovery.

▼ To Relocate Recovered Data

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 – You must provide the full pathname of the relocation directory.

3. Click OK.
4. Perform a data recovery as described in [“To Recover Data” on page 288](#).

▼ To View Versions of a Directory or File

You can use the Versions option in the `nwrecover` program to view the backup details of a directory or file. The versions feature lists the day, date, time, and size of the file that the Sun StorEdge EBS server backed up.

To view the version details of backed-up data:

1. In the `nwrecover` program, select the item whose backup version you want to view.
2. From the View menu, select Versions.

The Versions window displays the history of the selected file or directory, including every time the item was modified and backed up. Versions are sorted according to the backup time (the most recent backup first).

▼ To Search for a Recoverable File

Users can locate files or directories for recovery using the `nwrecover` program's Search feature by entering the name of the item being sought and the directory where the search will begin. The `nwrecover` program then displays the first file that matches the search pattern.

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. You can use wildcards to expand the search.
3. Enter a path to begin the search.

Viewing Details of Recoverable Files

You can use the File Details option in the `nwrecover` program to display additional information about files and directories. This feature displays detailed information regarding files and subdirectories, such as the date and time the file was last backed up, file permissions, owner of the file, and the date last modified.

▼ To View Details of Recoverable Files

To view details of recoverable files:

1. In the `nwrecover` program, select the item whose details you want to view.

2. Select File Details from the View menu to display this information.

Viewing the Volumes Required for Data Recovery

You can use the Volumes option in the `nwrecover` program to view a list of the volumes needed to recover files and directories. This feature displays a window with the names of all the volumes the Sun StorEdge EBS server needs to access for the data recovery.

▼ To View Volumes Required for Data Recovery

To view a list of the volumes required for data recovery:

1. In the `nwrecover` program, select the item whose volume information requirements you want to view.
2. Select Volumes from the View menu.

Disaster Recovery

Recovering data from Sun StorEdge EBS servers 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 *Legato NetWorker Disaster Recovery Guide*. For the most recent disaster recovery information, refer to the *Sun StorEdge Enterprise Backup Software Release Supplement*.

Note – You should keep a printed copy of this documentation in a safe, convenient place.

Directed Recovery

Directed recovery allows a Sun StorEdge EBS administrator to recover backed-up data from a Sun StorEdge EBS client (the source client) and direct it to another Sun StorEdge EBS client (the target client) within an enterprise. This recovery can be done remotely.

Since a user performing a directed recovery need not be physically present at the source client, the target client, or the Sun StorEdge EBS server, there is a high degree of flexibility and convenience in recovering and redirecting data. Directed recovery includes all of the Sun StorEdge EBS server's standard file handling features.

The target client receiving the recovered data can be the source client, another Sun StorEdge EBS client within the enterprise, or the administering client computer, subject to the following guidelines:

- The source and target clients must be of the same platform (UNIX-to-UNIX or Microsoft Windows-to-Microsoft Windows).
- The source and target clients must use the same filesystem type (for example, UXFS-to-UXFS, NTFS-to-NTFS).
- If the administering client is just administering—if it is neither the source client nor target client—it need not have the same platform and filesystem type as the other clients.



Caution – Sun StorEdge EBS software does not support directed recovery of data across platforms and filesystem types. In addition, it does not support directed recovery of SYSTEM or VSS SYSTEM save sets.

You cannot use a Sun Solstice Backup release earlier than 6.0 to recover data that was backed up using Sun Solstice Backup release 6.0 or later, due to the difference in save stream format.

Use the Sun StorEdge EBS command line `recover` program to administer directed recoveries. The `nwrecover` program does not support directed recovery tasks.

[FIGURE 7-1 on page 296](#) illustrates an authorized user on client *saturn* browsing client file indexes and subsequently performing a directed recovery of data to target client *mars*.




FIGURE 7-1 Browsing the Client File Index, Followed by Directed Recovery from a Remote Client

Browsing:

- Helps select files as part of a directed recovery.
- Requires user to have Remote Access All Clients privilege or to be listed in the source client's remote access list.

Directed recovery:

- Does only index-based recoveries (directed recovery using the save set ID is not supported) to a different Sun StorEdge EBS client (a target client) within the enterprise.
- Can recover from existing backups, including those prior to Sun StorEdge EBS release 6.0.
- Requires having Remote Access All Clients privilege.

Uses of Directed Recovery

Directed recovery is useful when recovering data for a shared server or other computer where users cannot recover the data themselves. Directed recovery is generally used to help in disaster-recovery situations. With a Sun StorEdge EBS client installed and the appropriate access rights, the directed recovery feature allows data to be recovered from a computer remotely, rather than from the source computer itself.

Directed recovery also permits the recovered data to be directed to other Sun StorEdge EBS client computers besides the requesting client computer. This allows file- and information-sharing with a broad range of Sun StorEdge EBS client computers. In a directed recovery, the client that originally backed up the data (source client) need not be involved.

Access Requirements for Directed Recoveries

Directed recovery is a restricted Sun StorEdge EBS function available only to users who have the Remote Access All Clients privilege. Consequently, most users cannot perform directed recoveries.



Caution – Be selective in granting the Remote Access All Clients privilege! A user with this privilege can browse all clients and can recover data to any client.

If you anticipate performing directed recoveries 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 [“Server Access Management” on page 59](#).

Access requirements to perform directed recoveries are as follows:

- The user must have the Remote Access All Clients privilege.
- The hostname of the Sun StorEdge EBS client used to administer the directed recovery must be listed in the `servers` file of each Sun StorEdge EBS client computer designated as a potential destination client.
- The `user@target_client_hostname` must have at least the Recover Local Data privilege.
- Users in the Administrators group on the Sun StorEdge EBS server are automatically granted the necessary privileges.
- The user must be a Sun StorEdge EBS Administrator while working with Sun StorEdge EBS servers running Sun StorEdge EBS software earlier than release 7.0.

TABLE 7-2 shows the access permissions and authorizations associated with performing various combinations of directed recoveries.

TABLE 7-2 Access Permissions Used in Directed Recoveries

Administering System	Source Client	Target Client	Who May Perform a Directed Recovery
UNIX	UNIX	UNIX	• User has the Remote Access all clients privilege. • The <i>user@target</i> must be listed in the source client's Remote Access list. • The <i>user@target</i> must have the Recover Local Data privilege.
Windows	Windows	Windows	• User has the Remote Access all clients privilege. • The <i>user@target</i> must be listed in the source client's Remote Access list. • The <i>user@target</i> must have the Recover Local Data privilege.
Windows	UNIX	UNIX	• The recovery must be launched by a system administrator on the administering host. • User has the Remote Access all clients privilege. • The <i>user@target</i> must be listed in the source client's Remote Access list • The <i>user@target</i> must have the Recover Local Data privilege.
UNIX	Windows	Windows	• The recovery must be launched by a system administrator (root) on the administering host. • User has the Remote Access all clients privilege. • The <i>user@target</i> must be listed in the source client's Remote Access list. • The <i>user@target</i> must have the Recover Local Data privilege.

Note – The original file ownership and permissions are always retained.

The following sections describe how to set the various server and client administration rights required to enable directed recovery.

Adding Privileges to Members of the User Group

To perform directed recoveries, users require certain privileges, such as the Remote Access All Clients and the Recover Local Data. For information about adding privileges to members of the User group, see [“Sun StorEdge EBS User Groups” on page 61](#).

Note – Members of the Administrator’s group automatically have the necessary privileges.

The Disable Directed Recover Attribute

The ability to perform directed recovery can be turned on or off for individual clients, 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 the ability for a remote host to direct a recovery to a client, this attribute should be set to Yes. If a client has this attribute set to Yes, recoveries directed to that client will not work.

For information about editing the `nsrla.res` database, see [“Editing a Client’s nsrla.res Database” on page 115](#).

▼ *To Add a Client Computer to the Remote Client’s servers File*

To add a remote computer name to the remote Sun StorEdge EBS client’s *servers* file:

1. Shut down the Sun StorEdge EBS daemons by entering the `nsr_shutdown` command at the command prompt on the remote client:

```
# nsr_shutdown
```

```
A list of Sun StorEdge EBS daemons to be shut down appears, and
you are prompted to confirm whether you want to continue with the
nsr_shutdown command.
```

2. Using a text editor, open the *servers* file, located in the `/nsr/res` (UNIX) or `<NetWorker_install_path>\res` (Microsoft Windows) directory of the Sun StorEdge EBS server.

3. Enter the Sun StorEdge EBS client computer's name in the *servers* file by using the following format:

```
mercury.planet.com
```

4. Save and close the *servers* file.
5. Start the Sun StorEdge EBS daemons by entering the following at the command prompt:

```
# /etc/init.d/networker start
```

Browsing Files Before Performing a Directed Recovery

Before performing a directed (index-based) recovery, you must browse the client file indexes to decide what to recover. Browsing is supported across platforms and filesystem types.

For information about browsing client file indexes, see [“To Recover Data” on page 288](#).



Caution – Remember that cross-platform and cross-filesystem-type recoveries are not supported. After browsing such indexes, do not initiate a recovery that mixes platforms or filesystem types (for example, UNIX-to-Windows or UXFS-to-VXFS) after browsing such indexes, or an error message results:

```
Recover: Permission denied by server <server>: Cross platform
recovery not supported
nwrecover: error recovering ./nsr/logs/daemon.log
nwrecover: recover failed: Error receiving files from NSR
server '<server>'
```

Note – Only the browsable filesystems on the source client can be recovered in a directed recovery. If a save set has exceeded its browse policy, it would still be recoverable locally, but could not be recovered by means of a directed recovery.

Performing a Directed Recovery

Use the **recover** command and its various options at the command prompt to perform directed recoveries. The **nwrecover** program does not support directed recovery. Remember that using the **recover** command is an interactive process; once the recovery is invoked, you must still interactively select and recover files chosen from the source client file index.

For information about performing directed recoveries for Microsoft Windows clients, see the *Legato NetWorker Administrator's Guide, Microsoft Windows Version*.

Using the recover Command in Directed Recoveries

The **recover** command has a variety of options you can add to modify the basic browse and recover functionality. [TABLE 7-3 on page 301](#) lists the options available with the **recover** command.

TABLE 7-3 Recover Options

Option	Meaning
-d <i>destination</i>	Specifies the name of the directory where you wish to relocate the recovered files. With this option, include a full path for destination, and use the proper syntax for the target client's platform. The use of partial paths is not supported.
-S <i>server</i>	Specifies the name of the source client's backup server.
-C <i>client</i>	Specifies the name of the source client.
-R <i>recover_target</i>	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. The Sun StorEdge EBS server automatically appends each conflicted filename with a .R.
<i>directory_name</i>	Specifies the name of the initial directory in which to begin browsing.

For a complete listing of the **recover** command and all its options, refer to the **recover** man page or the *Legato Command Reference Guide*.

Example: recover Command

The source client's backup server was named **venus**, the source client was named **mars**, and the destination client was named **jupiter**. You want the recovered files to overwrite any existing files with the same name. To do so, enter the following command at the prompt:

```
% recover -s venus -c mars -R jupiter -iY directory_name
```

Directed Recovery of Specific Files and Directories

You must use a set of procedures to specify files and directories you want to recover.

To perform a directed recovery of specific files and directories:

1. Enter the **recover** command using the following syntax:

```
% recover -s server -c source_client -R destination_client -i recover_option  
[directory_name]
```

2. Navigate to the directory containing the file or directory you want to select for directed recovery (if you did not specify the initial directory to browse by including a *directory_name*). For example:

```
recover> cd /var/adm
```

3. Select the file or directory for directed recovery by entering the following command at the prompt:

```
recover> add file_name
```

4. Repeat this procedure for each file or directory you want to select for directed recovery.
5. View the file or directory you have selected for recovery by entering the following command at the prompt:

```
recover> list
```


6. Start the directed recovery by entering the following command at the prompt:

```
recover> recover
```

Save Set Recovery

You can use the save set recover feature to recover backed-up data that has passed the period of its browse policy, but is still in the media database. When you perform a save set recover operation, you must recover the last full backup first, then recover level 1-9 and incremental backups in the chronological order that they were backed up.

The following sections discuss save set recovery:

- [“About Save Set Recovery” on page 303](#)
- [“To Recover an Entire Save Set” on page 304](#)
- [“To Recover Individual Files or Directories from a Save Set” on page 305](#)
- [“To View Volumes Required for Save Set Recovery” on page 306](#)
- [“To Relocate Recovered Save Set Data” on page 306](#)
- [“To Resolve Recovered Save Set Data Conflicts” on page 307](#)
- [“To Recover a Save Set from a Volume to the Sun StorEdge EBS Server” on page 308](#)
- [“To Recover One File Directly from a Volume” on page 309](#)

About Save Set Recovery

Save set recovery allows you to recover data from a save set rather than by using the `nwrecover` program to select items. The data recovered can be individual files, directories, or the entire save set.

When entire save sets are recovered—for example, an entire disk volume or partition on a disk—the recovery process is accelerated because the Sun StorEdge EBS server does not have to search for and locate each individual file in the index. Instead, the server only has to locate the save set containing all of the files needed for recovery.

You might want to perform a save set recovery under the following circumstances:

- The files are no longer browsable from the `nwrecover` program because the browse policy applied to the client file index entries has reached its expiration date, or the entries have been removed from the client file index through another means (such as by using the Remove Volume command).
- A disk crash during the backup produced an incomplete backup, so that it is not possible to recover the files by using the `recover` command.
- Use of the `recover` command to recover data fails, or is exceedingly slow (often indicative of insufficient computer memory).

Save set recovery is most efficient when restoring a large volume of data. Save set recovery is always an option for data that has passed its browse policy and is no longer accessible through the client file index. Save set recoveries avoid the overhead of the Sun StorEdge EBS file handling features.



Caution – Only `root` has the permission to perform a save set recovery. 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; `root@<ebs_server>` automatically has this privilege.

▼ To Recover an Entire Save Set

Note – Save set recovery recovers data to the machine on which the Sun StorEdge EBS Administrator program is launched. If you want to recover the data to the original client machine, you must perform the save set recovery on that computer.

To perform a full save set recovery of a client:

1. In the Sun StorEdge EBS Administrator program, select Recover from the Save Set 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 is more than one version.
5. Click Recover.
6. In the Save Set Recover Status dialog box, click Start to begin the recovery.

The recovery status appears in the Status box of the Save Set Recover Status dialog box.

You can also recover a save set from the command prompt using the `recover` program with the following syntax:

```
# recover -s server -S ssid -d directory_path
```

For more information, refer to the `recover` man page.

▼ To Recover Individual Files or Directories from a Save Set

To recover specific data from a save set.

1. In the Sun StorEdge EBS Administrator program, select Recover from the Save Set 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 is more than one version.
5. Click Recover to open the Save Set Recover Status 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 Paths to Recover text box. Click Add after each entry to add the entry to the list box.

Note – The save set recovery process parses a directory name and restores all instances that match the leading string, unless limited by a slash and assuming the overwrite option in cases of a recovered data conflict. Therefore, when specifying directory paths for recovery, 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.

Note – Click Start to begin the save set data recovery.

The recovery status appears in the Status box of the Save Set Recover Status dialog box.

▼ To View Volumes Required for Save Set Recovery

You can view a list of the volumes required to recover save set data while in the Save Set Recover Status dialog box. This dialog box displays each volume that needs to be available to the backup device to complete the save set recovery.

To view the volumes required for the save set recovery:

1. In the Sun StorEdge EBS Administrator program, select Recover from the Save Set 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 is more than one version.
5. Click Recover to open the Save Set Recover Status dialog box.
6. Click Volumes to display the volumes required for the recovery.

If no volumes appear, that indicates that no entries for the volume exist in the media database. Use the scanner program to re-create the client file index entries or media database entries. The scanner program can read the storage volume directly, without assistance from the Sun StorEdge EBS server. For more information about using the scanner program to re-create online indexes, see [“Recovering Online Indexes” on page 312](#).

▼ 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 filesystem.

To relocate recovered save set data:

1. In the Sun StorEdge EBS Administrator program, select Recover from the Save Set 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 you want to recover.
4. From the Instances list, select the particular version of the save set, if there is more than one version.

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.

For more information on relocating recovered data, see [“Relocating Recovered Data” on page 292](#).

▼ To Resolve Recovered Save Set Data Conflicts

While in the Save Set Recover Status dialog box, you can instruct the Sun StorEdge EBS server to perform a particular action when a data recovery conflict occurs during the save set recovery.

To resolve data recovery conflicts with save set data:

1. In the Sun StorEdge EBS Administrator program, select Recover from the Save Set 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 you want to recover.
4. From the Instances list, select the particular version of the save set, if there is more than one version.
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 7-4 on page 307](#) lists the options and the respective actions taken for the Duplicate File Resolution.

TABLE 7-4 Naming Conflict Resolution Options

Option	Action
Rename the Recovered File	Renames the recovered files by appending each conflicting name with .R.
Discard the Recover File	Discards the recovered file in order to prevent the Sun StorEdge EBS server from overwriting an existing file.
Overwrite the Existing File	Replaces the file on the filesystem with the recovered version.

8. Click OK.

For more information on resolving data recovery conflicts, see [“Resolving Recovered Data Conflicts” on page 291](#).

▼ To Recover a Save Set from a Volume to the Sun StorEdge EBS Server

To recover an entire save set directly from a volume to the Sun StorEdge EBS server’s filesystem, use the **scanner** program by entering the following command at the prompt:

```
# scanner -S save_set_id device_name | /usr/lib/nsr/uasm -rv
```

This command reads all the information associated with the *ssid* from the volume and places an exact copy of this data on the Sun StorEdge EBS server in the exact way that it is stored on the backup volume. In other words, the backup volume may contain files for a client, but is recovered to the Sun StorEdge EBS server’s filesystem.

To test the command before you perform it, add the **-n** flag to the **uasm** command. The **-n** flag sends the output from **scanner** to */dev/null* and lists all the filenames contained in the save set.

If the save set originated on a Sun StorEdge EBS client instead of the Sun StorEdge EBS server, you can use **rsh** (or its equivalent) in conjunction with the following command to recover the save set to the client:

```
# scanner -S ssid device_name | rsh client "cd destdir; /usr/lib/nsr/uasm -rv"
```

▼ To Recover One File Directly from a Volume

To recover a single file from a volume, enter *one* of the following commands at the prompt:

```
# scanner -S save_set_id device_name | /usr/lib/nsr/uasm -rv filename  
  
# scanner -S save_set_id device_name | /usr/lib/nsr/uasm -rv -m source=  
dest filename
```

The `-m` option of `uasm` maps (relocates) the recovered file from the *source* to the *dest* (destination) directory.

Comparison of Recovery Methods

This section compares the two types of data recovery methods that the Sun StorEdge EBS software provides:

- Index-based recovery, which includes directed recovery and standard client file index-based recovery
- Save set recovery

Index-Based Recovery

Standard recovery and directed recovery both use the client file index to locate and recover data.

Advantages of Index-Based Recovery

Index-based recovery provides the following key benefits:

- Directed recovery is well-suited for centralized administration. A large network can be administered for direct recovery from a single administering computer.
- Directed recovery is useful if a computer is inoperable, a hard disk has crashed, or a network no longer recognizes a computer. You can easily configure the replacement computer by recovering a set of files for the client that were backed up from another client on the network.

- Directed recovery is convenient if you want to transfer files between two clients. For example, if one client's `.profile` file is perfect for a new client, you can recover it from the backups of the old client to the new client's hard disk.

A standard, client file index-based recovery differs from a directed recovery only in the destination of the recovered data. With a standard recovery, the Sun StorEdge EBS server recovers the data to the data's original client. Thus, both kinds of client file index-based recoveries can be contrasted with a save set recovery.

Disadvantages of Index-Based Recovery

Index-based recovery has the following disadvantages:

- Entries for data that has passed its browse policy period are not listed in the client file index. Entries for these files can be recovered and restored to the client file index using the `scanner` program, which can be time-consuming or inconvenient. For more information on `scanner`, see [“The scanner Program” on page 316](#).
- It takes time to mark each file for recovery.
- It requires memory proportional to the number of files marked for recovery.

For these reasons, if you plan to recover a significant number of files, you might prefer to use save set recovery.

Save Set Recovery

Save set recovery involves recovering an entire save set, rather than selecting individual files within the save set for recovery.

Advantages of Save Set Recovery

Save set recovery provides the following benefits:

- Save set recovery is useful 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.
- When the browse policy period for data has expired, entries related to the data are removed from the client file index. In this case, save set recovery must be used to recover the data.
- You can also run `scanner` to re-create the client file index.

Disadvantages of Save Set Recovery

Save set recovery has the following disadvantages:

- You must locate the full save set and its relevant level saves and incrementals, and restore all of those save sets in the correct order.
- The recovery can restore more data than was requested. Files that were deleted after the last full save set backup are also recovered.

A save set recovery follows the basic steps of an index-based recovery, but without the advantage of client file index handling that the Sun StorEdge EBS server provides automatically with index-based recovery.



Caution – If you are recovering a number of save sets that were backed up at different levels, you must be careful to execute the save set recovery requests in the correct order, beginning with the full backup. For example, if file X existed at time A, but was deleted prior to time B, file X will be recovered during a save set recover at time B. Sun StorEdge EBS file handling is not available for save set recovery, so file X remains on your Sun StorEdge EBS server during the recovery operation.

It is possible that disk space could become scarce if too many previously deleted files or directories are recovered. If the save sets you recover are full rather than incremental, running out of disk space will be far less of an issue because you will not recover multiple copies of the same files or directories.

Typically, during a recovery you want to overwrite existing files. Overwriting files is even more important with save set recovery, because the same file might be recovered multiple times, with each successive version of the file coming from a later save set. For information about instructing the Sun StorEdge EBS server to overwrite existing files during save set recovery, see [“To Resolve Recovered Save Set Data Conflicts” on page 307](#).

Deciding between Index-Based and Save Set Recovery

This section summarizes the reasons for using either index-based or save set-based recovery.

You might prefer index-based recovery over save set recovery for the following reasons:

- Users can browse the index to mark data for recovery—users and administrators do not need to know the exact filename, location, or date of the specific files they want to recover.

The Sun StorEdge EBS server supports multiple users (sessions) concurrently browsing a client file index, where users can expand directories to view files and subdirectories. Thus, needed files usually can be deduced from information obtained during browsing. Users do not have to remember explicit filenames from previous backups.

- Users can query the client file index with a target time for which they want to rebuild the file hierarchy. The Sun StorEdge EBS server constructs a view of the filesystem that was written closest to the browse target time without exceeding it.
- The Sun StorEdge EBS server automatically re-creates the data exactly as it existed at the time of the backup.

For example, if a user requests recovery of a level 5 backup, the Sun StorEdge EBS server first recovers the full backup on which that level 5 depends, then applies the level 5 backup. Any renaming of files or deletions that occurred between the full and level 5 backups are handled as well.

Additionally, client file index-based recovery is most appropriate under the following conditions:

- You cannot determine which save sets to recover.
- Recovering extra files is not desired.
- Using save set recovery, you may have to perform multiple recoveries to restore the file to the correct state. For example, if a file is first backed up as full, then backed up at an incremental level, you have to restore two save sets—first the full backup, then the incremental backup—to restore the file to the state it was in after the incremental backup. Using index-based recovery, only one recovery is necessary. The Sun StorEdge EBS server handles the incremental backup issue internally, prompting you, if necessary, for the appropriate volumes.

Save set recovery is most appropriate under the following conditions:

- You can determine the correct save sets to recover.
- There are only a few save sets to recover for each filesystem.
- Recovering extra files is acceptable.
- Free disk space is not a problem.
- Your most recent save sets are from full, not incremental backups.
- The browse policy has expired, and you do not want to use the `scanner` program to re-create the client file index entries.

Recovering Online Indexes

This section describes different procedures for recovering online indexes.

▼ To Recover Client File Indexes

Unlike regular backed-up data, client file index information is recovered by using the `nsrck` program with the `-L7` option to perform a level 7 consistency check. However, a level 7 check merges recovered data into the existing client file index; therefore, if the existing client file index has become corrupt, you must first purge the corrupt save set information from the index using a level 5 consistency check before recovering the index from a backup volume.

To recover a client file index:

1. Log in as root.
2. To purge corrupted save set information from the index, run the `nsrck` command with the `-L5` option:

```
# nsrck -L5 client_name
```

3. To recover the client file index, run the `nsrck` command with the `-L7` option:

```
# nsrck -L7 client_name
```

Note – You can also completely remove a client file index by entering `nsrck -R -Y client_name`, then entering `nsrck -L7 client_name` to recover the index.

Recovering Client File Indexes from a Specific Time and Date

You can recover a client file index from a specific time and date using the `-t time date` option with `nsrck`. The time and date formats are in `nsr_getdate` format. For details on the `nsrck` and `nsr_getdate` commands, refer to the *Legato Command Reference Guide* or the appropriate man pages.

Index recovery is nondestructive, so 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 index contents. Because of this, the client file index becomes much larger. Index backups subsequent to this will also contain the added indexes if they are not explicitly removed using the `nsrmm` with the `-d -P` options. For more information, refer to the `nsrmm` man page.

Recovering the client file index from a specific time and date allows you to browse save sets that have passed their browse policy but are still recoverable. The save sets referred to by the recovered client file index are marked as browsable. They remain browsable for the length of time they were originally browsable.

To recover a client file index for a specific date, enter the following command at the prompt:

```
# 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, 2003:

```
# nsrck -t "04 pm 03/02/2003" -L7 client_name
```

▼ To Restore a Save Set Entry in the Online Indexes

The following sections describe the process of restoring save set entries into the only client file index, or into both the client file index and the media database.

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.
2. Run the `mminfo` command

```
# mminfo -a -v volume_name
```

3. From the `mminfo` output, find the `ssid` that you believe contains the file you want. Make sure it is not the bootstrap `ssid`.

4. Restore the save set entry in the client file index by entering the `nsrmm` command with the `-w` option to set a new browse time, and then the `scanner` program:

```
# nsrmm -S ssid -w browse_time
```

where `browse_time` is a time in the future.

Then enter:

```
# scanner -i -S ssid
```

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 both the client file index and the media database:

1. Log in as root.
2. Run the `scanner` program on the backup volume that you believe contains the appropriate file:

```
# scanner device_name
```

For instructions on finding the volume name, see [“To Find a Volume Name” on page 241](#).

3. Use the output from the `scanner` program to determine:
 - Whether the save set you want to rebuild is on this volume.
 - Whether to reintroduce the contents of this volume into the online indexes. (You must locate *all* the volumes that contain this `ssid`.)
4. After you have determined which volumes to reintroduce into the online indexes, run the `scanner` program:

```
# scanner -i device_name
```

The `scanner` program prompts for a new volume until you terminate it. To rebuild the indexes completely, you must scan in *all* the volumes that contain the *ssid*.

5. Use the `nwrecover` program to browse the client file index for the file you want to recover.

The scanner Program

The `scanner` program reads the volume directly, without accessing information that might be available in the media database. If no entries exist in the media database, you must take an educated guess at which volume contains the data to recover.

The `scanner` program is used to import data that has passed its browse and retention time, or was backed up by a different Sun StorEdge EBS server. 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. You can use `scanner -B` to locate bootstrap save sets on a volume, and then use `mmrecov` to restore the bootstrap. For more information, refer to the *Legato NetWorker Disaster Recovery Guide*.

Mount the volume and execute the **`scanner`** command with any of its applicable options. The most appropriate option depends upon the situation, your level of knowledge about the location of the data, and the nature of the recovery.

[TABLE 7-5 on page 316](#) lists the most common **`scanner`** options.

TABLE 7-5 Common scanner Options (1 of 2)

Option	Meaning
<code>-m device</code>	Rebuilds the media database entries for everything on the storage volume.
<code>-i device</code>	Rebuilds the media database and client file index entries for everything on the storage volume.
<code>-S save_set_id -m device</code>	Rebuilds the media database entries for a specific save set on the storage volume.
<code>-S save_set_id -i device</code>	Rebuilds the media database and client file index entries for a specific save set on the storage volume.
<code>-C client_name</code>	Processes only save sets from the specified Sun StorEdge EBS client machine.
<code>-N save_set_name</code>	Processes only save sets specified by <i>save_set_name</i> .

TABLE 7-5 Common scanner Options (2 of 2)

Option	Meaning
-S <i>save_set_id</i>	Sends a specific save set to standard output.
-x	Introduces further processing commands on a per-save set basis.
-B	Locates the bootstrap save sets on the volume.

Note – For information about using the `scanner` program to import save set information from a volume backed up by a different Sun StorEdge EBS server, see [“Importing Save Set Information with the scanner Program” on page 317](#).

For detailed information about `scanner` program syntax and how to customize `scanner` options for a particular recovery situation, refer to the `scanner` man page or the *Legato Command Reference Guide*.

The entries placed in the client file index assume the browse and retention policies of the original save set. If a save set originally had a browse time of one month and a retention time of three months, and that save set expired, when you scan the save set back in, it then remains browsable for one month and recoverable for three months.

The entries placed in the media database also assume the browse and retention policies of the original save set, and the volume mode is listed as “read-only.” For more information on managing the client file index and media database entries, see [“Storage Management Operations” on page 231](#).

If you upgrade a Sun StorEdge EBS server from a version prior to release 6.0, the current browse and retention policies are applied to save sets that had been scanned into the client file index and the media database.

Importing Save Set Information with the scanner Program

The `scanner -i` program can be used to import client file index entries for save sets on a volume that was backed up by a different Sun StorEdge EBS server. However, the Sun StorEdge EBS software uses a unique client ID for each client, and during backup, the client ID is recorded with each save set. When clients with the same name are configured on different Sun StorEdge EBS servers, they will have different client IDs.

The **scanner** program uses the client ID, rather than the client name, to identify clients. As a result, even when imported save sets belong to a client that has the same name on two servers, the **scanner** program is not able to associate imported data with the appropriate client.

To import the client data in this case, you must create a new client on the server to which the data is being imported. The new client must be created with the client ID that was used by the previous server.

You cannot simply delete the old client and re-create it with the new client ID. The Sun StorEdge EBS server retains the old client ID even after the client is deleted, so that all save sets indexed for the old client can be recovered. When a new client is created by using the same name as a deleted client, the new client will automatically have the client ID of the deleted client.

To resolve this problem, you must do one of the following:

- Create a new client by using the client ID and a unique name for the client. The client's name must have the format `~hostname-#`.
- Rename the existing client by following the procedure found at [“To Change a Client Name” on page 114](#) and then create a new client by using the client ID from the previous server.

▼ To Create a New Client Based on the Client ID

To create a new client based on the client ID:

1. Determine the client ID on the old server. For information about determining the client ID, see [“The Client ID” on page 113](#).
2. In the Sun StorEdge EBS Administrator program, select Client Setup from the Clients menu.
3. Click Create.
4. 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 a client that already exists on the new server.

If a client with the appropriate name already 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 for the client name:

`~jupiter-1`

5. From the View menu, select Details to display hidden attributes.
6. In the Client ID attribute, enter the client ID determined in [Step 1](#).
7. Complete other attributes as necessary, and click Apply.

Once the new client is created, import the client data using `scanner -i`. You can now perform directed recoveries to restore the client data. For more information about directed recoveries, see [“Directed Recovery” on page 294](#).

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 4](#). You can then create the client based on the client ID by using the preceding procedure. The entries in the media database will be associated with the correct client. However, you must run `scanner -i` again after creating the Client resource in order to import save set information into the client file index.

Data Archiving

This chapter describes the features and functionality of the optional Sun StorEdge Enterprise Backup Software archive feature.

The Sun StorEdge EBS archive feature is an optional feature that you must purchase and license separately from other Sun StorEdge EBS software components. The functionality only works when the archive feature is properly licensed in the Sun StorEdge EBS server. The evaluation conditions and procedures are the same for the archive feature as for other Sun StorEdge EBS products. For more information on licensing procedures, see the Sun StorEdge EBS Roadmap.

The main sections in this chapter are:

- [“Overview of Sun StorEdge EBS Archiving” on page 321](#)
- [“Permissions for Archiving” on page 325](#)
- [“About Archive Pools” on page 326](#)
- [“Archiving Data Procedures” on page 328](#)
- [“Retrieving and Recovering Archived Data” on page 333](#)
- [“Archive Request Management” on page 335](#)

Overview of Sun StorEdge EBS Archiving

The Sun StorEdge EBS archive feature provides archive services to Sun StorEdge EBS client computers. The archive process captures files or directories as they exist at a specific time, and writes the data to special 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 principal difference is that archive save sets have no expiration date. By default, the archive backup level is always set to full. Because the Sun StorEdge EBS server treats archive data differently than backup save set data, archive data must be written to separate pools. Browse and retention policies do *not* apply to archive data.

Some of the 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, and retained in archive volumes for quick retrieval in the future.
- Archived data is never subject to automatic recycling, so it *cannot* be accidentally overwritten.
- 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, make sure 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 (`nsrchive`), which is initiated by the client's `nsrexecd` service. You can schedule archives from the server or client using the Archive Requests resource in the Sun StorEdge EBS Administrator program.

FIGURE 8-1 on page 324 illustrates how the Sun StorEdge EBS software archives data.




FIGURE 8-1 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)

Data archived by using indexed archiving has entries added to the client file index. As a result, you can browse and select individual files for recovery.

To take advantage of indexed archiving, you must do either 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

The following sections describe permissions required to use the archive feature:

- [“Enabling Archive Services for the Client” on page 325](#)
- [“Enabling or Restricting Archive Access” on page 326](#)
- [“Enabling Public Archive Access” on page 326](#)

Enabling Archive Services for the Client

After the archive service is appropriately 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, make sure that the Archive Services attribute is enabled in the Client resource for the server. For instructions, see [“Enabling Archive Services for a Sun StorEdge EBS Client” on page 328](#).



Caution – When you enable the Archive Services attribute for a Client resource, you also enable the Archive Services attribute for all other clients of the same name on that server. For example, if the Legato 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. The Archive Services attribute is either enabled for both or disabled for both.

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 user names in this attribute. They must have read permissions to the data and must have the Backup Local Data privilege in order 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 61](#)

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

Note – Only members of the Administrators group or users with the Change Security Settings permission enabled can change the Archive Users attribute. For more information, see [“Sun StorEdge EBS User Groups” on page 61](#).

Enabling Public Archive Access

To allow users listed in the Archive Users attribute of the Client resource to retrieve any archived file from that client, enable the Public Archives attribute in the Server resource. 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.

You *cannot* omit specific clients from the retrieval permissions when you enable the Public Archives attribute.

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 327](#).

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.



Caution – 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 [“Pools” on page 166](#).

Archiving Data Procedures

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

This section outlines the following procedures:

- [“Enabling Archive Services for a Sun StorEdge EBS Client” on page 328](#)
- [“Manually Archiving Data” on page 328](#)
- [“Scheduling Archives” on page 331](#)

Enabling Archive Services for a Sun StorEdge EBS Client

To enable archive services for a Sun StorEdge EBS client:

1. In the Sun StorEdge EBS Administrator program, select Client Setup from the Clients menu.
2. To create a new client, click Create and complete the client information. To edit an existing client, select the client from the list provided.
3. For the Archive Services attribute, select Enabled to enable archive services for the client.
4. In the Archive Users attribute, enter all appropriate users. For details about setting this attribute, see [“Enabling or Restricting Archive Access” on page 326](#).
5. Click Apply.

Manually Archiving Data

You can manually archive data at any time. Manually archiving data is similar to manually backing up data.

You can designate the entire filesystem, specific directories, or individual files for manual archiving. Note that if you select a directory for archiving, all of its nested subdirectories and files are also selected. If you want to exclude certain files from manual archiving, you must unselect those items using the `nwarchive` program.

This section outlines the following procedures:

- [“To Request a Manual Archive from a Sun StorEdge EBS Client” on page 329](#)
- [“Manually Archiving from the Command Prompt” on page 330](#)

▼ To Request a Manual Archive from a Sun StorEdge EBS Client

To request a manual archive:

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 &
```

The `nwarchive` program displays a representation of the client’s filesystem, similar to the `nwbackup` program used for manual backups.

2. Select Mark from the Mark menu or click the Mark button 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 Mark menu, or click the Unmark button.
3. For the File menu, select Start Archive.

Note – You can also click the Start button.

4. Enter a comment in the Annotation attribute. This annotation is used to identify each archive save set during retrieval.
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 for the Compress attribute. Compression and decompression of data typically uses 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 for the Clone attribute.

If you enable cloning, enter or select an archive clone pool in the Clone Pool attribute. Only pools with their Pool Type attribute set to Archive Clone are listed.

- **Verify:** To instruct the Sun StorEdge EBS server to check the integrity of the data on the storage volume, select Yes for the Verify attribute.
- **Grooming:** To instruct the Sun StorEdge EBS server to remove the archived files from the disk, select Yes for the Grooming attribute.

7. Click Archive.

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: 2-21-03 5:18p

8. If the Groom option was selected in [Step 6](#), the Remove Archived File dialog box prompts you for confirmation before Sun StorEdge EBS software deletes archived files from your local disk.

If you do *not* want to be prompted for each file, select Suppress Further Prompting and all archived files will be removed from your local disk without further prompting.

Manually Archiving from the Command Prompt

You can also use the command line program `nsrarchive` to manually archive data. By default, when using the `nsrarchive` program, you are prompted to determine if each archived file should be groomed. To groom files automatically, include the `-y` option with the `nsrarchive` command.

For more information, see the man page for this command, or refer to the *Legato Command Reference Guide*.

Scheduling Archives

Scheduled archives of client data are configured through the Sun StorEdge EBS Administrator program and are performed by the Sun StorEdge EBS server. Unlike scheduled backups, scheduled archives run only once. The advantage of a scheduled archive is that you can run the archive when network traffic and computer use are low.

This section outlines the following procedures:

- [“To Schedule an Archive” on page 331](#)
- [“To Edit an Archive Request” on page 332](#)
- [“To Delete an Archive Request” on page 333](#)

▼ To Schedule an Archive

Note – Before you can schedule an archive request, you *must* enable the Archive Services attribute in the Client resource. For details, see [“Enabling Archive Services for a Sun StorEdge EBS Client” on page 328](#).

To schedule an archive:

1. In the Sun StorEdge EBS Administrator program, select Archive Requests from the Customize menu.
2. Click Create.
3. For the Name attribute, enter a name for the archive request (64-character limit).
4. For the Annotation attribute, enter a comment. This annotation is used to identify each archive save set during retrieval.
5. For the Status attribute, select one of the following options:
 - Start Now: Starts the archive as soon as you click Apply in the Archive Request window.
 - Start Later: Starts the archive at the time you specify in the Start Time attribute.
 - Enter a time for the archive request to start in the Start Time attribute.
 - Enter the time value in 24-hour time format, such as 23:59 for 11:59 p.m. The default start time value is 3:33.
6. For the Client attribute, enter the archive client’s hostname. If you want the Sun StorEdge EBS server to archive its own files, enter the server’s hostname.
7. For the Save Set attribute, enter the pathnames of the files or directories you want to archive.

Note – If you archive all of the client’s save sets, select None for the Grooming attribute (see). If you select Remove for the Grooming attribute, the Sun StorEdge EBS server deletes all of the client’s archived save sets.

8. For the Directive attribute, enter or select a directive. For more information on directives, see [“Directives” on page 144](#).
9. From the Archive Pool list, select the appropriate pool.
10. Select the appropriate response for the following data archiving attributes:
 - **Verify:** To instruct the Sun StorEdge EBS server to check the integrity of the data on the storage volume, select Yes for the Verify attribute.
 - **Clone:** To write a copy of each archive save set to a volume from an archive clone pool, select Yes for the Clone attribute.

If you enable cloning, enter or select an archive clone pool in the Archive Clone Pool attribute.
 - **Grooming:** To instruct the Sun StorEdge EBS server to remove the archived files from your disk, select Yes for the Grooming attribute.
11. For the Archive Completion attribute, enter a notification for the Sun StorEdge EBS server to use after completing the data archive. For information on notifications, see [“Notifications” on page 65](#).
12. Click Apply.

To view information about the status of an archive request, open the Archive Request Details window. See [“To View Details of a Scheduled Archive” on page 337](#).

▼ To Edit an Archive Request

To edit an archive request:

1. In the Sun StorEdge EBS Administrator program, select Archive Requests from the Customize menu.
2. Select the archive request from the list in the Archive Requests attribute.
3. Edit the attributes of the archive request and click Apply.

▼ 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 Sun StorEdge EBS Administrator program, select Archive Requests from the Customize menu.
2. In the Archive Requests list, select the appropriate archive request.
3. Click Delete.
4. When prompted for confirmation, click OK.

Retrieving and Recovering Archived Data

This section describes how to use the Sun StorEdge EBS archive feature to retrieve archived data:

- [“Retrieval Permissions” on page 333](#)
- [“To Retrieve Nonindexed Archived Data” on page 334](#)
- [“To Recover Indexed Archive Data” on page 335](#)

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 325](#).

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.

▼ To Retrieve Nonindexed Archived Data

Data archived using nonindexed archiving does *not* maintain information about the archived save set in the client file index. Therefore, data is retrieved for an entire save set.

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: Enter the name of the client whose archived data you want to retrieve.
 - Save Set Name: Enter the name of the archived save set you want to retrieve.
 - Search Annotation For: Enter all or part of the annotation assigned to the save set when it was archived. To aid the search, use wildcard symbols in this attribute.
4. Click Query to search for the save sets that match the criteria specified in [Step 3](#). 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 you want to retrieve and select Start Retrieve from the File menu or click the Start toolbar button. The Retrieve Status window appears.

To relocate the files to a different location, enter the new location in the Relocate To attribute of the Retrieve Status window. To automatically overwrite existing files, enable the Overwrite Existing Files attribute.

6. Click OK.

The retrieval status appears in the Status attribute of the Retrieve Status window. You may be prompted to indicate an appropriate action for the Sun StorEdge EBS server to perform when it encounters files with duplicate filenames during retrieval.

Note – You can also recover archived save sets using save set recovery. For information, see [“Save Set Recovery” on page 303](#).

▼ To Recover Indexed Archive Data

Data archived 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 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 that you wish to recover by selecting the files and clicking the Mark button on the toolbar.
4. Click the Start button on the toolbar.

Archive Request Management

This section describes how to accomplish basic tasks involving Sun StorEdge EBS archive requests:

- [“To Start a Scheduled Archive at Any Time” on page 336](#)
- [“To Stop a Scheduled Archive While in Progress” on page 336](#)
- [“To Disable a Scheduled Archive” on page 336](#)
- [“To View Details of a Scheduled Archive” on page 337](#)

▼ To Start a Scheduled Archive at Any Time

When you want to start a scheduled archive right away, rather than waiting for the scheduled time, the Sun StorEdge EBS server overrides the archive request's scheduled start time and immediately archives the data specified in the archive request.

To start a scheduled archive:

1. In the Sun StorEdge EBS Administrator program, select Archive Request Control from the Server menu.
2. Select the name of the archive request in the Archive attribute.
3. Click Start.

▼ To Stop a Scheduled Archive While in Progress

To stop an archive request in progress:

1. In the Sun StorEdge EBS Administrator program, select Archive Request Control from the Server menu.
2. Select the name of the archive request you want to stop in the Archive attribute.
3. Click Stop.

When you stop a scheduled archive in progress, the archiving action stops immediately and the archive request is listed as stopped in the Archive Requests attribute of the Archive Requests Control window. To restart the archive request, see [“To Start a Scheduled Archive at Any Time” on page 336](#).

▼ To Disable a Scheduled Archive

To disable an archive request:

1. In the Sun StorEdge EBS Administrator program, select Archive Request Control from the Server menu.
2. Select the name of the archive request you want to disable in the Archive attribute.
3. Click Disable.

▼ To View Details of a Scheduled Archive

To open the Archive Request Details window:

1. In the Sun StorEdge EBS Administrator program, select Archive Request Control from the Server menu.
2. Select the name of the archive request whose details you want to view in the Archive attribute.
3. Click Details to open the Archive Request Details window.

The Archive Request Details window provides information about the completion of an 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 success of the archive request is either completed, failed, or partial.

Cloning

The Sun StorEdge Enterprise Backup Software server can make duplicates, or clones, of complete volumes of stored data as well as individual save sets. This allows for secure offsite storage, easy transfer of data from one location to another and verification of backups.

The main sections in this chapter are:

- [“Volume Cloning” on page 339](#)
- [“Save Set Cloning” on page 345](#)
- [“Other Information and Procedures” on page 356](#)

To perform cloning operations, you must be a Sun StorEdge EBS administrator.

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

How Volume Cloning Works

Volume cloning is a two-part operation:

1. The Sun StorEdge EBS server reads data from the source volume.
2. The Sun StorEdge EBS server writes the data to a clone volume (a volume that belongs to a [clone pool](#)).



Caution – Cloning requires at least two active devices: one for reading the source volume, and one for writing the new, cloned data to another volume. This means there must be at least two devices enabled on the server in order for cloning to succeed.

During cloning, data is copied from the source volume to the clone volume. When you clone data, different capacities of storage media may mean that more or fewer clone volumes are required than were required for the original data. When you specify that a particular volume is to be cloned, the Sun StorEdge EBS software uses the save sets on the specified volume as the source data.

The Sun StorEdge EBS Administrator program displays cloning information in the Messages window. Clone-related messages are also logged to the Sun StorEdge EBS message file `/nsr/logs/messages`.

To clone volumes, you must be a member of the Sun StorEdge EBS Administrators group. For more information, see [“Sun StorEdge EBS User Groups” on page 61](#).

Volume Cloning Requirements

The following are required for cloning volumes with the Sun StorEdge EBS server:

- A minimum of two storage devices: one to read the existing data and one to write the cloned data.
- 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.

Volume Cloning Performance

In general, a volume 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 scheduled backups that follow.

The Sun StorEdge EBS server usually 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 any automatic cloning is complete. Therefore, if another group starts its backup while the previous group's clone operation is under way, there may be a contention for `nsrmmd` resources or specific volumes.

To avoid this problem, you may decide to refrain from automatic cloning and instead initiate a single cloning operation using the `nsrclone` command as part of a scheduled task that runs at a nonpeak time after all backups are complete. For more information, refer to the `nsrclone` man page or the `nsrclone` entry in the *Legato Command Reference Guide*.

Cloning Compared to Duplication of Volumes

When a volume is cloned, all save sets that begin on the volume are cloned. As a result, it is possible that more than one volume is required for a single volume clone. This is because the last save set on the volume may continue onto another volume. When this is the case, the save set is still cloned and an additional volume is required to complete the clone.

Also, the cloning operation leaves traceable information 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 an operating system or hardware device copy operation.

Problems with Volume Duplication When Using a Copy Operation

If you rely on an operating system or hardware device copy operation to duplicate volumes, you must first ensure that the destination can hold all the save sets contained on the source volume. If the last save set continues onto another volume, this save set is not automatically duplicated onto a second volume (as it is with cloning). This could result in lost data.

In addition, 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 the duplicated volume is left in the autochanger, the unlabeled volume is considered eligible for labeling and reuse during a scheduled backup. For more information on auto media management, see [“Auto Media Management” on page 226](#).

Similarly, it is possible to make an exact copy of a Sun StorEdge EBS archive volume. However, the annotation information associated with each archive save set is stored in the Sun StorEdge EBS server’s media database, not on the archive volume itself. Therefore, a duplicate of an archived save set does not include this annotation information. Furthermore, if the entry for the original archive save set is removed from the media database, the annotation information that describes it is also removed.

Volume Cloning and Data Tracking Information

A cloning operation does not insert entries into the client file index. Cloned save sets are only tracked through the media database. During cloning, the location of a cloned save set is added to the existing save set entry in the media database. That is, each save set clone shares the same *save set ID (ssid)* as the source save set. All characteristics that are true for the source save set are also true for the clone save set:

- If the source save sets are still browsable, the clone status is also browsable.
- If the source save sets have passed their browse policies, the clone status is recoverable.

Volumes that belong to a clone pool are also tracked through volume entries in the media database. The fact that all save sets share the same media database save set entry has implications for the following actions, which are executed on a *per-save-set* basis and not on a *per-volume* basis:

- Changing the mode of a cloned volume (of save sets)
- Purging a volume (of save sets) from the client file index
- Deleting a volume (of save set locations) from the media database

Changing the Mode of a Volume Clone to Recyclable

In some cases, you might want to manually change the mode of a clone volume to *recyclable*, with the intent of reusing that clone volume. However, keep in mind that the mode of a volume only changes to recyclable when all the save sets on that volume are recyclable. Therefore, when you change the mode of the clone volume to recyclable, you effectively change the status of all save sets on the volume to recyclable. Because the save sets share the same entry in the media database, there is no distinction between original and clone save sets. The end result is that all the save sets that reside on the now-recyclable volume, or on any other volume, become candidates for immediate recycling.

If you want to reuse a particular clone volume and still protect the save sets that exist on other volumes from inadvertent data loss, first change the mode of the volumes that you want to protect to manually recycle. This way, the Sun StorEdge EBS software cannot automatically recycle the volume. Then, you can safely change the volume that you intend for reuse to recyclable mode.

Purging a Clone Volume

Similarly, if you *purge* a clone volume, you effectively remove from the client file index all the file entries associated with all the save sets that reside on the clone volume (in whole or in part).

Deleting a Clone Volume

If you *delete* a clone volume, the `nsrim` index management program locates the entry in the media database for each save set that resides on the clone volume. The `nsrim` program then marks the information about the location of one of the save set clones for deletion from the entry. This action is performed for each save set entry. In addition, `nsrim` marks the entry for the particular clone volume (identified by its volume ID number) for deletion from the database.

Volume Cloning and Recovery

A clone volume is used for recovery any time 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*.

For more information about changing the status of a save set, see [“To Change the Status of a Save Set” on page 119](#).

You can execute 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 on the `scanner` program, see [“The scanner Program” on page 316](#). For information on how to recover data with the `scanner` program, refer to the *Legato NetWorker Disaster Recovery Guide*.

Creating Clone Volumes

You can manually clone all of the data on a volume in the Clone Volume window, as shown in [FIGURE 9-1 on page 344](#). This window provides information about the volumes for the server currently in use, and is a starting point for the cloning process.




FIGURE 9-1 Clone Volume Window

▼ To Create a Clone Volume

To create a clone volume:

1. In the Sun StorEdge EBS Administrator program, select Clone Volume from the Media menu.
2. In the Clone Volume window, select the volume to clone from the Volumes list.
3. Click Clone to open the Clone Volume Status dialog box.
4. Enter the name of the clone pool to which the clone volume will belong. Cloned volumes must belong to a clone pool, either one of the preconfigured clone pools or one you create. If you do not select a clone pool, the Sun StorEdge EBS server automatically selects the Default Clone Pool.
5. Click Start to begin the cloning.

Once the cloning process begins, its progress appears in the Status field of the Clone Volume Status window.

▼ To View Clone Volume Details

You can view the details of a clone volume, including its estimated capacity, amount of space used, number of files, and mode, by opening the Volume Details window.

To view clone volume details:

1. In the Sun StorEdge EBS Administrator program, select Clone Volume from the Media menu.
2. In the Clone Volume window, select the volume whose details you want to view in the Volumes attribute.
3. Click Details to open the Volume Details window.

Save Set Cloning

Cloning save sets allows you to create identical copies of backed-up data quickly and easily. The save set cloning feature provides a high degree of flexibility for determining which data to copy. You can clone save sets automatically or manually. You can also clone a single save set, or configure the software to clone only full backups for an entire group.

Overview

The Sun StorEdge EBS cloning function ensures the safety of data by making a copy of each selected save set and writing it to a destination volume. Source and destination media can be of different types and capacities. If necessary, cloned save sets span destination volumes just as they span source volumes. Save sets that fit on one original volume may spill over onto two clone volumes simply because the destination volumes may be a different size or may contain different data than the original. However, the opposite is also possible: The original save sets might require two volumes, while the cloned save sets require only one volume.

The Sun StorEdge EBS server allows only one clone of a save set per volume. Therefore, if you specify three clones of a save set, each clone is written to a separate volume.

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 a save set exist, clones of save sets on volumes in an autochanger are

generally selected as the source data, rather than a volume that is manually mounted. You can use the command prompt to specify the precise save set clone to use as the source.

Retention Policies

By default, save sets on the source and destination volumes have the same retention policies. However, you may override the preselected retention policy for the cloned volume by changing the mode of the volume to manual recycle in the Volumes window. For more information about changing the mode of a volume, see [“To Change the Mode of a Volume” on page 241](#).

When you apply the manual recycle mode to a volume, Sun StorEdge EBS software does not overwrite the volume unless you manually change the mode again to recyclable.

Performance Levels

To maintain backup and cloning performance levels, the Sun StorEdge EBS server can be configured to clone save sets automatically after completing a scheduled backup. You should allow time to complete the cloning process before scheduling another backup.

Online Index Entries

The Sun StorEdge EBS server creates index entries for the cloned save sets and volumes in the media database. However, both the original and clone volumes use the same set of entries in the client file index. If you eliminate either the clone or the original volume from the media database, the Sun StorEdge EBS server does not let you remove the client file index entries as long as one of the volumes still exists.

Devices

You must have at least two storage devices to clone data. One storage device is used for the source volume and the other for the destination volume.

- If you use libraries with multiple devices, the Sun StorEdge EBS server mounts the volumes required for cloning automatically.
- If you use stand-alone devices, you need to mount the volumes manually. The Sun StorEdge EBS server displays a message in the Messages window indicating which volumes to mount.

How Save Set Cloning Works

There are two different methods for cloning save set data:

- *Automatic* cloning of save sets can be set in the Group resource. For more information, see [“Automatically Cloning a Save Set” on page 350](#).
- *Manual* cloning of save sets is performed by entering the appropriate information in the Save Set Clone dialog box. For instructions on how to manually clone a save set, see [“To Manually Clone a Save Set” on page 347](#).

▼ To Manually Clone a Save Set

To manually clone a save set, you must 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. In the Sun StorEdge EBS Administrator program, select Clone from the Save Set menu.

Note – The text boxes in the Save Set Clone dialog box, shown in [FIGURE 9-2 on page 348](#), are case-sensitive.




FIGURE 9-2 Save Set Clone Dialog Box

2. For the Client text box, enter the name of the client associated with the save set.
If you do not enter a client name, Sun StorEdge EBS software locates and displays all save sets for all clients.
3. (Optional) For the Save Set text box, enter the name of the save set.
4. (Optional) For the Volume text box, enter the name of a volume.
5. Click More to display additional fields in the Save Set Clone window.
6. (Optional) For the Pool text box, enter the name of a volume pool.
7. For the Start Date and End Date, enter dates corresponding to when you believe the save set was created.

By default, *yesterday* is used for the start date, and the end date is left blank. This means that save sets backed up between yesterday at 12:01 a.m. and the current time will be displayed. For the Start and End Date fields, any of the following formats are acceptable:

- Written out completely (for example, November 1, 2003)
- Numerically as mm/dd/yy (for example, 11/01/03)
- As a phrase (for example, two weeks ago, three days ago)

8. (Optional) Select the appropriate status for the save sets. Browsable, Recyclable, Scanned-in, and Recoverable are already chosen by default. [TABLE 9-1](#) lists the seven save set status variations used by the Sun StorEdge EBS software.

TABLE 9-1 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 using the <code>scanner</code> 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.

9. (Optional) Select the appropriate level for the Maximum Level attribute to specify the backup levels of the save sets to be included in the query.

Incremental is specified by default. The Sun StorEdge EBS server searches for and displays all the levels up to and including the selected level. For example:

- If you select level 5, the software searches for and displays save sets backed up at levels Full, 1, 2, 3, 4, and 5.
- If you select level *Full*, the software searches for and displays only save sets backed up at level Full.
- If you select *Incremental*, the software searches for and displays save sets for all levels.

You must always have at least one level selected, and you may not select more than one level at a time.

10. (Optional) Use the Sort By field to select what order the save sets appear in the Save Sets list box.
11. Click Query. The save sets that fit your filtering criteria appear in the Save Sets list.

To view details of a selected save set, click Details to open the Save Set Clone Details box. This box provides information about the selected save set, such as the name and date, the number of files it contains, and its size and status. Click Close to close the box.
12. Select the save sets you want to clone from the Save Set list.
13. Click Clone. The Save Set Clone Status window appears, showing information about the status of the cloning operation.

Automatically Cloning a Save Set

The Sun StorEdge EBS server can automatically clone save set data at the time of a scheduled group backup. This is configured in the Group resource. For general information on configuring groups, including cloning, see [“Backup Groups” on page 250](#).

The backed-up save sets you identify for automatic cloning are cloned without any direct involvement by the administrator. Whenever you know of a save set cloning requirement in advance, try to use automatic save set cloning.

Automatic volume cloning is performed after all backup operations are complete. The savegroup completion report that is issued after a scheduled backup also reports on the success or failure of the cloning operation for each save set.

The location of the devices where the clone data is written is established 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 Cloning Save Sets

This section contains three different examples of save set cloning:

- [“Example: Manually Cloning a Save Set” on page 351](#)
- [“Example: Automatically Cloning One Client” on page 352](#)
- [“Example: Automatically Cloning Several Clients” on page 353](#)

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

Example: Manually Cloning 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 needs to send the data to another company located out of state that also uses Sun StorEdge EBS software. The Sun StorEdge EBS administrator needs to clone the most recent full backup, and any incrementals since the last full backup, to make sure that the most current data is sent.

The company receiving the cloned data needs to use the `scanner` program to retrieve the data. For more information about the `scanner` program, see [“The scanner Program” on page 316](#).

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. In the Sun StorEdge EBS Administrator program, select Clone from the Save Set menu.
2. In the Save Set Clone dialog box, enter the client name for the Client text box.
3. For the Save Set text box, enter the save set name.
4. 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.

5. Click Query.

All save sets that meet the browsing criteria appear in the Save Sets list.

6. Select the save sets to clone.

7. Click Details to determine the size of the data and the original volume that contains the data to be cloned.
 8. Mount the original volume.
 9. Click Clone.
-

Example: Automatically Cloning One Client

In this example, 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 the group `Documents`.
For more information on creating groups, see [“To Create a Group” on page 254](#).
2. Select the appropriate choices in the Groups dialog box. The required choices are:
 - Autostart = Enabled
 - Clones = Yes
 - Clone pool = Default Clone (or a custom clone pool)
3. Create the client `jupiter`.
For more information on creating clients, see [“To Create a Client” on page 110](#).
4. 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: Automatically Cloning Several Clients

FIGURE 9-3 on page 353 illustrates the level of control you can achieve when cloning data with Sun StorEdge EBS software. You can create an unlimited number of automatic cloning configurations using the Group, Schedule, and Client resources.




FIGURE 9-3 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 individual 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 needs to accomplish the following tasks:

- 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 9-2](#)).

TABLE 9-2 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 9-3](#)).

TABLE 9-3 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 in this example, the system administrator must:

1. In the Sun StorEdge EBS Administrator program, select Schedules from the Configure menu.
2. Create the *Accounting Incremental* schedule to perform incremental backups every day except Sunday, when it skips the backup (see [TABLE 9-2 on page 354](#)).
For details on creating schedules, see [“To Create a Schedule” on page 271](#).
3. Create the *Accounting Full-Clone* schedule to perform full backups every Sunday and skip backups the rest of the week (see [TABLE 9-3 on page 354](#)).

4. Create the *Accounting Incremental* group. The required attributes are:

- Autostart = Enabled
- Clones = No

For details on creating groups, see [“To Create a Group” on page 254](#).

5. Create the *Accounting Full-Clone* group. The required attributes are:

- Autostart = Enabled
- Clones = Yes
- Clone = Default Clone

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

For details on creating clients, see [“To Create a Client” on page 110](#).

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

- Schedule = Accounting Incremental
- Group = Accounting Incremental

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

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

Recovering Cloned Save Sets

When recovering cloned save sets, Sun StorEdge EBS software always looks for the original volumes by default. The volumes required for a save set recovery appear in the Required Volumes window. For more information, see [“Viewing the Volumes Required for Data Recovery” on page 294](#).

In order to recover save set files from a clone volume, rather than the original volume, you must first designate the original save set as *suspect*. Once the Sun StorEdge EBS server determines that the original save set is marked as suspect, it

looks for a cloned version. 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.

If you cannot determine which specific save set on the volume is the one you want to recover (in order to make it suspect), change the status of *all* of the save sets on the volume to suspect. This ensures that Sun StorEdge EBS software cannot find the save set it needs on the original volume during the recovery process.

Note – The Sun StorEdge EBS server displays only one set of volumes needed for recovery at a time. Therefore, even if more than one set of clones is available for recovery, only one set at a time appears in the Required Volumes window.

For information about changing the status of a save set, see [“To Change the Status of a Save Set” on page 119](#).

Other Information and Procedures

This section contains miscellaneous information and procedures that may be useful when cloning data with Sun StorEdge EBS software.

▼ To Clone Archived Data

To clone an existing archive volume:

1. In the Sun StorEdge EBS Administrator program, select Clone from the Save Set menu.
2. In the Save Set Clone dialog box, make selections for browsing save sets.
For more information, see [“To Manually Clone a Save Set” on page 347](#).
3. In the Pool text box, enter Archive.
4. Click Query to display the save sets matching the criteria.
5. Select the save sets to be cloned.
6. Click Clone.

Clone Storage Node Affinity

The link between storage nodes used to manage backup data and storage nodes used to clone that data is called *clone storage node affinity*.

The Clone Storage Nodes attribute, in the storage node's or server's Client resource, is used to define the storage nodes used for cloning data. Data is cloned from media that contains the original save sets to media on the specified clone storage node.

Note – The Clone Storage Node attribute applies to clients that are also serving as Sun StorEdge EBS storage nodes or servers. No information should be entered for clients that are not storage nodes or servers.

The Clone Storage Nodes attribute allows you to specify a network interface for storage nodes that perform cloning operations different than the one specified for the storage node's remote device. The server uses the exact hostname you specify for the attribute, instead of using the hostname prefix for the remote device name configured in the Device resource.

When a volume is being cloned, the Sun StorEdge EBS server checks the value of the Clone Storage Nodes attribute for that storage node client:

- If the Clone Storage Nodes attribute is a null value, then the value listed for the server's Clone Storage Nodes attribute is used.
- If the server's Clone Storage Nodes attribute also has a null value, then the server's Storage Node attribute is used.

Compatibility is maintained with the existing clone function that follows the server's Storage Node attribute.

To independently direct clones from each storage node, add the hostname of the storage node that you want to receive the directed clones to the Clone Storage Nodes attribute in the Client resource configured for the storage node. The first entry made on the list that has a functional, enabled device is selected to receive the cloned data from the storage node.

To direct clones from all storage nodes to the same destination, leave the Clone Storage Nodes attribute blank for the Client resources you configure for the storage nodes, and configure only the Sun StorEdge EBS server's Clone Storage Nodes attribute. This provides a single source of control for clone destination.

The client file index and media database entries for the save sets cloned to media on a remote device on a storage node still reside on the Sun StorEdge EBS server. This enforces the browse and retention policies in the same manner as for any cloned save sets that reside on the media in a device that is locally attached to the server.

Autochangers and Silos

This chapter provides information about how to configure and operate Sun StorEdge Enterprise Backup Software software with autochangers and silos. Working with autochangers and silos requires knowledge of installing and setting up storage devices. For additional information pertinent to all devices, including mounting and unmounting volumes, managing media, using pools, and cloning volumes, see [Chapter 5](#).

Using Sun StorEdge EBS software with an autochanger requires that you purchase and enable the Sun StorEdge EBS Autochanger Software Module. For more information on this product, contact a Sun StorEdge EBS sales representative or authorized Sun reseller. For general information on licensing Sun StorEdge EBS products, see the Sun StorEdge EBS Roadmap.

The main sections in this chapter are:

- [“Autochanger Installation and Configuration” on page 360](#)
- [“Editing Autochanger Configurations” on page 376](#)
- [“Dynamic Drive Sharing” on page 378](#)
- [“Sharing Autochangers among Sun StorEdge EBS Hosts” on page 384](#)
- [“Autochanger Operations” on page 391](#)
- [“Labeling Volumes with Barcodes” on page 404](#)
- [“Using Autochangers with a Volume Import/Export Capability” on page 409](#)
- [“Tips for Using Autochangers” on page 411](#)
- [“Silos” on page 415](#)
- [“Media Management in a Silo” on page 421](#)

Autochanger Installation and Configuration

Autochangers automate the task of loading, mounting, and labeling backup media. The term *autochanger* refers to a variety of storage devices, including the following: jukebox, carousel, library, near-line storage, datawheel, and autoloader. Note that tape “stackers” are not considered autochangers by the Sun StorEdge EBS software. Stackers access tapes sequentially, while autochangers can access media in any order.

Before you back up to an autochanger, perform the following tasks:

1. Install and enable the Sun StorEdge EBS software.
2. Install and enable the Sun StorEdge EBS Autochanger Software Module.
3. Load and label the volumes.

Some autochanger configuration occurs when you install and enable the Autochanger Software Module. After installing the autochanger, you can make changes to the available slots, devices, barcode choices, auto media management, and cleaning devices.

The Sun StorEdge EBS software does not support emulation mode for any storage device. It communicates with devices only in their native mode. For example, do not run an ADIC DLT in Exabyte 440 mode.

This section includes the following information:

- [“To View Autochangers” on page 360](#)
- [“To Add an Autochanger” on page 361](#)
- [“To Test the Autochanger Connection” on page 367](#)
- [“Testing Device Configuration Using the jbverify Program” on page 368](#)
- [“Configuring Autochanger Attributes” on page 371](#)
- [“To Add a Nonshared Device to an Autochanger” on page 371](#)
- [“Remote Autochanger Management” on page 374](#)
- [“Adjusting Sleep Periods for Autochanger Tasks” on page 375](#)

▼ To View Autochangers

To view the autochangers the Sun StorEdge EBS server recognizes:

1. In the Sun StorEdge EBS Administrator program, select Jukeboxes from the Media menu.
2. All autochangers connected to the Sun StorEdge EBS server appear in the Jukeboxes list box.

▼ To Add an Autochanger

You cannot use the Sun StorEdge EBS Administrator program to add or create an autochanger; instead you must use the `jbconfig` program at the command prompt. However, once the autochanger has been added and configured, you can use the Sun StorEdge EBS Administrator program to modify its configuration.

Note – In releases of the Sun StorEdge EBS software prior to 7.0, the `jbconfig` program listed all autochangers that were connected to the computer on which `jbconfig` was run, whether or not these autochangers were already configured in Sun StorEdge EBS. With the Sun StorEdge EBS 7.0 software, the `jbconfig` program lists only those autochangers that are connected but not configured.

To add an autochanger:

1. Log in as root on the Sun StorEdge EBS server.
2. Enter the following at the command prompt:

```
# jbconfig
```

The `jbconfig` program displays the following installation script and prompts you for responses. To accept the default answers, press [Enter].

For an example of the `jbconfig` program and appropriate responses, see [“Example: Configuring a SCSI Autochanger” on page 362](#) and [“Example: Configuring an SJI Autochanger” on page 364](#).

For more information about the autochangers that are connected to the server or related information that might be required to answer the `jbconfig` prompts, use the `inquire` and `sjisn` commands. For more information about using these commands, refer to the man pages for these commands, or to the *Sun StorEdge Enterprise Backup Software Installation Guide*.

3. If you plan to use a cleaning cartridge, enter the following at the command prompt:

```
# nsrjb -U number
```

where *number* is the number of times you want the Sun StorEdge EBS software to use the cleaning cartridge. Typically, the cartridge package lists the number of times a cleaning cartridge can be used.

Example: Configuring a SCSI Autochanger

The following example shows the required responses to configure a SCSI autochanger. The driver software detects and displays the information for all SCSI autochangers attached to the system. Examples of the answers to the questions are included in bold formatting.

```
# 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] 2

Removable Storage service not available.

Installing 'Exabyte 400 Series' jukebox - scsidev@1.0.0.

What name do you want to assign to this jukebox device? saturn

Turn NetWorker auto-cleaning on (yes / no) [yes]? yes

The following drives have been detected in this auto-changer:

    1> 8mm @ 1.1.0 ==> /dev/rmt/0cbn
    2> 8mm @ 1.2.0 ==> /dev/rmt/1cbn

These are all the drives that this auto changer possesses.

Do you want to change the model(s) or configure them as shared or
NDMP drives? (yes / no) [no] no

Should this device be configured as a dedicated storage node
(yes/no)? [no]
```

Note – All devices created on storage nodes (with the exception of servers) have a new Resource Administration Platform (RAP) attribute: Dedicated Storage Node. For information about the Dedicated Storage Node, see [“Dedicated Storage Node” on page 231](#).

Jukebox has been added successfully

The following configuration options have been set:

- > Jukebox description to the control port and model.
- > Autochanger control port to the port at which we found it.
- > Networker managed tape autocleaning on.
- > Barcode reading to on.
- > Volume labels that match the barcodes.
- > Slot intended to hold cleaning cartridge to 1. Please insure that a cleaning cartridge is in that slot

- > Number of times we will use a new cleaning cartridge to 12.
- > Cleaning interval for the tape drives to 6 months.

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

Example: Configuring an SJI Autochanger

The following example shows the responses to configure an SJI autochanger.

```
# 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) Breece Hill Saguaro
- 5) Breece Hill
- 6) Philips Blackjack
- 7) DLI Libra Series
- 8) Quantum DLT/Digital DLT
- 9) Exabyte 10e or 10h
- 10) Exabyte 10i
- 11) Exabyte 18D
- 12) Exabyte 60
- 13) Exabyte 120
- 14) Exabyte 210
- 15) Exabyte 218
- 16) Exabyte 220
- 17) Exabyte 230D
- 18) Exabyte 400 Series
- .
- .
- .

Choice? **18**

Installing an 'Exabyte 400 Series' jukebox.

Name you would like to assign to the jukebox device? **exb**

Pathname of the control port for the jukebox device? scsidev@1.0.0
[Return]

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]
[Return]

Is any drive going to have more than one path defined? (yes / no)
[no] [Return]

Please enter the device path information in one of the following
formats:

```
    /dev/rmt/0cbn --for local path or  
    host:device-path --for remote node or  
    host:drive-letter:directory path --for Windows  
    disk file
```

Drive 1, element 1

Drive path ? /dev/rmt/0cbn

Drive 2, element 2

Drive path ? /dev/rmt/1cbn

Please select the appropriate drive type number:

1) 3480	17) 9490	33) optical
2) 3570	18) 9840	34) qic
3) 3590	19) 9940	35) SD3
4) 4890	20) adv_file	36) sdlt
5) 4mm	21) dlt	37) sdlt320
6) 4mm 12GB	22) dlt1	38) SLR
7) 4mm 20GB	23) dlt7000	39) tkz90
8) 4mm 4GB	24) dlt8000	40)
travan10		
9) 4mm 8GB	25) dst (NT)	41) tz85
10) 8mm	26) dtf	42) tz86
11) 8mm 20GB	27) dtf2	43) tz87
12) 8mm 5GB	28) file	44) tz88
13) 8mm AIT	29) hint	45) tz89
14) 8mm AIT-2	30) logical	46) tz90
15) 8mm AIT-3	31) LTO Ultrium	47) tzs20
16) 8mm Mammoth-2	32) LTO Ultrium-2	48) VXA

Enter the drive type of drive 1? **10**

Are all the drives the same model? (yes / no) [yes][Return]

Jukebox has been added successfully

The following configuration options have been set:

```
>Autocleaning off.
>Barcode reading to on.
>Volume labels that match the barcodes.
>Slot intended to hold cleaning cartridge to 1.
>Please insure that a cleaning cartridge is in that slot
>Number of times we will use a new cleaning cartridge to
>12.
```

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] [Return]

Configuring Autochangers on a Storage Node

To configure an autochanger on a storage node, follow the instructions for [“To Add an Autochanger” on page 361](#), except execute the `jbconfig` program on the storage node, using the `-S server` option, where *server* is the host on which the Sun StorEdge EBS server resides. This command creates the `nsrjb.res` file. For information about adding devices to the autochanger, see [“To Add a Remote Autochanger Device” on page 375](#).

Note – To configure on autochangers on a storage node, you must be logged in as root and you must have the Configure Sun StorEdge EBS privilege. For more information about the Configure Sun StorEdge EBS privilege, see [“Sun StorEdge EBS User Groups” on page 61](#).

▼ To Test the Autochanger Connection

After configuring the autochanger, test the connection. To test the autochanger connection:

1. Log in as root on the Sun StorEdge EBS server or storage node.
2. Insert two volumes, one each into the first and last slots of the autochanger. Make sure that the devices are empty and that any device doors are open.
3. Enter the following command at the prompt:

```
# jbxercise
```

Specify the control port and the device type. The control port for SCSI autochanger models is typically expressed in the format `scsidev@n.n.n`. You can obtain the exact control port pathname from the response displayed by the `jbconfig` program script:

```
These are the SCSI Jukeboxes currently attached to your system:
```

- ```
1) scsidev@1.2.0: DLI Libra Series
2) scsidev@0.2.1: Quantum DLT/Compaq DLT
```

For example, the following command runs the `jbexercise` program on the Quantum DLT/Compaq DLT autochanger detected by the `jbconfig` program:

```
jbexercise -c scsidev@0.2.1 -m "Quantum DLT/Compaq DLT"
```

For more information about `jbexercise`, refer to the *Legato Command Reference Guide* or to the `jbexercise` man page.

## Testing Device Configuration Using the `jbverify` Program

You can use the `jbverify` program to check the devices used in the Sun StorEdge EBS environment. This program ensures that each device is configured properly by checking for accessibility and usability.

By default, `jbverify` **checks all the devices in the Sun StorEdge EBS database. However, you can use the command line options described in [TABLE 10-1 on page 368](#) to limit the devices that the command checks:**

**TABLE 10-1** The `jbverify` Command Line Options for Devices

| Option | Devices that Are Checked                                  |
|--------|-----------------------------------------------------------|
| -j     | Autochangers                                              |
| -d     | Stand-alone drives that are not contained in autochangers |
| -l     | Local devices                                             |
| -J     | Specific autochangers                                     |
| -D     | Specific drives                                           |

---

**Note** – The `jbverify` program does not test Legato GEMS SmartMedia devices.

---

For further information regarding the `jbverify` program, refer to the *Legato Command Reference Guide* or to the `jbverify` man page.

## Requirements for Running the `jbverify` Program

The following sections describe the requirements of the `jbverify` program.

## Requirements for Servers and Clients

The following requirements apply to running `jbverify` on Sun StorEdge EBS servers and clients:

- To run the `jbverify` program on Sun StorEdge EBS servers, the server daemon, `nsrd`, must be running.
- To run the `jbverify` program on Sun StorEdge EBS clients, the client daemon, `nsrexecd`, must be running. You must also use the `-S server` command line option to identify the server:

```
jbverify -s server
```

## Requirements for Storage Nodes

The `jbverify` program can be run on any storage node and can be used to test any device on that storage node. The device, however, must have been configured in the Sun StorEdge EBS software.

The `jbverify` program can also be used to remotely test devices from a storage node.

In such a case, you must complete *one* of the following procedures:

- The target storage node must have the invoking storage node listed in its `servers` file. For information about editing the `servers` file, see [“To Edit the servers File” on page 154](#).
- When the Sun StorEdge EBS client daemon, `nsrexecd`, is started on the target storage node, the invoking storage node must be included in the target storage node’s `servers` file. For information about editing the `servers` file, see [“To Edit the servers File” on page 154](#).

## ▼ To Test an Autochanger

For autochangers, the `jbverify` program performs the load and unload operations on each drive without performing any write operations, unless the `-t` option is used. If no slot is specified on the command line, the `jbverify` program checks all the slots that are available to the Sun StorEdge EBS server and then loads the first available slot.

To test only autochangers, enter the `jbverify` command with the `-j` option:

```
jbverify -j
```

To use a specific slot for the test, use the **-S** option:

```
jbverify -S slot
```

To test a specific jukebox, use the **-J** option:

```
jbverify -J jukebox_name
```

## ▼ To Test a Stand-Alone Device

To test a stand-alone device, without testing any other devices, use the **-d** option:

```
jbverify -d
```

To test a specific device, use the **-D** option:

```
jbverify -D device_name
```

## ▼ To Test a Local Device

To test a local device, without testing any other devices, use the **-l** option.

```
jbverify -l
```

## Testing Tapes with the jbverify Program

The **jbverify** program can test tapes that are loaded into the drives that are being tested. The program tests these tapes by using the Sun StorEdge EBS tape exercise program (**tapeexercise** on UNIX or **tapeexer.exe** on Microsoft Windows).

For further information regarding the **tapeexercise** and **tapeexer.exe** programs, refer to the *Legato Command Reference Guide* or to the **tapeexercise** man page.

To test a tape using the **jbverify** program, use the **-t** option:

```
jbverify -t
```

# Configuring Autochanger Attributes

Attributes for a specific autochanger are configured in the Jukebox resource. To view Jukebox resources, in the Sun StorEdge EBS Administrator program select Jukeboxes from the Media menu.

## Key Autochanger Attributes

[TABLE 10-2](#) lists some of the key autochanger attributes and their configuration.

**TABLE 10-2** Autochanger Attributes

| Attribute             | Configuration                                                                                                                                                                                                                                                      |
|-----------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Control Port          | Enter the pathname of the autochanger's command control port.                                                                                                                                                                                                      |
| Devices               | Enter the names of each device in the autochanger.                                                                                                                                                                                                                 |
| Barcode Reader        | Select Yes to enable this function for an autochanger with barcode-reading capability. In some cases, <code>jbconfig</code> sets this attribute to Yes, even though the autochanger does not have a barcode reader. If this is the case, set this attribute to No. |
| Match Barcode Labels  | Select Yes for a barcode-enabled autochanger to match barcode labels and volume labels.                                                                                                                                                                            |
| Cleaning Slots        | Enter the slot number containing the cleaning cartridge.                                                                                                                                                                                                           |
| Default Cleanings     | Enter the number of cleanings the cleaning cartridge can handle before the Sun StorEdge EBS server notifies you to replace the cartridge.                                                                                                                          |
| Auto Clean            | Select Yes to enable autocleaning.                                                                                                                                                                                                                                 |
| TapeAlert             | Enable common device interface (CDI).                                                                                                                                                                                                                              |
| Auto Media Management | Select Yes to enable auto media management, which includes automatic labeling of blank volumes.                                                                                                                                                                    |
| Available Slots       | Enter the range of autochanger slots available for the Sun StorEdge EBS server to use.                                                                                                                                                                             |

## ▼ To Add a Nonshared Device to an Autochanger

Before you can add devices to an autochanger, you must first configure the Sun StorEdge EBS server to recognize the new autochanger device. For instructions on creating an autochanger device, see [“To Add a Device” on page 192](#).



---

**Caution** – When you are creating an autochanger device, you must set the Jukebox Device attribute to Yes. The Jukebox Device attribute is a hidden attribute; for information about displaying hidden attributes, see [“Hidden Attributes” on page 42](#).

---

To add a nonshared device to a library:

1. Check the autochanger to make sure that all cartridges have been removed from all the drives.

---

**Note** – Do not continue with this procedure until all cartridges have been removed.

---

2. Select Jukeboxes from the Media menu.
3. In the Jukeboxes dialog box, select the appropriate autochanger.
4. In the Device attribute, enter the new device pathname of the autochanger’s tape devices, making sure that it is listed in the correct location and click Add.

If an autochanger has more than one device, you must list the device pathnames in the same order as their physical location in the autochanger. If you are unsure of their order, do one of the following:

- Refer to the autochanger hardware manual.
- Log in as root on the Sun StorEdge EBS server or storage node computer and enter the `inquire` command at the shell prompt. The `inquire` command returns a list of attached SCSI devices.

For an example of autochanger device pathname naming, see [FIGURE 10-1 on page 373](#).

5. In the Number of Devices attribute, increase the current number by one.
6. In the Number of Drives attribute, increase the current number by one.
7. Enter an additional empty line for each of the following attributes:
  - Loaded Volumes
  - Loaded Barcodes
  - Loaded Slots

The number of lines in these attributes must equal the Number of Devices attribute (specified in [Step 5](#)). This should be equivalent to the number of device pathnames in the Devices list.

You can make other modifications to the autochanger configurations at this time. For instance, you might select barcode labels or decide to change the range of available slots.



---

**Caution** – You must reset an autochanger every time changes or modifications are made to the device. For information about how to reset an autochanger, refer to [“To Reset an Autochanger” on page 374](#).

---

### *Example: Autochanger Device Pathnames*

FIGURE 10-1 shows an example of pathname naming.




**FIGURE 10-1** Autochanger Device Pathnames

## ▼ To Edit Autochanger Attributes

To edit an autochanger's attributes:

1. In the Sun StorEdge EBS Administrator program, select Jukeboxes from the Media menu.
2. In the Jukeboxes dialog box, select the appropriate autochanger.
3. Edit the attributes. To learn more about each attribute, select On Window from the Help Menu.
4. Click Apply when complete.

## ▼ To Delete an Autochanger

To delete an autochanger:

1. In the Sun StorEdge EBS Administrator program, select Jukeboxes from the Media menu.
2. In the Jukeboxes dialog box, select the appropriate autochanger.
3. Click Delete.

When prompted, click OK to confirm the deletion.

4. Click Apply.

## ▼ To Reset an Autochanger

You must reset an autochanger every time you modify the device. Use the `nsrjb -HE` or `sjiielm` commands to reset an autochanger.

For example, the autochanger inventory must be correct after adding drives to an SJI-compliant autochanger, 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 autochanger. This must be done before running `jbconfig`. The `-E` option reinitializes the autochanger's element status. Some autochangers can keep track of whether there is media in a component in the autochanger. This feature is known as an *element status* capability.

## Remote Autochanger Management

Autochangers that are connected to storage nodes require additional configuration and management steps.

You can control most operations on remote autochanger devices from the Sun StorEdge EBS Administrator program. For some remote autochanger operations, such as reset, you must use the `nsrjb` or `jbconfig` programs on the storage node. You can issue the commands as root, either from the local computer or through a remote login session.

After you install the storage node binaries on the storage node, define the storage node's devices. For instructions on defining devices, see [“To Configure a Storage Node” on page 228](#). For instructions on adding autochanger devices, see [“To Add a Remote Autochanger Device” on page 375](#).



## ▼ To Add a Remote Autochanger Device

To add a remote autochanger device:

1. In the Sun StorEdge EBS Administrator program, select Server Setup from the Server menu.
2. Add the storage node's hostname to the Administrator attribute in the Server resource using the following syntax:

```
root@storage_node_hostname
```

3. Run the `jbconfig` program on the storage node computer to define each device in the autochanger.

For more information on installing and configuring autochangers, see [“Autochanger Installation and Configuration” on page 360](#), and refer to the *Legato Command Reference Guide* or the `jbconfig` `man` page.

---

**Note** – The device names for remote devices begin with a prefix of “rd=” and the name of the storage node computer. For example, `rd=jupiter:/dev/rmt/1mbnis` is a device called `/dev/rmt/1mbn` on a storage node computer called `jupiter`.

---

## Adjusting Sleep Periods for Autochanger Tasks

The Jukebox resource includes attributes that specify the number of seconds that the autochanger is inactive after loading, unloading, or ejecting a volume. This function is known as *sleeping*, and while sleeping, the autochanger is restricted from receiving or performing other operations.

This period of inactivity is sometimes necessary when loading, unloading, and ejecting volumes on some autochangers whose devices take longer than the usual time to perform the task. Without the sleep period, the loading or unloading of volumes might fail.

When you configure a Jukebox resource using the `jbconfig` program, Sun StorEdge EBS software automatically configures default sleep periods. Only change these values if you are troubleshooting the autochanger's performance, or if a Sun StorEdge EBS technical support specialist instructs you. Typically, the higher the sleep values specified in the attributes, the longer it takes the autochanger to perform the task. Be cautious when changing these values.

[TABLE 10-3 on page 376](#) lists the attributes related to the autochanger sleep function.

**TABLE 10-3** Autochanger Resource Sleep Attributes

| Attribute    | Description                                                                                                                                                                                                                                                                                                                                       |
|--------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Eject Sleep  | Specifies the number of seconds the Sun StorEdge EBS software sleeps after ejecting a volume. The Sun StorEdge EBS software uses this attribute if the autochanger is not configured for volume autoejection. Autochangers configured for volume autoejection do not perform a separate eject operation, so this attribute has no effect on them. |
| Unload Sleep | Specifies the number of seconds an autochanger sleeps between unloading a volume and moving the volume from the device to a slot. The Sun StorEdge EBS software uses this attribute for all autochangers when a volume is unloaded.                                                                                                               |
| Load Sleep   | Specifies the number of seconds the Sun StorEdge EBS software sleeps between loading a volume and trying to access that volume's data.                                                                                                                                                                                                            |

## ▼ To Adjust Sleep Periods for Autochanger Tasks

To adjust sleep periods for autochanger tasks:

1. In the Sun StorEdge EBS Administrator program, select Jukeboxes from the Media menu.
2. In the Jukeboxes dialog box, select Details from the View menu to display hidden attributes.
3. Select the appropriate autochanger.
4. Adjust one or more of the sleep-related attributes.

---

## Editing Autochanger Configurations

In a large SAN environment, the ability to add or remove devices—or storage nodes that share a drive—without having to reconfigure the autochanger saves time and facilitates backups. The `jbedit` program enables such changes without having to delete the existing configuration and rerun the `jbconfig` program to reconfigure the autochanger. The `jbedit` program 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 autochanger.

The `jbedit` program, which works with Sun StorEdge EBS release 6.x or later, supports all direct-attached SCSI/SII, SAN, and NDMP autochangers. It does not currently support AlphaStor, SmartMedia, or silo autochangers.

The `jbedit` program is not intended to be a full-fledged editor of the autochanger resource. It primarily serves to facilitate adding or deleting drives or device definitions. The editing of all autochanger resource attributes—except for the device paths—is still done through the Sun StorEdge EBS Administrator program. The `jbedit` menu provides selection lists that make it easy to find the drives or devices you want to add or delete.

---

**Note** – Running the `jbedit` program requires Configure Sun StorEdge EBS privileges.

---

You *must* run the `jbedit` program on the computer with the attached autochanger. You can, however, use the `-o` option to run it a storage node that is not directly attached to the autochanger. This option prevents the `jbedit` program from trying to communicate with the autochanger.

## Using the `jbedit` Program

[TABLE 10-4 on page 377](#) lists some of the options most commonly used with the `jbedit` program.

**TABLE 10-4** 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               |
| -o     | Overrides information reported by the autochanger. |

For a detailed description of the `jbedit` command, its options, and associated diagnostic messages, refer to the `jbedit` man page.

## Limitations

Under certain conditions, you cannot modify the autochanger resource. For example:

- A corrupted Sun StorEdge EBS database must be repaired before the `jbedit` program can function properly.

- The autochanger configuration cannot be edited while an autochanger operation (load, unload, inventory, reset, or move) is in process.

Note: The **save** and **recover** programs are not autochanger operations.

- Drives or devices cannot be deleted if they contain volumes. This is to ensure that the autochanger can be left in a stable and usable state after a deletion.

---

## Dynamic Drive Sharing

Dynamic drive sharing (DDS) is a feature that allows Sun StorEdge EBS software 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.

This section contains the following topics:

- [“To Enable DDS When Configuring an Autochanger” on page 378](#)
- [“To Stop Sharing a Device within an Autochanger” on page 382](#)
- [“To Change the Idle Device Timeout Attribute” on page 383](#)

For further information about the DDS feature, see [“Dynamic Drive Sharing” on page 198](#).

For information explaining how to enable DDS with the `jbconfig` program, see [“To Enable DDS When Configuring an Autochanger” on page 378](#).

### ▼ To Enable DDS When Configuring an Autochanger

You can enable the DDS feature when you configure an autochanger through the `jbconfig` program.

---

**Note** – Before you enable DDS through the `jbconfig` program, you must ensure that each node (or server) that is to share a drive is licensed. Without licensing before configuration, the `jbconfig` program fails to configure the autochanger. For licensing information, the Sun StorEdge EBS Roadmap.

---

To enable DDS when you configure an autochanger:

1. Log in as root.

2. Enter the following command:

```
jbconfig
```

3. Select the appropriate autochanger for the system and complete the prompts as required.
4. Enter Yes when prompted whether the drive is to be shared.
5. Complete the prompts as required for the first host to have access to the shared drive.
6. Enter Yes when prompted whether the drive is to be shared with another host
7. When prompted, enter the name of the second device to have access to the shared drive.
8. Complete the prompts for the second device.

## Sample jbconfig Output

This sample `jbconfig` output is for the following configuration:

- Server (jupiter)
- A storage node (sn\_1)
- A shared IBM 3590 jukebox with one drive (drive\_1)

```
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] 2
```

```
Installing 'Standard SCSI Jukebox' jukebox - scsidev@5.0.6
```

```
What name do you want to assign to this jukebox device? IBM
```

```
Turn NetWorker auto-cleaning on (yes / no) [yes]? [yes]
```

```
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]**

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/0cbn --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 ? **/dev/ntape/tape34\_d1**

Device path 2 ? [] **jupiter:/dev/rmt1.1**

Device path 3 ? []

Drive 2, element 2

Device path 1 ? **/dev/ntape/tape33\_d1**

Device path 2 ? [] **jupiter:/dev/rmt0.1**

Device path 3 ? []

Please select the appropriate drive type number:

|                   |                     |             |
|-------------------|---------------------|-------------|
| 1) 3480           | 17) 9490            | 33) optical |
| 2) 3570           | 18) 9840            | 34) qic     |
| 3) 3590           | 19) 9940            | 35) SD3     |
| 4) 4890           | 20) adv_file        | 36) sdlt    |
| 5) 4mm            | 21) dlt             | 37) sdlt320 |
| 6) 4mm 12GB       | 22) dlt1            | 38) SLR     |
| 7) 4mm 20GB       | 23) dlt7000         | 39) tkz90   |
| 8) 4mm 4GB        | 24) dlt8000         | 40)         |
| travan10          |                     |             |
| 9) 4mm 8GB        | 25) dst (NT)        | 41) tz85    |
| 10) 8mm           | 26) dtf             | 42) tz86    |
| 11) 8mm 20GB      | 27) dtf2            | 43) tz87    |
| 12) 8mm 5GB       | 28) file            | 44) tz88    |
| 13) 8mm AIT       | 29) himt            | 45) tz89    |
| 14) 8mm AIT-2     | 30) logical         | 46) tz90    |
| 15) 8mm AIT-3     | 31) LTO Ultrium     | 47) tzs20   |
| 16) 8mm Mammoth-2 | 32) LTO Ultrium-248 | VXA         |

Enter the drive type of drive 1? **32**

Are all the drives the same model? (yes / no) **[yes]**

A Dedicated Storage Node can backup only local data to its devices.

Should jupiter be configured as a Dedicated Storage Node? (yes / no) **[no]**

Jukebox has been added successfully

The following configuration options have been set:

```

> Jukebox description to the control port and model.
> Autochanger control port to the port at which we found
> it.
> Networker managed tape autocleaning on.
> 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. Your jukebox does not report
> that it has a barcode reader, but it does report that
> it can handle volume tags. Some jukeboxes that have
> barcode readers report this way.
> Volume labels that match the barcodes.
> Slot intended to hold cleaning cartridge to 24.
> Please insure that a cleaning cartridge is in that
> slot
> Number of times we will use a new cleaning cartridge
> to 50.
> Cleaning interval for the tape drives to 6 months.

```

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]**

## ▼ To Stop Sharing a Device within an Autochanger

To stop sharing a DDS-configured device within an autochanger:

1. Delete the autochanger:
  - a. In the Sun StorEdge EBS Administrator program, select Jukeboxes from the Media menu.
  - b. Select the appropriate autochanger.
  - c. Click Delete.
  - d. When prompted, click OK to confirm the deletion.



- e. Click Apply.
2. Remove the value from the Hardware ID attribute in the Device resource:
  - a. In the Sun StorEdge EBS Administrator program, select Devices from the Media menu.
  - b. Select the appropriate autochanger.
  - c. Delete the value in the Hardware ID attribute.
  - d. Click Apply.
3. Reconfigure the autochanger using the `jbconfig` program.



---

**Caution** – When you stop sharing a DDS-configured drive, this does not mean that the drive is disabled. This means that Sun StorEdge EBS software is no longer aware that a specified drive is shared.

To avoid affecting other Sun StorEdge EBS operations when you stop sharing a device, remove the other device instances.

---

## ▼ To Change the Idle Device Timeout Attribute

You might notice that a tape remains mounted in its drive for 10 to 15 minutes after a backup has completed. Other requests for the drive from another device path wait during this period.

You can modify this timeout value by adjusting the Idle Device Timeout attribute. To access this attribute:

1. In the Sun StorEdge EBS Administrator program, select Jukeboxes from the Media menu.
2. In the Jukeboxes dialog box, select the appropriate autochanger.
3. Select Details from the View menu to display hidden attributes.
4. Scroll down to the Idle Device Timeout attribute and select a time that is appropriate for the system configuration.

The default Idle Device Timeout value is 10 minutes.

---

**Note** – Selecting zero means that the device never times out. Consequently, you would have to eject the tape manually.

---

---

# Sharing Autochangers 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 an autochanger, a feature known as *autochanger sharing*.

The presence of a SAN within the datazone is not required for autochanger sharing.

## Autochanger Sharing Licensing and Availability

Licensing and availability of the Sun StorEdge EBS autochanger sharing feature varies with the type of Sun StorEdge EBS license you have.

- Sun StorEdge EBS Power Edition and Network Edition — Autochanger sharing is automatically enabled as part of Sun StorEdge EBS Power and Network Editions. However, you must still have a valid Sun StorEdge EBS autochanger license for each autochanger.
- Sun StorEdge EBS Business Edition and WorkGroup Edition — Autochanger sharing is not available for Sun StorEdge EBS Business and WorkGroup Editions. If the datazone environment requires autochanger sharing, upgrade the Sun StorEdge EBS license to Network Edition or Power Edition.

Licensing autochanger sharing is similar to licensing other Legato NetWorker Modules and add-on products. Only a single autochanger sharing license is required for all autochangers in the Sun StorEdge EBS datazone, regardless of the number of autochangers present. However, you have to license each autochanger to enable the autochanger functionality. For general information about licensing Sun StorEdge EBS products, the Sun StorEdge EBS Roadmap.

---

**Note** – Only Sun StorEdge EBS Network Edition requires you to license autochanger sharing, since this feature is already enabled in Power Edition, and is not available with Business and Workgroup Editions.

---

# How Autochanger Sharing Works

Autochanger sharing allows one Sun StorEdge EBS host to control the autochanger's robotic arm, while other Sun StorEdge EBS hosts (as well as the host controlling the robotic arm) can each control and use specific autochanger devices. A specific device can be controlled only by a single Sun StorEdge EBS host.

FIGURE 10-2 on page 385 shows how multiple Sun StorEdge EBS hosts can share autochanger devices.




FIGURE 10-2 How Autochanger Sharing Works

## Autochanger Sharing Configuration

You can configure an autochanger for sharing during the initial autochanger setup with the `jbconfig` program. The `jbconfig` program must be run from the Sun StorEdge EBS host that controls the autochanger (controls the robotic arm). The `jbconfig` program prompts you to enter the Sun StorEdge EBS hostname controlling the first autochanger device. The default hostname value is the local Sun StorEdge EBS host on which `jbconfig` is running. However, you can change the default hostname value to another Sun StorEdge EBS host within the datazone. The `jbconfig` program prompts you to enter the controlling host for each device in the autochanger.

## Example: The jbconfig Program Additions

The following example shows the additions to the `jbconfig` program for the Sun StorEdge EBS server `neptune` and a storage node `pluto`. In this example, `jbconfig` is being run on the server `neptune`, which also controls the robotic arm. However, the first media device of the autochanger is being assigned to the storage node `pluto`. The user responses are in bold.

```
Enter hostname that owns media drive 1: ? [neptune] pluto
Enter pathname of media drive 1: ?
/dev/rmt/1cbn

using 'rd=pluto:/dev/rmt/1cbn' as device name
```

---

For more information on the `jbconfig` program, refer to the `jbconfig` man page or the *Legato Command Reference Guide*.

## Device Naming during Configuration

The `jbconfig` program automatically generates the device names for a shared autochanger according to the following conditions:

- If a device is controlled by the Sun StorEdge EBS server, only local device syntax is used (`device_path`).
- If the device is controlled by any other host, the hostname and device path are used to construct a device name that follows the existing remote device syntax (`rd=hostname:device_path`).

## Host Device Selection for Recoveries

The Read Hostname attribute in the Jukebox resource designates a host for use in recoveries when the host controlling a required device is not listed in the Storage Nodes attribute of the Client resource. The default value for the Read Hostname attribute is the Sun StorEdge EBS host that is controlling the first device in the shared autochanger. However, you can change the host to another value contained in the autochanger.

This functionality designates an alternative device to use for recoveries and cloning if a storage node's preferred hosts are unavailable. Normally, when a recovery is requested, the Sun StorEdge EBS server matches the host that is controlling the

device required for the recovery with the hosts listed in the Storage Nodes attribute in the Client resource. After the first match, the recovery or clone operation takes place. If there is no match, then the request reverts to the host listed in the Read Hostname attribute.

## Example: Data Recovery Using a Shared Autochanger

FIGURE 10-3 on page 387 displays an example of a Sun StorEdge EBS datazone environment that consists of a server `venus`, a client `earth`, and a storage node `mars`. There is a shared autochanger with the first device connected to `venus`, and a second device connected to `mars`.




FIGURE 10-3 Data Recovery Using a Shared Autochanger

The client `earth` requests a data recovery from a volume in the shared autochanger. The needed volume is not currently mounted in any device. To initiate the recovery, the server `venus` searches the hostname entries listed in the Storage Nodes attribute of `earth's` Client resource for a match with one of the hostname devices in the autochanger. A hostname match is made for storage node `mars`. The needed volume is mounted on that device, and the recovery occurs.

If a hostname match was not found, then `venus` would refer to the hostname value listed in the Read Hostname attribute of the Jukebox resource. By default, this value is `venus`. The needed volume is then mounted in the device controlled by `venus`, and the recovery occurs.

---

## Host Device Selection for Cloning

When a cloning request is received, the Sun StorEdge EBS server uses the host listed in the Read Hostname attribute of the Jukebox resource to read the volume being cloned. If the required volume is already mounted in a device, then the Sun StorEdge EBS server accesses that particular device.

## Labeling, Loading, and Unloading Volumes on Shared Autochanger Devices

Autochanger sharing also affects tasks related to volumes present in autochanger devices, such as:

- Labeling volumes
- Loading and unloading volumes for labeling, client backups, and data recoveries

The Sun StorEdge EBS command to manage autochanger volumes, `nsrjb`, has been modified to allow management of volumes on shared autochanger devices controlled by other hosts. The `nsrjb` command must always be run from the Sun StorEdge EBS host controlling the autochanger (such as having control of the robotic arm) regardless of whether the autochanger is shared or not.

---

**Note** – When you use the `nsrjb` command with the `-f devicename` option and the device name contains backslashes, you must use the backslash “escape” character with the backslash. For example, a device with the name `rd=sn1:\\.\\Tape0` would be entered at the command prompt as:

---

```
rd=sn1:\\\\.\\Tape0
```

## Managing Volumes on Server Devices

For devices controlled by the Sun StorEdge EBS server, you manage autochanger volumes using the `nsrjb -f` command with the device path, even if the `jbconfig` program for the shared autochanger was run from another host. For example, to load a volume into a shared autochanger device that is controlled by the Sun StorEdge EBS server, you would enter the following command at the prompt:

```
nsrjb -l -f device_path volume_name
```

---

**Note** – You can also use the `-J host_name` option instead of the `-f device_path` option to identify the hostname of the computer to which the jukebox is attached. The `-J` option is only supported when used in conjunction with the `-l` and `-L` options, and cannot be used with the `-f device_path` option. For more information, refer to the `nsrjb` man page or the *Command Reference Guide*.

---

## Managing Volumes on Storage Node Devices

When a shared autochanger is defined within the datazone, you must specify the full remote device syntax name when using the `nsrjb -f` command for devices controlled by a storage node:

```
nsrjb -l -f rd=host:device_path volume_name
```

where `host` is the hostname of the appropriate node that controls the device.

Notes:

- In a datazone with no shared autochangers, the Sun StorEdge EBS server automatically adds the remote device host (`rd=host`) information; therefore, you only need to input the local device syntax (`device_path`). However, in a datazone with one or more shared autochangers, you must always enter remote device syntax (`rd=host:device_path`) information with the device path.
- You can also use the `-J host_name` option instead of the `-f device_path` option to identify the hostname of the computer to which the jukebox is attached. The `-J` option is only supported when used in conjunction with the `-l` and `-L` options, and cannot be used with the `-f device_path` option. For more information, refer to the `nsrjb` man page or the *Command Reference Guide*.

## Automatic Unmounting of Volumes

At times, a volume that is mounted in one device may be needed by another device on the same autochanger. For example, data being recovered by one device may span more than one volume, and the required volume is mounted on another device. To address this need, you can define a value in the Idle Device Timeout attribute in the Jukebox resource for that autochanger. The Idle Device Timeout attribute, which applies only to shared autochangers or autochangers that are managed by Legato GEMS SmartMedia 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. For information about changing the Idle Device Timeout attribute, see [“To Change the Idle Device Timeout Attribute” on page 383](#).

## Setting Compatible Block Size for UNIX and Microsoft Windows

There are different SCSI hardware limitations between UNIX and Microsoft Windows operating systems, which can lead to compatibility problems. For example, if you have a device defined in UNIX and it 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 this device. If you are going to use both operating systems, you must determine a common block size acceptable to both.

### ▼ To Determine the Allowable Block Size

Use the scanner program to determine the allowed block size. Enter the following command at the prompt:

```
scanner -v device_name
```

where *device\_name* is the name of the device as defined for the server. The block size is included in the following output from this command:

```
scanner: volume id 2641870849 record size 98304
```

The record size entry is the block size of the device, in bytes.



## Solutions

To solve problems with block size compatibility:

- For UNIX, you can 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 of the Device resource in the Sun StorEdge EBS Administrator program). 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, you can install a later model HBA or upgrade to drivers that can support up to 128 KB blocks for the previous UNIX example. Windows also accepts the same environment variable format for setting block size as listed for UNIX.

---

**Note** – You must restart the Sun StorEdge EBS server for the environment variables to take effect.

---

---

## Autochanger Operations

This section includes the following information for performing autochanger operations:

- [“To Designate Available Autochanger Slots” on page 392](#)
- [“Maintaining Autochanger Devices” on page 393](#)
- [“Auto Media Management with Autochangers” on page 399](#)
- [“Mounting Autochanger Volumes” on page 400](#)
- [“To Label an Autochanger Volume” on page 402](#)

You may need to perform some of these tasks on a daily basis and others less frequently. The frequency with which you perform these tasks depends on:

- The number of volumes you label at one time.
- How often you change the contents of the autochanger.
- How often you need to mount volumes for manual backups.

## ▼ To Designate Available Autochanger Slots

The available slots feature allows you to control which volumes the Sun StorEdge EBS server uses for backing up data. The server uses all of the volumes in the autochanger to perform recoveries. However, you can control the volumes automatically selected for backups by designating a range of available slots in the autochanger.

For example, you may designate slots 1 through 5 for Sun StorEdge EBS backups in an autochanger that contains 10 slots. The entries can be a range of slot numbers or a single slot number. The range of available slots you enter is the same for all of the storage devices in the autochanger. You cannot enter a different range for each device located in an autochanger.



---

**Caution** – Make sure you have placed volumes in all the available slots of the autochanger so the Sun StorEdge EBS server can proceed uninterrupted with an automatic backup.

---

To enter the number of available slots:

1. In the Sun StorEdge EBS Administrator program, select Jukeboxes from the Media menu.
2. In the Jukeboxes dialog box, select the appropriate autochanger.
3. In the Available Slots attribute, enter the range of available slots.
4. Click Apply.

---

**Note** – With two-sided media, the number of available slots is effectively doubled. For example, if you have 32 optical disks, labeled “jupiter.001.a” to “jupiter.032.b,” you have a total of 64 sides, and therefore 64 slots from which to choose.

---

# Maintaining Autochanger Devices

Maintaining a storage device in good working order requires periodic cleaning. The Sun StorEdge EBS server provides automatic cleaning of devices located in an autochanger. The server does not support automatic cleaning for stand-alone devices. Cleaning autochanger devices requires that you configure attributes in both the Jukebox and Device resources of each appropriate device. For more information on cleaning autochangers, see [“To Configure an Autochanger Resource for Cleaning” on page 394](#).

The service mode feature allows a device to be taken offline temporarily for cleaning or other maintenance. For more information about the service mode feature, see [“Service Mode” on page 393](#).

The TapeAlert feature provides diagnostic information regarding devices for which hardware cleaning is enabled. To use this feature, first enable automatic cleaning. For more information about the TapeAlert attributes, see [“TapeAlert” on page 397](#).

## Service Mode

Use the Service Mode setting to take a device offline temporarily. Service mode differs from the disabled state in that the `nsrmm` process is not stopped.

While a device is in service mode, `save` or `recover` sessions that are running or pending are completed. No more 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 manual operations using the `nsrjb` or `nsrmm` command with the `-f` option. For more information, refer to the man pages.

The device may 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 attribute. This means that if there are no hardware issues, the tape can be ejected and used in other drives. For more information on setting the maximum consecutive error count, see [“To Limit Repeated Device Errors” on page 226](#).

---

**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*

To put a device in Service Mode:

1. In the Sun StorEdge EBS Administrator program, select Devices from the Media menu.
2. In the Devices list, select the device you want to take offline.
3. Select Service for the Enabled attribute.

### ▼ *To Put a Shared Device in Service Mode*

To put a shared device in Service Mode:

1. In the Sun StorEdge EBS Administrator program, select Devices from the Media menu.
2. In the Devices dialog box, select View>Details.
3. In the Devices list, select the device you want to take offline.
4. Select Service All for the Shared Devices attribute.

## ▼ To Configure an Autochanger Resource for Cleaning

---

**Note** – Some autochangers have special hardware to enable autocleaning. If your autochanger has this feature and you wish to use the autochanger's automatic cleaning, skip this section.

---

To configure an autochanger for cleaning:

1. In the Sun StorEdge EBS Administrator program, select Jukeboxes from the Media menu.
2. In the Jukeboxes dialog box, select the appropriate autochanger.
3. In the Cleaning Slots attribute, enter the slot number of the autochanger that contains the cleaning cartridge.  
  
If more than one slot contains a cleaning cartridge, specify the range with a hyphen, for example, "8-9."
4. In the Default Cleanings attribute, enter the number of times the cleaning cartridge can be used before it needs to be replaced.
5. Select Yes for the Auto Clean attribute to enable autocleaning.
6. Click Apply.

7. Configure each Device resource, as follows:

- a. In the Sun StorEdge EBS Administrator program, select Devices from the Media menu.
- b. In the Devices dialog box, select the appropriate device.
- c. Enter a value for the Cleaning Interval attribute.  
This attribute indicates the amount of time between cleanings. The time may be specified in days, weeks, or months. Note that the last time the device was cleaned is listed in the Date Last Cleaned attribute.
- d. Click Apply.
- e. Repeat [Step a](#) through [Step d](#) for each device included in the autochanger.

To ensure there is no interference with other autochanger operations, the Sun StorEdge EBS server only cleans devices before mounting a volume in a device. For recommendations on the frequency and cleaning methods for autochanger devices, refer to the autochanger manufacturer's documentation.

[TABLE 10-5](#) lists the notification messages the Sun StorEdge EBS software displays to inform you about cleaning cartridge operations.

**TABLE 10-5** Notification Messages

| Notification Message        | Meaning                                                                      |
|-----------------------------|------------------------------------------------------------------------------|
| Device cleaning required    | Auto Clean is disabled and the device needs to be cleaned.                   |
| Device cleaned              | Auto Clean is enabled and the device has been cleaned.                       |
| Cleaning cartridge required | Auto Clean is enabled and there are no usable cleaning cartridges available. |
| Cleaning cartridge expired  | The cleaning cartridge has been used the specified number of times.          |

## ▼ To Use a Nondefault Slot for the Cleaning Cartridge

To use a nondefault slot for the cleaning cartridge:

1. Insert the cleaning cartridge in the slot you want.
2. In the Sun StorEdge EBS Administrator program, select Jukeboxes from the Media menu.
3. In the Jukeboxes dialog box, select the appropriate autochanger.
4. Select Details from the View menu to display hidden attributes.
5. Select Yes for the Auto Clean attribute.

6. In the Cleaning Slots attribute, specify the slot number to use for the cleaning cartridge.
7. The Default Cleanings attribute lists the number of uses left on the cleaning cartridge. Record this number for use in [Step 10](#).
8. Change the Available Slots attribute to reflect the range of slots available for data volumes.

If the slot for the cleaning cartridge is not the first or last slot in the autochanger, you must specify two ranges of slots in the Available Slots attribute. This is because the inventory of the autochanger must be performed in two steps, once for each range of slots.

For example, if the autochanger contains 11 slots and slot 6 is used for the cleaning slot, specify:

**1-5**

**7-11**

9. Click Apply.
10. Enter the following at the command prompt:

```
nsrjb -U number_of_uses -j autochanger_name -S slot
```

where:

- *number\_of\_uses* is the number of uses listed in the Default Cleanings attribute (the number you wrote down in [Step 7](#)).
- *autochanger\_name* is the name of the Jukebox resource.
- *slot* is the slot you now use as the cleaning slot. You can omit the -j option if there is only one autochanger.

Every time you replace the cleaning cartridge in the autochanger, you must run the nsrjb program to specify the number of uses left.



---

**Caution** – If the autochanger does not support the barcode labeling option, you must issue the command shown in [Step 10](#) to tell the autochanger to add the cleaning cartridge to its inventory.

---

## ▼ To Replace Cleaning Cartridges

When a cleaning cartridge expires, the Sun StorEdge EBS software displays the following message:

```
Date server_name: Sun StorEdge EBS Cleaning cartridge expired:
(notice) Cleaning cartridge required in Jukebox 'autochanger_name'
```

Replace the expired cartridge and enter the following command:

```
nsrjb -U number_of_cleanings
```

If you do not enter the `nsrjb -U` command along with the number of cleanings, the Sun StorEdge EBS software does not recognize the new cartridge and continues to send the "cleaning cartridge required" notification.

If the autochanger does not support barcode labeling, you must enter the following command at the prompt to tell the autochanger to add a cleaning cartridge to its inventory:

```
nsrjb -U number_of_uses -j autochanger_name -S slot
```

where:

- *number\_of\_uses* is the number of cleanings the cartridge supports.
- *autochanger\_name* is the name of the Jukebox resource.
- *slot* is the slot number where you inserted the cleaning cartridge.

## TapeAlert

When CDI is enabled in the Device resource on the server, the TapeAlert attributes, which provide tape drive status, are displayed. Devices that are capable of TapeAlert perform constant self-diagnostics and, by using `nsrmmmd`, communicate the diagnostic information through messages posted to the Device resource. The messages are also entered into the daemon log.

---

**Note** – If CDI cannot be enabled, the platform is not supported.

---

TABLE 10-6 lists the three types of TapeAlert severity messages.

TABLE 10-6 TapeAlert Severity

| Severity    | Urgently Requires User Intervention | Risks Data Loss | Explanatory |
|-------------|-------------------------------------|-----------------|-------------|
| Informative |                                     |                 | X           |
| Warning     |                                     | X               | X           |
| Critical    | X                                   | 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, such as:

- 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, such as:

- 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 may be disabled, for example:

- Unrecoverable read or write errors
- Tape marked read-only
- Drive cleaning required immediately
- Drive prediction of hardware failure

---

**Caution** – 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. If the message is not cleared, simply delete it by clicking the Delete button.

Critical messages relating to hardware errors are not cleared by `nsrmmd`; they remain in the log as reminders, as they may indicate intermittent hardware problems. These messages can also be manually deleted by clicking the Delete button.



# Auto Media Management with Autochangers

Auto media management further automates the backup procedure by freeing administrators from the task of labeling volumes before backup. When auto media management is enabled in the Jukebox resource, the Sun StorEdge EBS software assumes that all the volumes loaded in the autochanger may be handled exclusively as Sun StorEdge EBS media. Volumes that appear unlabeled are considered blank and are automatically labeled, mounted, and overwritten with new data. The Sun StorEdge EBS software only uses volumes that it considers unlabeled if it cannot locate a writable volume or a volume eligible for automatic recycling.

The Sun StorEdge EBS software does not recognize the following volume labels and considers these volumes unlabeled:

- Volumes without a label.
- Volumes labeled with something other than a Sun StorEdge EBS label.
- Volumes with Sun StorEdge EBS labels that were written in a density different from the device in which the volume is currently loaded.

For example, if you performed Sun StorEdge EBS backups on a new 8 mm tape drive, an older 8 mm tape drive might not be able to read the volumes from the newer device because it used a different density to write the data.



---

**Caution** – Be careful about sharing Sun StorEdge EBS volumes among different storage devices if you have auto media management enabled. You could potentially relabel and overwrite valuable data.

---

The auto media management feature is located in both the Jukebox and Device resources. For autochangers, you only need to enable auto media management in the Jukebox resource. The Sun StorEdge EBS software does not allow you to enable auto media management in the Device resource if the device is located in an autochanger.

If you back up to a stand-alone device and want to use the auto media management feature, see [“Storage Management Operations” on page 231](#).

When auto media management is enabled, the Sun StorEdge EBS software assumes that the autochanger is exclusive to the server, and the volumes loaded in the device are handled as Sun StorEdge EBS media. Do not enable this feature if the autochanger or device is a shared resource. The Sun StorEdge EBS server assumes it can label and overwrite any media it considers unlabeled.

However, if the autochanger is a shared device, you may use the available slot range to restrict the volumes used for backup. If you place or store other volumes in the autochanger, you do not want to use slots that are not included in the available slot range.

## How the Sun StorEdge EBS Server Selects a Volume for Writing

The Sun StorEdge EBS server uses the following process to select a volume for writing:

1. If auto media management is enabled and the Sun StorEdge EBS server is connected to an autochanger, the autochanger searches for a mounted volume for writing.
2. If no mounted volume is a candidate for writing, the Sun StorEdge EBS server looks for an unmounted, appendable volume from the appropriate pool that can be mounted.
3. If no unmounted, appendable volume is found, the Sun StorEdge EBS server looks for a recyclable volume from the required pool that can be mounted. If one is found, the Sun StorEdge EBS server mounts and relabels the volume. As part of the relabeling operation, the Sun StorEdge EBS software removes all entries for the save sets and volume are in the media database, replacing them with a new entry for the relabeled volume.

The Sun StorEdge EBS software formats the label by obtaining the next volume name from the pool's associated label template. The server automatically creates a label that incorporates the next available unused label, which is not necessarily the next-in-sequence label, depending on previous labeling activity or volume deletions. Administrators are advised to affix an adhesive label on the outside of the volume for easier identification.

4. If no recyclable volumes from the required pool are found, the Sun StorEdge EBS server looks for an unlabeled volume. If one is found, the server mounts the new volume and labels it. An entry for the newly labeled volume is added to the media database.

## Mounting Autochanger Volumes

Typically, autochangers automatically mount volumes needed for backup and recovery if they are located in the autochanger. When a volume located in an autochanger is required for backup, the Sun StorEdge EBS software uses the following order to select a volume from the appropriate pools that are not currently in use:

- Already mounted, appendable volume
- Unmounted, appendable volume
- Unmounted, recyclable volume
- Unmounted, unlabeled volume (autochanger must have auto media management enabled)

## ▼ To Mount an Autochanger Volume

To mount an autochanger volume:

1. In the Sun StorEdge EBS Administrator program, select the appropriate autochanger device in the Devices window.
2. Select Mount from the Media menu.
3. Select a volume from the Slot list on the left and click OK.

---

**Note** – You can also mount volumes by dragging the volume from the Slot list box to the appropriate device in the Device list.

---

The Sun StorEdge EBS server automatically mounts the volume without any further prompting. The Devices window reflects the device's mounted status. Place an adhesive label on the autochanger identifying its device pathnames. This reminds you which device pathnames belong to the autochanger, and can be helpful when you back up to more than one autochanger.

## ▼ To Mount an Uninventoried Volume

This procedure describes how to mount a volume that is not included in the autochanger inventory, but which is a valid (properly labeled) Sun StorEdge EBS volume.

To mount an uninventoried volume:

1. Manually load the volume in an empty autochanger slot.
2. In the Sun StorEdge EBS Administrator program, select the appropriate autochanger device in the Devices window.
3. Select Inventory from the Media menu.
4. Inventory the slot by entering the same slot number in the First Slot and Last Slot attributes. For more information on inventorying, see [“Inventorying Autochanger Volumes” on page 407](#).

The Sun StorEdge EBS software then inventories the specified slot.

5. Mount the inventoried volume as described in [“To Mount an Autochanger Volume” on page 401](#).

---

**Note** – Unlabeled media is not considered “valid” for mounting in this case because media of this type can only be mounted for labeling, not inventorying. If you try to mount uninventoried volumes using unlabeled media, an I/O error occurs and the volume is ejected.

---

## ▼ To Unmount an Autochanger Volume

To unmount an autochanger volume:

1. In the Sun StorEdge EBS Administrator program, select the appropriate autochanger device in the Devices window.
2. Select Unmount from the Media menu.
3. Select the device with the volume you want to unmount from the Device list and click Unmount.

The Sun StorEdge EBS server automatically unmounts the volume without any further prompting and returns the volume to its appropriate slot. The Devices window reflects the appropriate device’s unmounted status.

## ▼ To Label an Autochanger Volume

The Sun StorEdge EBS server uses the label template selected in the Pool resource to label the volumes loaded in the autochanger. Before starting a labeling session, select a label template from the Pool resource, or create a custom label. For more information on label templates, see [“How the Sun StorEdge EBS Server Uses Label Templates” on page 180](#).

It takes the Sun StorEdge EBS server some time to label the volumes in an autochanger. You may want to perform the labeling process at a time when you do not need to back up or recover files.

To label an autochanger volume:

1. In the Sun StorEdge EBS Administrator program, select the appropriate autochanger device in the Devices window.
2. Select Label from the Media menu.
3. Enter a label in the Starting With field.

You can enter a custom label only if you are labeling a single volume in the autochanger. If you are labeling a range of volumes, the label is automatically assigned based on the pool you select.



---

**Caution** – If you label a specific range of volumes, the name that appears for the Label attribute must match the label template. If you label a single volume, you can use any name; it does not have to match the label template. The First and Last attributes for the Slot Range in the Autochanger Label dialog box should be identical when labeling a single volume.

---

4. Complete the Slot Range text boxes:

- To label a range of slots, enter the range in the First and Last text boxes.
- To label only one volume, enter the same number in both the First and Slot text boxes.

5. From the Pool list, select a pool for the volumes being labeled.

The Sun StorEdge EBS server uses the label template associated with the selected pool to label the volumes. Only pools that have been enabled in the Pool resource appear in the Jukebox Labeling dialog box.

6. To manually recycle the volume, select the Manual Recycle check box.

---

**Note** – Use this option only if you want to keep the data beyond the retention policies associated with the data and protect the volume from automatic recycling. When you no longer want to keep the data, you must change the volume mode to either *recyclable* or *appendable*.

---

7. Click OK.

If a valid Sun StorEdge EBS label already exists on the media you are trying to label, a confirmation message appears asking if you really want to relabel the volume. The confirmation is important because it prevents you from accidentally relabeling the media. After a volume is relabeled, you cannot recover its contents under the previous label.

## Verifying the Label When Volume Is Unloaded

If a SCSI reset is issued during a backup, the volume will rewind and the Sun StorEdge EBS server may overwrite the tape label. To ensure that the server does not overwrite the tape label in this circumstance, set the Verify Label On Unload attribute in the Jukebox resource to Yes. When this attribute is set to Yes, the Sun StorEdge EBS server will verify if a tape label exists before unloading the tape. If the tape label does not exist, all save sets on the volume are marked as suspect and the tape is marked as full.

---

# Labeling Volumes with Barcodes

Labeling media with barcode labels is beneficial in two ways:

- Barcode labeling significantly speeds up volume inventory.
- Barcode labeling provides improved labeling accuracy.

Inventorying with barcode labels is fast and efficient because you do not have to mount the volumes into the device. Instead, the autochanger scans the external barcode labels with an infrared light while the volumes remain in their slots. Inventorying with barcode labels greatly reduces the time it takes to locate a volume or determine the contents of an autochanger.

Barcode labels also provide greater labeling accuracy because the labels are adhered to the media prior to being loaded and scanned in the autochanger. Once the autochanger scans 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. It is important to update and maintain an accurate account of the autochanger contents by using the `inventory` command. Otherwise, locating a volume for backup or recovery (whether or not it uses barcode labels) is impossible because the Sun StorEdge EBS server does not know which slot contains the needed volume.

A volume must have a label, but it does not require a barcode label.



---

**Caution** – The autochanger includes hardware to read the 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 autochanger documentation or hardware vendor.

---

## ▼ To Label a Volume with Barcodes

To label a volume with barcodes:

1. Adhere the barcode labels to the volumes.
2. Place the volumes with the barcode labels in the autochanger.
3. In the Sun StorEdge EBS Administrator program, select Jukeboxes from the Media menu.
4. In the Jukeboxes dialog box, select the appropriate autochanger.

5. Select Yes for the Barcode Reader attribute.
6. If you want the volume label to match the barcode label, select Yes from the Match Barcode Labels list.

Each volume has two labels: the volume label (which may or may not match the barcode label, depending upon the barcode choices) and the barcode label.

7. Click Apply and close the Jukebox window.
8. In the Sun StorEdge EBS Administrator program, select the appropriate autochanger device in the Devices window.
9. Select Label from the Media menu to open the Jukebox Labeling dialog box.

If you enabled the Match Barcode Labels attribute in the Jukebox resource, the Volume Label field in this dialog box displays the following label:

|               |
|---------------|
| Barcode label |
|---------------|

Otherwise, the Volume Label attribute displays a label template name.

After choosing barcode labels and labeling the volume, both the volume label name and barcode label for the volume appear in the Volumes window.

## Barcode Labeling Tips

The Sun StorEdge EBS server uses the volume label and barcode label to identify volumes. Both labels are recorded in the media database. However, 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. You can require that the volume label match the barcode label by making the appropriate selections for barcode labeling in the Jukebox resource.

The following are some 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. You can only change a label name if you relabel the volume 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.

- Give every volume a unique volume label. If you try to label a second volume with an identical barcode label and you have enabled the Match Barcode Labels attribute in the Jukebox resource, the Sun StorEdge EBS server displays an error message and does not allow you to label the second volume. The error message identifies the autochanger slots containing the two volumes with identical labels and the barcode label.

To correct the problem, you must either apply a different label to one of the volumes and begin the labeling process again, or you must disable the Match Barcode Labels attribute in the Jukebox resource 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. This is because you will not inventory these volumes often, if ever.
- Before using barcode labels on existing volumes, first adhere the barcode labels to them. Then, load and mount each volume individually, so the Sun StorEdge EBS server can match the barcode label with the existing volume label.
- Record the volume label on the tape.
- You may purchase a variety of barcode labels from a third-party vendor. You can choose numeric labels, alphanumeric labels, or a special combination of numbers and characters. Furthermore, you can order barcode labels that match the current volume labeling scheme.
- Use a consistent labeling scheme. If you label volumes 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 you need. Labeling instructions for barcode labels should be provided with the autochanger hardware documentation. If you have questions about barcode labels, contact the hardware manufacturer.

A consistent labeling scheme helps you better organize and track volumes. It also aids the inventory process if all the volumes use barcode labels rather than just a limited number of them.

## ▼ To Use Unmatching Volume and Barcode Labels

---

**Note** – If you choose not to match the volume label to the barcode label, attach labels to the outside of the media.

---

To use unmatching volume and barcode labels:

1. Apply the barcode labels to the volumes.
2. Place the volumes with the barcode labels in the autochanger.



3. In the Sun StorEdge EBS Administrator program, select Jukeboxes from the Media menu.
4. Select the appropriate autochanger.
5. For the Barcode Reader attribute, select Yes.
6. For the Match Barcode Labels attribute, select No.
7. Click Apply and close the Jukebox window.

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.

8. In the Sun StorEdge EBS Administrator program, select the appropriate autochanger device in the Devices window.
9. Select Inventory from the Media menu to open the Jukebox Inventory window.  
Inventorying the volumes ensures that the Sun StorEdge EBS server has the most current volume information.
10. Click OK to inventory the volumes.
11. Select Volumes from the Media menu to determine the volume label and barcode label for each volume.

Use the information in the Volumes window to match the correct volume labels to the barcode labels. This is the easiest way to make sure that you adhere the correct volume labels to the volumes with the existing barcode labels. You may want to make a list of the name correlations.

---

**Note** – If you enabled the barcode functionality and did not adhere a barcode label to the volume, an error message alerts you that a barcode label does not exist.

---

## Inventorying Autochanger Volumes

When the Sun StorEdge EBS software labels the contents of an autochanger, it registers the location of the volumes in the autochanger slots when it assigns the volume label. *This is called taking inventory.* When you inventory the volumes in the autochanger, Sun StorEdge EBS software reads the label of each volume and records its slot number. If you do not change the volumes in the autochanger after labeling them, the Sun StorEdge EBS server can access the volumes because each volume label is assigned to a specific slot.

However, if you change the contents of the autochanger without performing the labeling process, or if you move volumes into new slots, you must inform the Sun StorEdge EBS software that the autochanger now holds a different set of labeled volumes or that the volumes are in a different order. For example, if you have more than one magazine for an autochanger, you must inventory each time you remove one magazine and load another one into the autochanger.

When you label the volumes in a new magazine, you do not need to inventory them because the Sun StorEdge EBS software automatically records the slot number in which the newly labeled volume is located.

The Sun StorEdge EBS software can use barcode labels to speed up the inventory process. If the autochanger supports the use of barcode labels, you may want to use them if you have a large number of volumes or change the autochanger contents often. For more information on using barcode labels, see [“Labeling Volumes with Barcodes” on page 404](#).

## ▼ To Inventory Autochanger Volumes

To inventory autochanger volumes:

1. Load the new volumes in the autochanger.
2. In the Sun StorEdge EBS Administrator program, select the appropriate autochanger device in the Devices window.
3. Select Jukeboxes from the Media menu.
4. For the First Slot and Last Slot attributes, enter the range of the slots you want to inventory  
To inventory all the slots, use the default values.
5. In the Autochanger Inventory dialog box, enter the slots you want to inventory:
  - To inventory a single slot, enter the same slot number for the First and Last slot attributes.
  - To inventory a range of slots, enter the first and last slot for the First and Last slot attributes.
6. Click Apply.

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 may take some time to complete. After completing an inventory, the Sun StorEdge EBS Administrator program displays the results in the Volumes window.



---

**Caution** – If you inventory the tapes in an autochanger and a tape is not labeled, the Sun StorEdge EBS server returns the following error message:

*volume tape label read, I/O error*

The process by which an inventory attempts to verify the volume label for each tape in an autochanger cannot distinguish between an I/O error caused by reading an unlabeled tape and an I/O error caused by reading a bad tape. Once the volume is loaded into the drive, `nsrjb` makes a request to `nsrd` to load the volume and verify its label. The `nsrd` command allocates `nsrmmd` to perform the verify label operation. The `nsrd` command then attempts to position the volume to read the label and generates the error.

To recover from running an inventory command on an autochanger filled with new, unlabeled tapes:

1. Enter `nsrjb -HE` (hardware reset with initialize element status) at the command prompt to synchronize the state of the library.
2. Label the volumes in the autochanger.
3. Inventory the volumes in the autochanger using either the steps described in this section or the `nsrjb -I` command.

---

## Using Autochangers with a Volume Import/Export Capability

The Sun StorEdge EBS software supports the use of the SCSI-II import/export feature found in many brands of autochangers. Depending on the autochanger model, this feature is also known as cartridge access port (CAP), mail slot, and loading port. The import/export feature deposits and withdraws volumes from slots in the autochanger. 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 autochanger. However, the Sun StorEdge EBS server does not automatically inventory the volume after a deposit and withdraw.

---

**Note** – The procedures in this section do not apply to silos. For information about import/export capabilities with silos, see [“Using Silos with a Volume Import/Export Capability” on page 422](#).

---

## ▼ To Deposit a Volume by Using the Import/Export Feature

To deposit a volume by using the import/export feature:

1. Log in as root on the Sun StorEdge EBS server or storage node.
2. Press the button on the front of the autochanger to move the cartridge holder forward and open the CAP.
3. Place the volume in the holder.
4. Press the button again to move the media into the autochanger and close the CAP.
5. Enter the following command at the prompt:

```
nsrjb -d -S slot
```

where *slot* is the slot number for the volume.

If you do not specify a volume name with the `nsrjb` command, the Sun StorEdge EBS server marks the slot with a “-\*” to indicate that the volume in that slot is unknown. Inventory the slot with the following command:

```
nsrjb -I -S slot
```

6. To verify that the volume was deposited in the correct slot, mount the volume.

## ▼ To Withdraw a Volume by Using the Import/Export Feature

To withdraw a volume by using the import/export feature:

1. Log in as root on the Sun StorEdge EBS server or storage node computer.

2. Enter the following command at the prompt:

```
nsrjb -w -S slot -P port
```

where *slot* is the slot number for the volume and *port* is the value assigned to the control port.

3. When the volume reaches the CAP, press the button to open the CAP.
4. Remove the volume and close the CAP.

---

## Tips for Using Autochangers

This section provides additional suggestions for using autochanger and backup media effectively and reliably. It includes the following sections:

- [“Autochanger Notifications” on page 411](#)
- [“Recycling Compared to Adding More Volumes” on page 412](#)
- [“Using Pools with Autochangers” on page 412](#)
- [“Device Calibration” on page 413](#)
- [“To Disable the Element Status” on page 413](#)
- [“Resetting an Autochanger” on page 413](#)
- [“Optical Autochangers” on page 414](#)

## Autochanger Notifications

The Sun StorEdge EBS server uses notifications to send messages about Sun StorEdge EBS events. A number of preconfigured notifications are included to inform you of the following situations:

- Volumes in the autochanger are 90% full
- Autochanger needs more volumes to continue
- Autochanger has a mechanical problem
- Autochanger device needs cleaning
- Cleaning cartridge needs attention

For more information on notifications, see [“Notifications” on page 65](#).

The Sun StorEdge EBS software automatically mounts a required volume as long as the volume is loaded in the autochanger. If the recovery requires a volume that is not loaded in the autochanger, the *Tape mount request 1* notification sends a message to the Messages window in the Sun StorEdge EBS Administrator program, informing you to mount a specific volume.

After you correct an autochanger problem, you might need to mount a volume so the Sun StorEdge EBS server can continue to back up or recover files. Check the Sun StorEdge EBS Administrator program for messages.

## Recycling Compared to Adding More Volumes

The Sun StorEdge EBS server saves files on volumes marked *appen* (appendable). If the volumes inside the autochanger are marked *full*, they cannot receive backups.

You can do one of the following tasks with volumes marked *full*:

- Remove the full volumes and replace them with new media if you need to keep the volumes for long-term storage.
- Change the volume mode to *recyc* (recyclable) if you do not need the data on the full volumes. The Sun StorEdge EBS server overwrites the data with new backups, but maintains the existing labels. This is the only instance in which you do not need to relabel a volume to make it eligible to be overwritten with new data.

When all of the save sets on the volume have passed the time period specified by the retention policy, the mode of a volume automatically changes to recyclable.

There are advantages to both recycling media and adding more media to a pool. By recycling, you reuse the same volumes and do not add new volumes to the pool. However, the media may wear out over time and have a higher failure rate.

On the other hand, if you need to store backups for a period of time, then you may have to add more media to the pool instead of recycling it. For example, an autochanger may need new volumes every three months if the company policy is to maintain the backups for a year. In this case, you need to keep adding new media to the pool until you can recycle the volumes that contain expired or old backups.

## Using Pools with Autochangers

If you plan to have 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 autochanger. This allows you to remove all of the volumes at the same time.

## Device Calibration

For information about the frequency and method for calibrating the loading mechanism for the device, refer to the autochanger manufacturer's documentation.

### ▼ To Disable the Element Status

Some autochangers support the element status feature. The Sun StorEdge EBS software uses the element status feature to determine whether a slot contains a volume. If volumes have barcodes, the element status feature speeds up the inventory process by eliminating the need to load and read volume labels. If the volume has a barcode label, the Sun StorEdge EBS software reads it without loading the volume.

To disable the element status option for an autochanger if the volumes do not have barcodes:

1. In the Sun StorEdge EBS Administrator program, select Jukeboxes from the Media menu.
2. Select Details from the View menu to display hidden attributes.
3. Select the appropriate autochanger.
4. For the Jukebox Options attribute, clear the Elements Status attribute.
5. Click Apply.

## Resetting an Autochanger

You might need to reset an autochanger if it is unable to complete an operation, or if it appears to be in a "confused" state. For example, the Sun StorEdge EBS server acts as if a slot is empty when a volume is actually loaded in the drive. The opposite might be also be true: the autochanger acts as if the slot is full, but the slot is actually empty. Use of the reset feature can return an autochanger with these characteristics back to normal functionality.

## ▼ To Reset an Autochanger

To reset an autochanger:

1. Enter the following command at the prompt:

```
#nsrjb
```

Make a note of the output. This command displays information about the volumes loaded in the devices, the volume labels, and the slots from which the volumes were loaded into the devices.

2. Manually unload the volume from the device. For information on how to do this, refer to the autochanger documentation.
3. Manually reinsert the volume into the CAP, flipping it over so that the “A” side faces up.
4. Reload the volume from the CAP into the slot from which it originally came, using the control panel on the autochanger.
5. Enter the following command at the prompt:

```
#nsrjb -H
```

The autochanger resets itself.

## Optical Autochangers

If an optical autochanger does not work with the Sun StorEdge EBS software, the operating system might not currently include support for the optical autochanger media device.

When you create the Sun StorEdge EBS device names for an optical autochanger, you should use the raw name of the device.

If you experience a power outage, the optical autochanger loses track of the state of its backup volumes. Issuing the `nsrjb -H` command might not successfully reset the autochanger. If you see output similar to the following, it means the `nsrjb -H` command failed to reset the autochanger:

```
nsrjb: Autochanger error, Invert operation not supported
```



---

# Silos

The Sun StorEdge EBS software manages silos and silo media similarly to autochangers. For best results, familiarize yourself with autochanger 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. Like autochangers, silos make data and media operations more automatic. Silos can load, change, and manage volumes, and clean the devices automatically.

## 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 an autochanger. For example, the silo management software keeps track of the slot where each silo volume resides, and may control the deposit and withdrawal of volumes, and automated cleaning of silo devices.

### ▼ 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), following the instructions provided by the silo vendor.

Do not install the STLI library on the following models, because all the necessary software is installed when you install the Sun StorEdge EBS software:

- IBM 3494 on Solaris and AIX
- StorageTek on Solaris, AIX, HP-UX, and DYNIX/ptx
- ADIC EMASS/Grau on Solaris, AIX, HP-UX, and Windows NT

3. Ensure that the Sun StorEdge EBS server is properly connected to the media devices in the silo.

4. Run the `jbconfig` program to configure the devices in the silo.

The `jbconfig` program prompts you to enter the following information:

- Device type to install. Enter **3** (for STLI Silo).
- Type of silo.
- Name of the silo. Enter any alphanumeric string.
- Hostname of the silo server.
- Whether to enable automatic device cleaning. Enter No. The silo server controls device cleaning in silos.
- Pathname of the STLI software. Enter the full path on the Sun StorEdge EBS server where the STLI software resides.
- Number of devices to configure.
- For each device configured, the program prompts you for the following information:
  - Pathname of the device
  - Media type of the device (if the device is not already configured)
  - Silo name of the device

For additional information, see [“To Add an Autochanger” on page 361](#).

5. License the Silo Support Module. For general information on licensing, the Sun StorEdge EBS Roadmap.

# Silo Device Naming Conventions

The `jbconfig` program prompts you for the silo name of the storage devices. The silo name is the name that the silo management software uses to refer to the storage device. Depending on the type of silo you have, 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 looks like 1,0,1,0.

You must ask the silo administrator for the drive names of the devices that the Sun StorEdge EBS server can use. There is no method to find 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 you accidentally swap operating system device names and silo names, you can only mount and unmount volumes; you cannot read or write to the volumes after they are mounted. To reconfigure the device names properly, use the Sun StorEdge EBS Administrator program to change the order of the device names in the STL Device Names attribute of the Jukebox resource.

## 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 `jbconfig` program obtains the name of the device from the device driver and displays the device name as the default value.

- On a Solaris system, you must use the IBM-supplied `mtlib` command (`mtlib -l library_name -D`) to determine the names of all the devices in the 3494. 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.

## EMASS/Grau

The silo management software for the EMASS (in North America) or Grau 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.

### ▼ To Set up DAS to Work with Sun StorEdge EBS Software

To set up DAS to work with Sun StorEdge EBS software:

1. Ask the silo administrator to configure DAS to accept commands from the Sun StorEdge EBS server or storage node computer.
2. Ask the silo administrator to 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 your Sun StorEdge EBS server or storage node as an administrator, so you can enter the `dasadmin allocd` command to allocate devices from the Sun StorEdge EBS server or storage node computer.

### ▼ *To Find the Assigned Names for 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 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
```

## ▼ To Define a Silo Device

If you add a new device to the silo and you want the Sun StorEdge EBS server to recognize and use the device, you must first delete the existing Jukebox resource, and then reconfigure the silo's Jukebox resource to access the new device.

To define a silo device:

1. Log on as root to the Sun StorEdge EBS server that uses the silo.
2. In the Sun StorEdge EBS Administrator program, select Jukeboxes from the Media menu.
3. Select the appropriate silo.
4. Click Delete.

Click OK to confirm the deletion.

5. Run the `jbconfig` program to reconfigure the silo. For more information about `jbconfig`, see [“To Add an Autochanger” on page 361](#). Use exactly the same name for the silo as before, and add the additional device.

If you define multiple devices in a silo for the Sun StorEdge EBS software to use, you do not need to specify them in any particular order in the Jukebox resource. The only restriction is that the list of device names in the Devices attribute must be in the correct order with respect to the list in the STL Device Names attribute.

## ▼ To Release a Silo Device

When a silo device is configured for use with a Sun StorEdge EBS server, you have the option of restricting silo access only to the Sun StorEdge EBS server. These restrictions allow increased availability to the silo for those with full access. However, these restrictions can be lifted by using the Release Device feature.

To release a silo device:

1. In the Sun StorEdge EBS Administrator program, select the appropriate silo in the Devices window.

2. Select Release Shared Device from the Media menu.
3. Repeat these steps for each device you want to release.

## Silo Device Cleaning

Do not enable automated cleaning for the silo 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 on how to clean devices in the silo, refer to the silo management software documentation.

### ▼ To Set Environment Variables for EMASS/Grau and StorageTek Silos

For EMASS/Grau and StorageTek silos, you need to set environment variables in the Sun StorEdge EBS startup script. [TABLE 10-7](#) lists the environment variables to set:

**TABLE 10-7** EMASS/Grau and StorageTek Silo Environment Variables

| Silo Model | Environment Variables                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| EMASS/Grau | <p> DAS_SERVER = <i>name_of_DAS_server</i><br/> DAS_CLIENT = <i>name_of_system_as_defined_to_DAS_server</i><br/> DAS_MEDIUM = <i>type_of_tape_drive_used</i><br/> ACI_MEDIA_TYPE = <i>type_of_tape_drive_used</i> </p> <p>For DAS_MEDIUM and ACI_MEDIA_TYPE, use one of the following values:</p> <ul style="list-style-type: none"> <li>• 3480</li> <li>• OD-Thick</li> <li>• OD-Thin</li> <li>• DECDLT</li> <li>• 8MM</li> <li>• 4MM</li> <li>• D2</li> <li>• VHS</li> <li>• 3590</li> <li>• CD</li> <li>• TRAVAN</li> <li>• DTF</li> <li>• BETACAM</li> <li>• AUDIOTAPE</li> </ul> <p>To specify the cartridge access port or range of ports for depositing or withdrawing volumes, use the following environment variables rather than the -P <i>ports</i> option with nsrjb:</p> <ul style="list-style-type: none"> <li>• DAS_INSERT_AREA_NAME=<i>port</i></li> <li>• DAS_EJECT_AREA_NAME=<i>port</i></li> </ul> <p>The default values for DAS_INSERT_AREA_NAME is I01, and the default value for DAS_EJECT_AREA_NAME is E01.</p> |

**TABLE 10-7** EMASS/Grau and StorageTek Silo Environment Variables

| Silo Model | Environment Variables                                                                                                                                                                                                                                                                |
|------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| StorageTek | <p>CSI_HOSTNAME = <i>name_of_ACSLS_system</i></p> <p>Additionally, the following commands should be running on the system and can be included in the Sun StorEdge EBS startup script:</p> <pre>&lt;ebs_binaries_path&gt;/mini_el &amp;<br/>&lt;ebs_binaries_path&gt;/ssi &amp;</pre> |

To set environment variables for EMASS/Grau and StorageTek silos:

1. Using a text editor, open the Sun StorEdge EBS startup script (/etc/init.d/networker) and add the environment variables after the lines:

```
(echo 'starting Sun StorEdge EBS daemons:') > /dev/console
```

2. Enter the environment variables using the format:

```
ENV_VAR = value
export ENV_VAR
```

**Note** – You must restart the Sun StorEdge EBS server for the environment variables to take effect.



# Media Management in a Silo

More than one software application can use a single silo. Therefore, media management in a silo requires some extra operations to prevent the Sun StorEdge EBS software from overwriting volumes other programs use.

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 an autochanger, the Sun StorEdge EBS software specifies many of the functions by a slot number. A silo uses the same idea. An autochanger has a fixed number of slots, and the Sun StorEdge EBS software uses the slot number to refer to a volume's physical location.

A silo, however, 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. You can 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 you allocate additional barcodes that fall earlier in the sort sequence, the slot numbers change for all the volumes later in the sequence. For more information, refer to the `nsrjb` man page.

Because the slot number is not a perfect identifier for a silo volume, operations that might change the slot number cannot accept slot numbers as arguments. For example, you cannot deallocate volumes based on slot numbers, because this operation can change the slot numbers of volumes being deallocated.

## Using Silos with a Volume Import/Export Capability

The Sun StorEdge EBS software supports the use of the SCSI-II 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 long operation of reinitializing the silo. Note, however, that Sun StorEdge EBS software does not automatically inventory the volume after a deposit.

You can use Sun StorEdge EBS commands or the silo management software to control the import/export feature on the currently supported silos to deposit and withdraw volumes in a silo. However, it is often more efficient to use the silo management software, especially to deposit or withdraw a large number of volumes.



Unlike autochangers, you cannot issue deposit and withdraw commands using the Sun StorEdge EBS Administrator program. To issue deposit and withdraw commands, enter the following Sun StorEdge EBS commands at the prompt:

- To allocate and deposit volumes:

```
nsrjb -a -T tags -d
```

- To deallocate and withdraw volumes:

```
nsrjb -x -T tags -w
```

where *tags* specifies the tags or barcodes of volumes in a remote silo.

On some silos (for example, StorageTek silos with the import/export feature set to automatic mode), the silo management software inserts volumes automatically. You cannot use the Sun StorEdge EBS software to insert volumes.

## Mounting and Unmounting a Silo Volume

Mounting and unmounting silo volumes is the same as for autochangers:

- You must mount a volume before you can label it, or read or write data on it. The robotic mechanism mounts volumes in the devices of a silo.
- You must unmount volumes before you can inventory the volumes in a silo or deallocate a volume from a Sun StorEdge EBS pool.

For details on mounting and unmounting volumes, see [“To Mount an Autochanger Volume” on page 401](#) and [“To Unmount an Autochanger Volume” on page 402](#).

To specify a barcode identifier or template for the volumes, use the `-T` option with the `nsrjb` command. For more information, refer to the `nsrjb` man page.

## 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 [“To Label an Autochanger Volume” on page 402](#).

The use of barcodes, and matching barcode labels and Sun StorEdge EBS volume labels, are both enabled for a silo. You can change the Match barcode Labels attribute, but do not set the Barcode Reader attribute to No. When both attributes are set to Yes, 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, you can use unlabeled volumes. The silo management software assigns a “virtual” barcode label to those volumes. Although you can use volumes without barcodes, it is very difficult to maintain integrity, since after you remove the volume from the silo, the information about the virtual barcode is lost. Any volume without a 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.

## Allocating a Silo Volume

When you allocate volumes, you direct the Sun StorEdge EBS server to those volumes it can use.



---

**Caution** – 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 problem, most silo management software includes methods to limit access to volumes based on the hostname of the computer on which Sun StorEdge EBS software and the other programs run. The Sun StorEdge EBS software does not provide any method for setting up this sort of protection; the silo management software must configure it.

---

Allocating a volume causes the Sun StorEdge EBS software to query the silo management software to verify that the volume you requested exists.

- If the volume exists, the volume is allocated to the Sun StorEdge EBS software.
- If the volume does not exist, the following message appears:

barcode *label* is not present in the silo and was not added

If you are allocating a range of volumes, the allocation continues after displaying the message. The message is informational and does not indicate a Sun StorEdge EBS error.

## ▼ To Allocate a Silo Volume

To allocate a silo volume:

1. In the Sun StorEdge EBS Administrator program, select the appropriate silo in the Devices window.
2. Select Add Library Volumes from the Media menu.
3. In the Add Library Volumes dialog box, select either Template or List for the method for entering volume names.
  - To use a template-oriented naming scheme, select Template. For example, to name four tapes A01B, A02B, A03B, and A04B, you would enter A0, 1-4, and B.
  - To enter the names for each volume separately, select List. For example, you would enter the name for each tape: A01B, A02B, A03B, and A04B.
4. Enter the appropriate name or naming scheme in the Barcodes field.
5. Click OK.

To deposit volumes into a silo and then allocate them (on silos that require manual depositing, like EMASS/Grau), place the volumes in the insert area, then enter the following command:

```
nsrjb -a -T tags -d
```

where *tags* specifies the tags or barcodes of volumes in a remote silo.

On StorageTek silos, the silo management software deposits volumes automatically.

For more information on STLIs, see [“How Sun StorEdge EBS Software Interacts with a Silo” on page 415](#).

## Deallocating a Silo Volume

If you no longer need an STL volume in a silo for the Sun StorEdge EBS software, you can deallocate the volume from the library. Deallocation is basically the same operation as removing a volume from an autochanger. Although the volume cannot be loaded by the robotic mechanism, the entries in the Sun StorEdge EBS media database remain intact. If you allocate the volume again later, Sun StorEdge EBS software can retrieve the data from it.

Use deallocation when the silo license limits the number of slots you can use, or when you move data offsite for safer storage. If the license limits the number of slots, you might be able to leave the volumes in the silo. This way, you can easily reallocate the volumes when you need to access the data on them.

The allocation operation is not automatic. You must manually allocate the volumes again and reinventory them to let the Sun StorEdge EBS server access the data. If you remove the volume from the silo for offsite storage, you must deallocate it with the Sun StorEdge EBS software and then use the silo management software to eject it from the silo.

## ▼ To Deallocate a Silo Volume

To deallocate a silo volume:

1. Unmount the volume from the device. For instructions on unmounting volumes, see [“To Unmount an Autochanger Volume” on page 402](#).
2. In the Sun StorEdge EBS Administrator program, select the appropriate silo in the Devices window.
3. Select Remove Library Volumes from the Media menu.
4. In the Remove Library Volumes dialog box, edit the values shown in the Barcodes list to add or exclude volumes from the deallocation process.
  - To use a template-oriented naming scheme, select Template. For example, to name four tapes A01B, A02B, A03B, and A04B, you would enter A0, 1-4, and B.
  - To enter the names for each volume separately, select List. For example, you would enter the name for each tape: A01B, A02B, A03B, and A04B.
5. Enter the appropriate name or naming scheme in the Barcodes field.
6. Click OK.

For more information on STLs, see [“How Sun StorEdge EBS Software Interacts with a Silo” on page 415](#).

## 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 an autochanger. The tasks for inventorying volumes in a silo are the same as for an autochanger. For information on inventorying an autochanger, see [“Inventorying Autochanger Volumes” on page 407](#).

The slot number of a silo volume is the number of the volume’s position in the list of volumes in a silo.

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 but 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. However, if 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 for a proper inventory.



# Sun StorEdge Enterprise Backup Software Cluster Support

---

This chapter discusses Sun StorEdge Enterprise Backup Software support for backing up and recovering data that resides on a cluster.

The main sections in this chapter are:

- [“Cluster Terminology” on page 429](#)
- [“Clustering” on page 430](#)
- [“Sun Cluster Environment” on page 432](#)
- [“Backing Up Data in a Cluster Environment” on page 432](#)
- [“Recovering Data in a Cluster Environment” on page 435](#)
- [“Tracking Scheduled Saves” on page 437](#)
- [“Using Autochangers and Stand-Alone Tape Devices” on page 440](#)
- [“Sun StorEdge EBS Cluster Licensing” on page 454](#)

---

## 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.

---

## Clustering

*Clustering* means that two or more nodes are connected and appear to network users as a single, highly available system. A highly available system allows the application services to continue despite most hardware or software failures. *Each cluster node has its own IP address.* Each cluster node might also have *private* (local) resources or disks that are available only to that computer. In addition to the nodes, a cluster can have one or more *virtual services*, each with its own IP address. A *virtual service* can have zero or more shared disks.

FIGURE 11-1 on page 431 illustrates a general cluster configuration consisting of two or more nodes and at least one Sun StorEdge EBS server. In this illustration:

- Node A is a computer with its own **Local Disk 1**
- Node B is a computer with its own **Local Disk 2**
- Highly available Sun StorEdge EBS Virtual Server C owns **Shared Disks 3**






**FIGURE 11-1** Basic Cluster Configuration

In this example, the Sun StorEdge EBS server can fail over between Node A and Node B; however, the Sun StorEdge EBS server only runs on one node at a time.

The Sun StorEdge EBS server carries various data services and inherits the virtual server's IP address. Although the Sun StorEdge EBS software is installed on all the physical nodes within the cluster, it only runs on one node at a time. If a hardware or software breakdown occurs on the node on which the Sun StorEdge EBS server is running, the Sun StorEdge EBS server moves to another node in the cluster.

## 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 *Legato 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.

---

## Sun Cluster Environment

The Sun StorEdge EBS software can be used for backups within a Sun Cluster environment. By configuring the Sun StorEdge EBS server software as a resource group within a Sun Cluster environment, it becomes a highly available application. For the complete configuration procedure, refer to the *Sun StorEdge Enterprise Backup Software Installation Guide*.

A Sun StorEdge EBS server installed within a cluster can migrate or fail over among other nodes in the same cluster. Failover allows another node in the cluster to take over the operations from the first node. When failover occurs during a Sun StorEdge EBS scheduled backup, the interrupted save operation restarts when the failover has completed, beginning with the interrupted save set.

Refer to the *Sun StorEdge Enterprise Backup Software Installation Guide*, for details on how to:

- Install Sun StorEdge EBS software in a cluster environment.
- Migrate an existing Sun StorEdge EBS server running on Solaris to a Solaris Sun Cluster server.

---

## Backing Up Data in a Cluster Environment

This section describes how to configure scheduled saves or data from a physical cluster client and a virtual cluster client, as well as how to perform a manual backup of data from a virtual cluster client. For more information, see the following topics:

- [“To Configure a Scheduled Save from a Physical Cluster Client” on page 433](#)

- “To Configure a Scheduled Save from a Virtual Cluster Client” on page 433
- “Performing a Manual Backup from a Virtual Cluster Client” on page 435

## ▼ 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 3](#).

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 Enterprise Backup Software Installation Guide*.

## ▼ 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:
  - a. Open the Sun StorEdge EBS Administrator program.
  - b. Select Groups from the Customize menu.
  - c. Select Details from the View menu to display hidden attributes.
  - d. Click Create to create a new save group.
  - e. Define the remaining attributes and click Apply.



---

**Caution** – 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.

---

- f. For each save group, click Enabled for the Autorestart attribute.
  - g. For each save group, clear the Manual Restart option in the Options check box.
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. In the Sun StorEdge EBS Administrator program, select Client Setup from the Client menu.
  - b. Click Create and enter the name of the physical client for the Name attribute.
  - c. Define the remaining attributes in the Client window and click Apply.
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. In the Sun StorEdge EBS Administrator program, select Client Setup from the Client menu.
- b. Click Create and enter the name of the virtual client for the Name attribute.
- 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.
- e. Define the remaining attributes in the Client window and click Apply.

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 standalone server on which the Sun StorEdge EBS server resides, enter the following command:

```
savegrp -pv -c virtual_client group_name
```



---

**Caution** – If the test probe does not display the appropriate scheduled backups and index, see [“Tracking Scheduled Saves” on page 437](#).

---

## Performing a Manual Backup from a Virtual Cluster Client

Backing up data from a virtual cluster client is similar to backing up data from noncluster clients. For example:

```
save saveset
```

or

```
save -c client saveset
```

## 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 `/nsr/res/servers` file. Only the specific servers listed in this file can back up the client. For more information about the `servers` file, see [“To Edit the servers File” on page 154](#).

---

## Recovering Data in a 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:

- [“To Recover Data from a Private Disk Belonging to a Physical Client” on page 435](#)
- [“To Recover Data from a Shared Disk Belonging to a Virtual Client” on page 436](#)

### ▼ 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 4](#).

## ▼ 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. In the Sun StorEdge EBS Administrator program, select Client Setup from the Client menu.
  - b. 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*

For details on how to configure scheduled backups, see [“To Configure a Scheduled Save from a Virtual Cluster Client” on page 433](#).

2. Make sure that the filesystems of the virtual client are mounted.
3. Recover data from a shared disk belonging to a virtual client. You can use either the `nwrecover` program or the command prompt.
  - To recover the data using the `nwrecover` program, make sure that you select the virtual client.
  - To perform a manual 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 further information regarding the `recover` command, refer to the `recover` man page.

---

# 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 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.



---

**Caution** – 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 438](#).

---

To override scheduled save rules (not path-ownership 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.

## ▼ To Force 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. In the Sun StorEdge EBS Administrator program, select Client Setup from the Clients menu.
2. Select the client that is using the incorrect index.

3. 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 137](#).

## Path-Ownership Rules

In a cluster environment, the Sun StorEdge EBS software must distinguish between filesystems associated with a physical client and filesystems 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 filesystem) 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 filesystem to be omitted and *not backed up* during a scheduled save:

- If a filesystem owned by a virtual client is defined in the Save Set attribute for a physical client resource.
- If a filesystem owned by a physical client is defined in the Save Set attribute for a virtual client resource.

### ▼ 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 filesystems *client\_name* owns. This procedure is part of the normal cluster installation setup. For detailed instructions, refer to the *Sun StorEdge Enterprise Backup Software 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 `/usr/sbin/nsr/` directory. This file causes the Sun StorEdge EBS software to back up the filesystem in question; however, the filesystem is saved under the client file index of its correct owner.

If the scheduled save uses the wrong index, see [“To Force Save Sets to Use a Specific Index” on page 437](#).



---

**Caution** – Creating the `pathownerignore` file is not recommended, but it might be required under special circumstances.

The `pathownerignore` file does not override the default path-ownership rules. It causes the path-ownership rules to be ignored when determining if a filesystem should be backed up during a scheduled save.

---

## 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.

### ▼ *To Override Path-Ownership Rules during a Manual Backup*

To override the default path-ownership rules for a manual backup, enter the following at the command prompt:

```
% save -c client_name
```

### ▼ *To Override Path-Ownership Rules during a Scheduled Backup*

To override the path-ownership rules during a scheduled backup:

1. In the Sun StorEdge EBS Administrator program, select Client Setup from the Clients menu.
2. Select the client that is using the incorrect index.

3. 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 137](#).

---

## 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 440](#).
- Shared autochanger devices, but *not* shared stand-alone tape devices, use [“Configuration Option 2” on page 443](#).
- Nonshared devices, or any configuration, use [“Configuration Option 3” on page 444](#), 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 11-2 on page 441](#) provides a graphical view of this configuration option.

---

**Note** – Only one autochanger is required, although you can add more.

---




**FIGURE 11-2** 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.



---

**Caution** – 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 script (`/usr/sbin/networker.start`) by adding any required node device-reassignment reservation commands (commands for the bridge) before the line that starts the `nsrd` daemon.



---

**Caution** – If you modify the `networker.start` script is done, it is no longer a Sun Cluster certified failover solution.

---

4. Test that the device-sharing infrastructure supports path isolation. 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. Use the Sun StorEdge EBS `jbconfig` program to configure the autochanger and tape devices. For a detailed description of the command, refer to the `jbconfig` man page.

## 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 manual intervention might be required during the data recovery process. [FIGURE 11-3 on page 443](#) 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 11-3** 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. Use the Sun StorEdge EBS `jbconfig` program to 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 within `jbconfig`, assign ownership to the node exclusively connected to that tape device. For a detailed description of the `jbconfig` command, refer to the `jbconfig` man page.

---

## 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 manual intervention might be required during the data recovery process. [FIGURE 11-4 on page 445](#) 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 11-4** 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
```

For a detailed description of the `jbconfig` command, refer to the `jbconfig` man page.

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 447](#)
- [“Controlling the Robotics from a Stand-Alone Host” on page 452.](#)

For information explaining how to activate the DDS feature, see [“Dynamic Drive Sharing” on page 198.](#)

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 *Sun StorEdge Enterprise Backup Software 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 11-5 on page 448](#) illustrates how DDS can be enabled within a cluster.




**FIGURE 11-5** 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 450.

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; consequently you should 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 11-6 on page 453](#) 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 11-6** DDS Enabled with the Robotics Controlled Outside the Cluster

## 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 11-7 on page 454](#). The virtual Sun StorEdge EBS server of cluster B controls the robotics for the physical storage nodes in cluster A.




**FIGURE 11-7** 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 455](#)
- [“Enabling Cluster and Virtual Clients” on page 455](#)
- [“Licensing Requirements for Backing Up Cluster Clients” on page 456](#)

For installation requirements, refer to the *Sun StorEdge Enterprise Backup Software Installation Guide*.



# Enabling the Highly Available Sun StorEdge EBS Server

Running the Sun StorEdge EBS server as a highly available application on a cluster requires one of the following modes or editions of the Sun StorEdge EBS software:

- Evaluation Mode — Unlimited cluster client support during the evaluation period
- Power Edition — Includes two cluster client licenses of the same platform type; after that, additional Cluster Client Connection licenses are required

---

**Note** – Network Edition does *not* support backup and recovery for a 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.

---

**Note** – Workgroup Edition does not allow backing up cluster clients. Business Edition only allows backing up two-node cluster clients.

---

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 444](#).

## 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 *instead of just* 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.

---

**Note** – Power Edition is only required for a highly available Sun StorEdge EBS server. A stand-alone Sun StorEdge EBS server can back up a physical or virtual client of a cluster with a Network Edition license and the appropriate number of Cluster Client Connection licenses.

---

## 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.

# NDMP Support for the Sun StorEdge EBS Server

---

This chapter provides instructions on how to enable and use the Sun StorEdge Enterprise Backup Software NDMP Connection with a Sun StorEdge EBS server.

This chapter includes the following sections:

- [“Overview of NDMP” on page 457](#)
- [“Supported NDMP Configurations” on page 458](#)
- [“General Requirements for and Limitations of NDMP Support” on page 463](#)
- [“Configuring the Sun StorEdge EBS Server for NDMP Operations” on page 466](#)
- [“Performing NDMP Backups” on page 477](#)
- [“NDMP Cloning” on page 481](#)
- [“Performing NDMP Recoveries” on page 482](#)
- [“Vendor-Specific Information” on page 489](#)

---

## Overview of NDMP

Network Data Management Protocol (NDMP) is a standards-based, storage management client/server protocol that allows the Sun StorEdge EBS server to perform NDMP backups, cloning, and recoveries.

# NDMP Connection Support

Sun StorEdge EBS backup and recovery software is NDMP-compliant. The NDMP interface enables the Sun StorEdge EBS server to communicate with other hosts on the network that support NDMP. With the NDMP interface, you can connect to hosts that have an active NDMP service or an NDMP data module installed. The hosts, except for the Sun StorEdge EBS server, do not need to have the Sun StorEdge EBS software installed.

To install an NDMP data module or enable NDMP support on an application server, refer to the respective vendor's documentation.

Sun StorEdge EBS software, release 7.0 and later, supports up to NDMP version 4.

## NDMP Connection Software and Installation Requirements

Before you enable the NDMP Connection, you must install the Sun StorEdge EBS server software on a computer designated for this purpose, and you must be using Sun StorEdge EBS Network Edition or Sun StorEdge EBS Power Edition. For more information on installing the Sun StorEdge EBS software, refer to the *Sun StorEdge Enterprise Backup Software Installation Guide*.

## 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. For details about the NDMP tiers, contact Legato Sales at [sales@legato.com](mailto:sales@legato.com).

Additionally, while a file type device is not supported for NDMP operations, the device can be used as a bootstrap device. A DiskBackup license is required for the file type device.

---

## Supported NDMP Configurations

The workflow for NDMP backup and recovery operations is similar to the workflow for traditional Sun StorEdge EBS backup and recovery operations.

The Sun StorEdge EBS server with NDMP Connection supports the following configurations:

- **Local** — The backup is performed to a tape device attached to the same server where the data resides.
- **Three-party** — The backup is performed from one NDMP server to another NDMP server.

## Local Backup

The local backup involves a tape device attached directly to the server where the data resides. The data being backed up is sent directly from the NDMP server to its attached autochanger. Configure the Sun StorEdge EBS server for NDMP operations as you would for any backup.

FIGURE 12-1 on page 459 illustrates a local configuration.




FIGURE 12-1 NDMP Local Configuration

## Three-Party Backup

A three-party backup, also known as a three-way backup, is completed using two NDMP servers. The data flows from the NDMP host to the NDMP tape server and then to the autochanger that is attached to the tape server.



**Caution** – 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—and restore to—another NetApp server.

FIGURE 12-2 on page 460 demonstrates a three-party configuration.




**FIGURE 12-2** NDMP Three-Party Configuration

## Remote Backup

In a remote backup, the data travels from an NDMP host to an NDMP tape server. By using a tape server, any NDMP host can be backed up to the same tape server, thus reducing the need for additional devices and libraries.

The Sun StorEdge EBS software with NDMP supports the following tape servers:

- Legato NetWorker SnapImage Module 1.6 software installed on the Sun StorEdge EBS server or another server (software)
- DinoStor TapeServer (hardware)

## Remote Backup with SnapImage Software

In a remote backup, the Legato NetWorker SnapImage Module 1.6 software acts as an NDMP tape server. This enables the backup of data from an NDMP host to a Sun StorEdge EBS server or storage node with the SnapImage Module software installed and tape drives attached. The SnapImage Module software does not actually perform the backup. The backup is controlled by the native NDMP utilities provided by the NAS box. This is not a block-level backup using the SnapImage Module software. The Sun StorEdge EBS server or storage node with the SnapImage Module software installed functions as a tape server, accepting the NDMP data stream and writing it to tape.

---

**Note** – Backing up data *from* the server or storage node with the SnapImage Module software *to* a tape drive attached to a NAS server is not currently supported.

---

FIGURE 12-3 on page 461 shows a possible configuration for backing up data on a NAS server to a Sun StorEdge EBS server. In this configuration, the data flows from the NAS server, over the LAN, and into the tape library directly attached to the Sun StorEdge EBS server. The SnapImage Module software is installed on the Sun StorEdge EBS server and functions as the NDMP tape server.




**FIGURE 12-3** Performing a Remote NAS NDMP Backup to a Sun StorEdge EBS Server

FIGURE 12-4 on page 462 shows a configuration where the data is backed up to a server or storage node, other than a Sun StorEdge EBS server, with the SnapImage Module software installed. The data flows from the NAS server, over the LAN, and into the tape library directly attached to the server or storage node. The backup is controlled by the native NDMP utilities provided by the NAS box.




**FIGURE 12-4** Performing an NDMP Backup to a SnapImage Module Server

For instructions on performing a remote backup with the SnapImage software, see [“To Perform a Remote Backup with the SnapImage Software” on page 509](#).

## Remote Backups with the DinoStor TapeServer

DinoStor manufactures the DinoStor TapeServer, which connects one or more libraries to the network. This allows any NDMP host to be backed up to one location instead of having a local backup device for each server.

[FIGURE 12-5 on page 462](#) demonstrates how the Sun StorEdge EBS server interacts with the DinoStor TapeServer.




**FIGURE 12-5** Remote Backups with the DinoStor TapeServer

For instructions on performing a remote backup with the DinoStor TapeServer, see [“To Perform a Remote Backup with a DinoStor TapeServer” on page 507](#).



# General Requirements for and Limitations of NDMP Support

The following sections contain information about general requirements and limitations of the Sun StorEdge EBS software for NDMP support. For information regarding the limitations, autochanger configuration, and application information variables of each vendor, see [“Vendor-Specific Information” on page 489](#). TABLE 12-1 lists the specific vendors and location of this information.

**TABLE 12-1** Vendor-Specific Information

| For information about... | See...                                            |
|--------------------------|---------------------------------------------------|
| Network Appliance        | <a href="#">“Network Appliance” on page 490</a>   |
| EMC Celerra              | <a href="#">“EMC Celerra” on page 496</a>         |
| EMC IP4700               | <a href="#">“EMC IP4700” on page 499</a>          |
| Auspex                   | <a href="#">“Auspex” on page 500</a>              |
| Procom NetFORCE          | <a href="#">“Procom NetFORCE” on page 503</a>     |
| Mirapoint                | <a href="#">“Mirapoint” on page 504</a>           |
| BlueArc                  | <a href="#">“BlueArc” on page 505</a>             |
| DinoStor                 | <a href="#">“DinoStor TapeServer” on page 506</a> |
| SnapImage™               | <a href="#">“SnapImage Module” on page 507</a>    |

## Sun StorEdge EBS NDMP Requirements

The Sun StorEdge EBS software has the following requirements when using NDMP:

- 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 for recovering the Sun StorEdge EBS server after a disk crash.



**Caution** – If data is to be cloned, a second non-NDMP device must be configured to receive the cloned bootstrap data.

- In order to perform either an NDMP backup or recovery operation, you must be logged in as root.

- A separate pool must be specified for NDMP devices. NDMP and non-NDMP devices cannot be in the same pool.
- The media device handle must be a nonrewind device handle.
- To perform parallel backups, either specify multiple save sets in a single Client resource and verify multiple tape drives are available, or specify multiple Client resources in separate groups for each volume you want to back up.
- 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

Adding NDMP support to a Sun StorEdge EBS datazone environment renders the following changes:

- You can browse other NDMP hosts for save set recoveries through both the command line and the Sun StorEdge EBS Administrator program. You cannot browse NDMP hosts for backups.
- The following are not supported:
  - Archiving
  - Multiplexing
  - Staging
  - File type devices
  - Save set consolidation
  - The `jbexercise` utility

To verify that the Sun StorEdge EBS server can see the autochanger on the NDMP server, you can either reset the autochanger with the `nsrjb -HE` command or inventory the autochanger with the `nsrjb -IE` command.
- The **scanner** program generates media database entries only for data backed up with NDMP. 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 jukebox, perform one of the following before recovering the generated volumes and save sets:
  - An inventory on the jukebox to synchronize the slots and volumes for automount/unmount
  - A manual mount operation

If the corresponding index save set is available, indexes can be brought back using **nsrck L7**, providing the client is 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, issue the **nsrjb -HEvv** command.

- Sun StorEdge EBS Server resource attributes (nsrmmd polling interval, nsrmmd restart interval, and nsrmmd control timeout) do not apply to NDMP hosts.
- 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, complete a full backup on day 1 and level 1 backups on days 2 through 7.
- NDMP and non-NDMP save sets on a single volume are not supported.

## Protocol Limitations

The following protocol limitations apply to versions 1 and 2 of the NDMP protocol:

- Direct Access Restore (DAR) is not supported.
- The Save Set attribute default option All is not supported for NDMP servers. The filesystems must be manually listed to be backed up.
- The autochanger handle must be determined to run **jbconfig**.

## Important Notes on Index Processing and Memory Requirements

Backing up or recovering a large number of files may require more swap space. If a backup or recovery fails due to inadequate swap space, an error message similar to the following appears:

Failed to store index entries

If this message appears, increase the swap space. If the backup or recovery still runs slowly after increasing the swap space, try increasing the amount of memory on the Sun StorEdge EBS server.

To determine the space required to back up or recover of a large number of files, use the following formulas:

- The */nsr/index* entries require approximately 150 bytes per file entry. For example, if there are 15 million files, */nsr/index* would require 2.25 GB of space.
- To determine the required physical space for the */nsr/tmp* directory per file entry, multiply the average filename length in bytes by 2 and add 202 (2 x average filename length bytes + 202).
- The swap space requirements are 3MB plus the higher of the values for *nsrndmp\_2fh* and *nsrdmpix* (binaries called during file history processing).
  - For *nsrndmp\_2fh*, multiply the number of entries by 64.
  - For *nsrdmpix*, multiply the average number of entries per directory by the average depth of the directory tree times (the average filename length + 16 + 24) + 32 times the number of entries.

Once you determine the two values, add 3 MB to the higher value for the required swap space total. If the backup still runs slowly after increasing the swap space, increase the amount of memory on the system.

The following sections provide requirements for HP-UX hosts.

---

## Configuring the Sun StorEdge EBS Server for NDMP Operations

Before you can perform an NDMP backup, you must configure the Sun StorEdge EBS server for NDMP operations. Use the Sun StorEdge EBS Administrator program to configure the resources listed in [TABLE 12-2 on page 467](#) on the Sun StorEdge EBS server before you back up data. When you configure the Sun StorEdge EBS server for NDMP operations, 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 476](#).

---

**Note** – Resources not listed in [TABLE 12-2](#) should be configured the same as for non-NDMP operations.

---

**TABLE 12-2** Configuring Resources for NDMP

| Resource                | Configuration                                                                                                                                                                                                                                                                              |
|-------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Autochanger             | Configure an autochanger with the <b>jbconfig</b> command. For instructions, see <a href="#">“Configuring Autochanger Support for NDMP” on page 467</a> . For vendor-specific requirements, see <a href="#">“Vendor-Specific Information” on page 489</a> .                                |
| Device                  | If you are using a stand-alone device, create a Device resource for NDMP backups. For instructions, see <a href="#">“To Configure an NDMP Device Resource” on page 473</a> . If you are using an autochanger, see <a href="#">“Configuring Autochanger Support for NDMP” on page 467</a> . |
| Client                  | Create a Client resource with specific NDMP attributes. For instructions, see <a href="#">“To Configure an NDMP Client Resource” on page 474</a> . For vendor-specific requirements, see <a href="#">“Vendor-Specific Information” on page 489</a> .                                       |
| Group                   | Add the NDMP client to a backup group. If you are backing up more than 5 million files, set the Inactivity Timeout attribute to 0. For additional issues, see <a href="#">“Additional Considerations for NDMP Operations” on page 476</a> .                                                |
| Schedule                | Select a schedule for the NDMP client.                                                                                                                                                                                                                                                     |
| 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 Autochanger Support for NDMP

Before you can perform NDMP backups and recoveries, you must configure the autochangers using the **jbconfig** program. Depending on the environment, the autochanger handle may not be required to configure the jukebox, as follows:

- If the NDMP host is running NDMP version 3 or later and supports NDMP\_CONFIG interfaces, the autochanger handle *is not* required to configure the jukebox through **jbconfig**.
- If the NDMP host is running NDMP version 2 or earlier or does not support NDMP\_CONFIG interfaces, the autochanger handle *is* required to configure the jukebox.

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 [“To Configure an NDMP Device Resource” on page 473](#).

---

This section discusses the following autochanger configuration topics:

- [“Obtaining Information about NDMP Autochangers” on page 468](#)
- [“Using the inquire Command with NDMP Devices” on page 468](#)
- [“To Configure an NDMP Autochanger” on page 469](#)
- [“Sharing Autochangers among NDMP Clients” on page 471](#)
- [“Running nsrjb to Load Tapes” on page 472](#)
- [“Silo Support” on page 473](#)

## Obtaining Information about NDMP Autochangers

Before the autochanger can be configured for NDMP backups and recoveries, you must obtain the pathname of the media device. For instructions on obtaining the device pathname for a particular vendor, see [“Vendor-Specific Information” on page 489](#).

## Using the inquire Command with NDMP Devices

Use the `inquire` command to search for available devices on the NDMP tape server.

The NDMP autochanger information can be obtained two ways:

- From the location on the local server.
- On a specific NDMP host.

To obtain the NDMP autochanger information:

1. Enter the following command at the prompt:

```
inquire -N
```

2. Respond to the prompts as follows:

- a. Enter the NDMP tape server name.
- b. Enter the NDMP user name.

- c. Enter the NDMP password on NDMP host *hostname* (characters will not be echoed).

The list of devices appears.

For more information on the `inquire` command, refer to the *Legato Command Reference Guide* or the `inquire` man page.

## ▼ To Configure an NDMP Autochanger

The Sun StorEdge EBS Autochanger resource configuration tool, `jbconfig`, helps you configure autochangers for NDMP operations. Once the autochanger is configured, use the `jbedit` utility to add or remove drives and devices. For more information about `jbedit`, see [“Editing Autochanger Configurations” on page 376](#).

To configure an autochanger:

1. Make sure you have the device pathname information for the NDMP autochanger. To obtain this information, see [“Vendor-Specific Information” on page 489](#).
2. Log in as root on the Sun StorEdge EBS server.
3. Enter the following command at the prompt:

```
jbconfig
```

**Note:** To use `jbconfig`, you must know the NDMP password for the machine to which the tape device is attached.

4. When prompted, select one of the following and answer the questions that appear:
  - If the autochanger 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 autochanger 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 autochanger for a NetApp filer by using the `jbconfig` program. The script uses the term *jukebox* instead of *autochanger*. The terms are interchangeable.

---

**Note** – This example might not be valid for your NDMP host. 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 you use the `jbconfig` program to configure an autochanger, a new resource is created with the name you specified. You can view the new Autochanger resource in the Media Management section of the Sun StorEdge EBS Administrator program.

For more information on the **jbconfig** command, refer to the *Legato Command Reference Guide* or the **jbconfig** man page.



## Sharing Autochangers among NDMP Clients

The Sun StorEdge EBS software permits the control and usage of individual devices within an autochanger to different Sun StorEdge EBS hosts within a datazone. If only NDMP hosts have access to an autochanger, the autochanger has to be defined as NDMP. In this situation, the tape drives within the autochanger can still be shared and controlled by one host, and shared by the other hosts.

### *Robotic Arm Control*

The autochanger is configured for sharing during the initial autochanger setup using 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 the following regarding robotic arm control when using the `jbconfig` program:

- If the Sun StorEdge EBS server owns the robotic arm, run **jbconfig** from that server and choose Autodetected SCSI Jukebox when prompted.
- If an NDMP host owns the robotic arm, run **jbconfig** from the Sun StorEdge EBS server and choose Autodetected NDMP SCSI Jukebox.
- If a standard storage node owns the robotic arm, run **jbconfig** from that node and choose Autodetected SCSI Jukebox.

While running `jbconfig`, you are prompted to enter the following:

- The Sun StorEdge EBS hostname and device path of the server that controls the first autochanger device.
- The controlling host for *each* device in the autochanger.

For example, in a three-drive library to be shared among two NDMP hosts and a Sun StorEdge EBS server, with *Host1* controlling the arm and the first drive, when prompted specify the *hostname:device path* for *Host1* as the host controlling the autochanger device. When prompted for the owner of *Drive 2*, specify the *hostname:device path* for *Host2*, and when prompted for the owner of *Drive 3*, specify the Sun StorEdge EBS server *hostname:device path*.

For non-NDMP information about sharing autochangers, see [“Sharing Autochangers among Sun StorEdge EBS Hosts” on page 384](#).

### *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 *Legato Software Compatibility Guide*.

You can enable the DDS feature when you are configuring an autochanger through the `jbconfig` program.

### ▼ *To Enable the DDS Feature*

To enable DDS:

1. When prompted by the **jbconfig** program whether the drive is to be shared, enter 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, enter Yes.
5. When prompted, enter 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 module:

- 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 [“Dynamic Drive Sharing” on page 198](#).

## Running `nsrjb` to Load Tapes

If you run `nsrjb` to load a tape, you need to 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, the syntax for loading a tape into that drive is as follows:

```
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.

### ▼ 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 [“Devices” on page 187](#). The following section provides instructions for configuring the Device resource information when using a stand-alone device with NDMP.

Note: The jbconfig program creates the NDMP device for NDMP devices in an autochanger.

To configure an NDMP device in the Device resource:

1. For the Name attribute, specify the device attached to the NDMP host using the following syntax:

`rd=NDMP_host:tape_device_pathname`

2. Set the Target Sessions attribute to 1. (Multiplexing is not supported on NDMP hosts.)
3. For the NDMP attribute, select Yes.

---

**Note** – This option can only be set when you create the device. You cannot change the NDMP attribute after the device is created. To change the device configuration, first delete the configuration, then re-create it.

---

4. For the Remote User attribute, enter the NDMP account user name for the host.
5. For the Password attribute, enter the password for the NDMP user account.  
This password should match the password for the NDMP user account for the host.
6. Complete any other attributes as you would for a regular Sun StorEdge EBS device configuration.

## ▼ To Configure an NDMP Client Resource

For information on creating a Client resource for non-NDMP operations, see [“To Create a Client” on page 110](#). The following 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. For the Name attribute, enter the client name.
2. For the Save Set attribute, list the filesystems to be backed up. You must specify the filesystem mount points or the raw partition names here. To back up all of the client’s filesystems, specify All.

Note: If the NDMP server is running a version earlier than 3, the All option is not supported. The filesystems must be manually listed to be backed up.

If you need to back up large client filesystems, you may want to schedule each filesystem to back up separately. For example, you might create two separate clients with the same name, but with different save sets.

3. For the Remote Access attribute, list all hosts that will store and retrieve entries for the client. In addition, list any host that is used to browse entries for the NDMP host. Since the Sun StorEdge EBS server stores the index entries for the NDMP host, be sure that `root@server` is listed.
4. For the Remote User attribute, enter the NDMP account user name for the NDMP host. If this is root, be sure that root has a password.
5. For the Password attribute, enter the password for the NDMP user account. This password should match the password of the NDMP user account for the NDMP host.
6. For the Backup Command attribute, set the value to `nsrndmp_save`.

The syntax for `nsrndmp_save` is:

```
nsrndmp_save -T backup_type
```

The value for the NDMP backup type depends on the type of NDMP host. For example, NetApp, EMC, and Procom all support dump, so the value for the Backup Command attribute is:

```
nsrndmp_save -T dump
```

For a list of backup values supported by a particular NDMP host, refer to the appropriate vendor documentation.

7. For the Application Information attribute, enter the NDMP environment variables. For a list of the variables and rules for entering their values, see [“Specifying Application Information” on page 475](#).
8. For the NDMP attribute, select Yes.
9. For the Storage Node attribute, specify the appropriate NDMP tape server hostname.
10. Complete any other attributes as you would for a regular 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

- Enter each environment variable in the Application Information attribute on a separate line.
- Click *Add* after each entry.

The values you enter depend on the NDMP host.

### *Required Variables for All NDMP Hosts*

The following application information variables are required for all NDMP hosts:

- HIST=y

HIST determines whether a file history is created, which is 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 504](#).

- UPDATE=y

UPDATE forces the timestamp information to be updated with the time, date, level, and filesystem being backed up. This value must be set to *y*. If a filesystem is backed up at the same level as a previous backup of the same filesystem, the old timestamp information entry for that filesystem is overwritten with the new information.



---

**Caution** – 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 [“Vendor-Specific Information” on page 489](#). These values were correct at the time of publication and *must* be set before a backup is performed. For the latest information, check with the NDMP vendor.

## Additional Considerations for NDMP Operations

The following sections provide information about various Sun StorEdge EBS settings to modify for NDMP operations.

### Force Incremental and Interval Attribute Settings

If the Force Incremental attribute is set to Yes and the Interval attribute is set to a value that is less than 24 hours, executing a level backup from the Sun StorEdge EBS Administrator program results in a full backup being performed.

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 performed, regardless of the level specified for the backup. Since NDMP does not support incremental backups, a full backup is performed. To avoid this problem, set the Interval attribute to 24 hours or higher, or set the Force Incremental attribute to No.

### Group Parallelism Setting

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 at the same time, divide the number of drives by the number of groups and use that value for each of the group Parallelism attributes.

---

# Performing NDMP Backups

The `nsrndmp_save` command is used in NDMP operations. This command supports most of the same command line arguments as `save(1m)`. The `-T backup_type` option within `nsrndmp_save` identifies the NDMP backup type.

The following sections describe how to perform an NDMP backup using the `nsrndmp_save` command. For information on the `nsrndmp_save` command, refer to the *Legato Command Reference Guide* or the `nsrndmp_save` man page.

## Prerequisites

Perform a test backup after you complete the following:

- 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, you can schedule differential level 1 backups 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. Filesystems 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 manually or through the schedule.

To perform the test backup, run a scheduled group backup.

# Performing an NDMP Backup from the Command Prompt

You can initiate NDMP backups from the command prompt using the `nsrndmp_save` command. If you perform a backup from the command prompt, you cannot browse the files or directories. To browse files or directories that have been backed up, you must use the `nwrecover` program or the `recover` command.

---

**Note** – You must configure the client in the Sun StorEdge EBS Administrator program before performing a command line backup.

---

To perform a backup from the command prompt, use the following syntax:

```
nsrndmp_save -T backup_type -s server_name -c client_name -l backup_level -N name -g save_group -W width -LL local
```

[TABLE 12-3 on page 478](#) describes the options you can use with the `nsrndmp_save` command.

**TABLE 12-3** Command Line Options for `nsrndmp_save`

| To:                                                                                                                             | Use this option:                                |
|---------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------|
| Display the usage statement for <b>nsrndmp_save</b> .                                                                           | None (type <b>nsrndmp_save</b> and press Enter) |
| Specify the type of NDMP backup. The <b>-T</b> must be capitalized.                                                             | <b>-T backup_type</b>                           |
| Specify the name of the Sun StorEdge EBS server with NDMP Connection enabled. The local host is the default.                    | <b>-s server_name</b>                           |
| Specify the client name for starting the save session.                                                                          | <b>-c client_name</b>                           |
| Specify the backup level.                                                                                                       | <b>-l backup_level</b>                          |
| Specify the symbolic name of the save set.                                                                                      | <b>-N name</b>                                  |
| Specify the group to be backed up.                                                                                              | <b>-g save_group</b>                            |
| Provide for backward compatibility. Used to restrict the width when formatting summary information output.                      | <b>-W width</b>                                 |
| Print additional output for <b>savegrp(8)</b> . Specify <b>-LL</b> . (The <b>-L</b> option alone has no effect on the command.) | <b>-LL local</b>                                |



---

**Note** – The **-n** and **-q** options are not used but are provided for compatibility reasons.

---

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 Administrator program.

---

**Note** – When backups are performed through the `nsrndmp_save` command, the bootstrap is not automatically backed up. Without the bootstrap, you cannot perform a disaster recovery of an NDMP host. To perform a bootstrap recovery of an NDMP host, you must 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 Administrator program. All environment settings specified for that group, such as clients, pools, and schedules, are used.

---

## ▼ To Perform a Three-Party Backup

Three-party backups involve two NDMP servers. For more information, see [“Three-Party Backup” on page 459](#).

---

**Note** – For instructions on performing a remote backup, see [“To Perform a Remote Backup with a DinoStor TapeServer” on page 507](#) or [“To Perform a Remote Backup with the SnapImage Software” on page 509](#).

---

To perform a three-party backup:

1. Set up the servers as demonstrated in [FIGURE 12-2 on page 460](#).
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.

---

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 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 466](#).

## Performing DAR Backups

DAR is an optimized 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 significantly reduces recovery time of large backups.

To determine whether DAR is supported with a particular NDMP vendor, refer to the *Legato Compatibility Guides* on the Legato web site.

To specify whether to use DAR, set the application information variables before the backup, as follows:

```
DIRECT=y
HIST=y
```

DAR is supported in the following NDMP configurations:

- Local
- Three-party
- Remote

A DAR backup is performed in the same manner as a regular NDMP backup. A DAR recovery is performed through the `recover` command or the `nwrecover` program.



---

**Caution** – To perform a DAR recovery, use the **recover** command or the **nwrecover** program. You cannot use **nsrndmp\_recover** for a recovery with DAR.

---

---

# NDMP Cloning

NDMP cloning is set through the same Sun StorEdge EBS clone interface as non-NDMP backups. For information on how to clone data, see [Chapter 9](#).



---

**Caution** – 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.

---

The following functionality applies to NDMP cloning:

- The Sun StorEdge EBS Administrator program and 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 tape devices.
- 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 manually recovered. For more information, see [“Volume Cloning and Recovery” on page 343](#) and [“Recovering Cloned Save Sets” on page 355](#).

## NDMP Index and Bootstrap Cloning

Index and bootstrap cloning is performed only through a non-NDMP cloning device. Index and bootstrap cloning can be achieved either through the command line using `nsrclone`, the Sun StorEdge EBS Administrator, or a scheduled group with automatic cloning. For more information on automatic cloning, see [“Automatically Cloning a Save Set” on page 350](#).

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 NDMP save sets or volumes.

- A non-NDMP device that is used for backing up bootstrap information and indexes.
- A non-NDMP cloning device that is used for cloning indexes and bootstrap from an NDMP backup.

---

**Note** – The indexes and bootstrap are non-NDMP save sets, so a non-NDMP device is needed to clone those non-NDMP save sets. The indexes and bootstrap both go to the same volume, unless they are configured to go to separate devices.

---

When cloning is performed manually, a minimum of two devices are required, one to read the existing data and one to write the cloned data. For more information, see [“To Manually Clone a Save Set” on page 347](#).

---

## Performing NDMP Recoveries

NDMP save set and directory-level recoveries are performed through the Sun StorEdge EBS Administrator program or the `nsrndmp_recover` command. NDMP file-by-file recoveries are performed through the `nwrecover` program 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 using the Sun StorEdge EBS Administrator program, the `nsrndmp_recover` command is selected automatically.



---

**Caution** – 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*.

---

## Performing a Destructive Save Set Recovery

The following sections describe how to perform a destructive save set recovery from the Sun StorEdge EBS Administrator program or from the command prompt. With a destructive save set recovery, the 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 486](#).

## ▼ To Perform a Recovery from the Sun StorEdge EBS Administrator Program

To perform a destructive save set recovery from the Sun StorEdge EBS Administrator program:

1. From the Save Set menu, select Recover.
2. Enter or 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 list box.

3. From the Save Set 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 Instances window.

4. Select the particular instance of the save set's backup to recover.
5. Click Recover.
6. If needed, specify a subdirectory in the Path to Recover attribute.
7. Click Start.

## ▼ To Perform a Recovery from the Command Prompt

To perform a save set recovery from the command prompt with the `nsrndmp_recover` command, use the following syntax:

```
nsrndmp_recover -c client -s server -S ssid/cloneid
```

TABLE 12-4 on page 483 describes the options you can use with the `nsrndmp_recover` command for a destructive save set recovery.

**TABLE 12-4** Command Line Options for `nsrndmp_recover` (1 of 2)

| To:                                                                                                                                                                         | Use This Option:     |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------|
| Specify the Sun StorEdge EBS server hostname. The default is the local host.                                                                                                | -s <i>server</i>     |
| Specify the name of the computer that saved the files, for example, the NDMP host.                                                                                          | -c <i>client</i>     |
| Specify the pathname of the raw device to recover the data from a SnapImage backup. This value is optional and should only be used to recover data from a SnapImage backup. | -r <i>raw_device</i> |

**TABLE 12-4** Command Line Options for `nsrndmp_recover` (2 of 2)

| To:                                                                                                                                                                                                                                                                  | Use This Option:              |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------|
| Specify the save set ID ( <i>ssid</i> ). To obtain the <i>ssid</i> , issue the <b><i>mminfo -avot</i></b> command.<br>When there are multiple clone instances of a save set, the cloneid can also be specified to select that particular clone to recover data from. | <b><i>-S ssid/cloneid</i></b> |
| Specify whether to verify the paths against the index database.                                                                                                                                                                                                      | <b><i>-v verify_flag</i></b>  |
| Recover a specific directory from the entire save set, using the directory-level save set restore feature. Specify the full path of the directory. If the entire save set is being recovered, do not use this option.                                                | <b><i>paths...</i></b>        |

For example, to recover data to a particular raw device on a client named *venus*, with a server named *mars*, issue the following command:

```
nsrndmp_recover -c venus -s mars -S 2485094913
```

For additional information on this command, refer to the *Legato Command Reference Guide* or the `nsrndmp_recover` man page.

## Performing a Directory-Level Recovery

The following sections describe how to perform a directory-level recovery from the Sun StorEdge EBS Administrator program or from the command prompt. For information about directed recoveries, see [“Performing a Directed Recovery” on page 486](#).

### ▼ To Perform a Directory-Level Recovery from the Sun StorEdge EBS Administrator Program

To perform a directory-level recovery from the Sun StorEdge EBS Administrator program:

1. From the Save Set menu, select Recover.
2. Enter or 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 list box.

3. In the Save Set list box, select the save set with the data to recover.  
The save set's backup instances appear in the Instances window.
4. Select the particular instance of the save set's backup to recover.
5. Click Recover.
6. To recover a specific directory, specify the full path of the directory in the Paths to Recover attribute. If you are recovering the entire save set, leave this attribute blank.  
**Note:** The permissions of the recovered files and directories match the corresponding original permissions.
7. If needed, specify a subdirectory in the Path to Recover attribute.
8. Click Start.

## ▼ To Perform a Directory-Level Recovery from the Command Prompt

To perform a directory-level recovery with the `nsrndmp_recover` command, use the following syntax:

```
nsrndmp_recover -c client -s server -S ssid -v verify_flag paths...
```

[TABLE 12-4 on page 483](#) lists the available options for the `nsrndmp_recover` command.

## Performing a File-by-File Recovery

A file-by-file recovery only recovers a select group of files or directories. The following sections describe how to perform a file-by-file recovery from the `nwrecover` program or the command prompt.

## ▼ To Perform a File-by-File Recovery from the Sun StorEdge EBS Recover Program

To recover files using the Sun StorEdge EBS recover program:

1. In the Sun StorEdge EBS recover program, select Recover from the Operations menu.
2. From the left pane, select the directory folder that contains the files to recover.
3. From the right pane, select the files or directories.

4. Click Mark to mark the selected file or directory for recovery.

A check mark appears next to the items you select. If you select a disk volume or directory, all of its nested subdirectories and files are also marked for recovery.

Click Unmark to unmark a file or directory.

5. When all the files to recover are selected, click Start.

## ▼ To Perform a File-by-File Recovery from the Command Prompt

To perform a file-by-file recovery from the command prompt, use the interactive `recover` program, as follows:

1. From the command prompt, enter the following:

```
recover -s server
```

2. When prompted, enter the directory to browse, such as **cd /**.
3. Issue the appropriate commands to add the necessary files to the recover list.

Note: For a complete list of command options, refer to the *Legato NetWorker Command Reference Guide*.

4. When all the required files have been added, enter the following:

```
recover
```

The files are recovered.

## 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.



---

**Caution** – 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 user name 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.

---

## ▼ To Perform a Directed Save Set Recovery

You can perform a directed save set recovery using either the Sun StorEdge EBS User Administrator program or the `nsrndmp_recover` command:

- Using the NetWorker User Administrator program:
  - a. From the Save Set menu, select 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. For the Relocate Recovered Data To This Raw Device text box, specify the mount point in the following format:

*target\_hostname::/mount\_point*

For example, *molokai::/usr*.

---

**Note** – For the mount point, specify the NAS box 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. Click OK.
  - f. In the Save Set window, click OK.
- Using **`nsrndmp_recover`**, enter the following command:

```
nsrndmp_recover -c client -m target_hostname::/mount_point -S ssid
```

**Note:** For the mount point, specify the NAS box 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.

---

TABLE 12-5 on page 488 lists the options available for a directed recovery from the command prompt.

TABLE 12-5 Command Line Options for nsrndmp\_recover

| To:                                                                                                                                                                                                                   | Use This Option:      |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------|
| Specify the Sun StorEdge EBS server hostname. The default is the local host.                                                                                                                                          | -s <i>server</i>      |
| Specify the name of the computer that saved the files, for example, the NDMP host.                                                                                                                                    | -c <i>client</i>      |
| Specify the save set ID ( <i>ssid</i> ). To obtain the <i>ssid</i> , issue the <b>mminfo -avot</b> command.                                                                                                           | -S <i>ssid</i>        |
| The NFS mount point of the raw device specified by the <b>-r</b> option.                                                                                                                                              | -m <i>mount_point</i> |
| Specify whether to verify the paths against the index database.                                                                                                                                                       | -v <i>verify_flag</i> |
| Recover a specific directory from the entire save set, using the directory-level save set restore feature. Specify the full path of the directory. If the entire save set is being recovered, do not use this option. | paths...              |

▼ To Perform a Directed Index Recovery

You can perform a directed index recovery using either the recover program or either the NetWorker User program or the recover command:

- Using the Sun StorEdge EBS recover program:
  - a. Select Relocate from the Operations menu.
  - b. For the Relocate attribute, enter:

target\_hostname : /mount\_point

- c. Complete the other attributes as you would for a regular recovery.
- Using the interactive **recover** program, perform the following:
  - a. From the command prompt, enter the following:

# **recover -s** server

- b. When prompted, enter the directory to browse, such as **cd /**.

- c. Issue the appropriate commands to add the necessary files to the recover list.

Note: For a complete list of command options, refer to the *Legato NetWorker Command Reference Guide*.

- d. To restore the files to a different location, enter the following:

```
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 box hostname. If you are using Data ONTAP, you might need to add a backslash (/) after the mount point; for example, *target\_hostname::/mount\_point/*.

- e. When all the required files have been added, enter the following:

```
recover
```

The files are recovered to the specified location.

---

## Vendor-Specific Information

The following sections cover limitations and configuration requirements for specific vendors:

- [“Filers That Act as Data Servers” on page 489](#)
- [“Filers That Act as Tape Servers Only” on page 506](#)

### Filers That Act as Data Servers

This sections contains configuration information for the following vendor products that act as NDMP data servers:

- [“Network Appliance” on page 490](#)
- [“EMC Celerra” on page 496](#)
- [“EMC IP4700” on page 499](#)
- [“Auspex” on page 500](#)
- [“Procom NetFORCE” on page 503](#)

- “Mirapoint” on page 504
- “BlueArc” on page 505

## Network Appliance

The following sections provide information about limitations, autochanger configuration, and application information variables for the NetApp filer. For a complete list of supported Data ONTAP versions, refer to the *Legato Software Compatibility Guide*.

### Features and Limitations

The following sections cover NDMP issues that are specific to NetApp hosts.

---

**Note** – NetApp passwords are limited to eight characters.

---

#### *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 using a NetApp filer and supported Data ONTAP version.

The supported configurations are:

- 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 Data ONTAP versions support which languages. 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.

The following 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 non-ASCII characters using NDMP, the UTF-8 settings on the filer and the Sun StorEdge EBS server must match.

To set the UTF-8 variable:

- If the data from the filer is not in UTF-8 format, set the UTF-8 environment variable as follows, or leave the value blank:

**UTF8=n**

The default is *n*.

- If the data from the filer is in UTF-8 format, set the UTF-8 environment variable as follows:

**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 (fr), Spanish (es), and Portuguese (pt), 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 you use an earlier NDMP version.

### ▼ *To Verify the NDMP Version*

To verify the NDMP version:

1. Log in to the NetApp host.
2. Enter the following command at the prompt:

```
ndmpd version
```

The NDMP version is displayed.

### ▼ *To Change the NDMP Version*

To change the NDMP version:

1. Stop the NDMP process by entering the following command at the prompt:

```
ndmpd off
```

The NetApp host confirms that the NDMP process has stopped.

2. To change the NDMP version, enter the following 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.

3. To restart the NDMP process, enter the following command:

```
ndmpd on
```

The NetApp host confirms that the NDMP process has started. You can now run DAR.

## 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.

### ▼ To Find the Device Pathname for a NetApp Host

To configure the autochanger for NDMP backups and recoveries, you must obtain the pathname of the media device. Use the pathname to configure the autochanger with the `jbconfig` program, as described in [“Configuring Autochanger Support for NDMP” on page 467](#).



---

**Caution** – 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. Verify that the tape device is installed on the NetApp filer by entering the following command at the prompt:

```
sysconfig -v
```

The host responds with:

```
NetApp Release 6.2: Sun Mar 12 13:29:37 PST 2000
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
```

2. Determine the device pathname by entering the following command at the prompt:

```
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 entering the following command at the prompt:

---

```
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.



**Caution** – When **jbconfig** prompts for the hostname and device path for a NetApp device, do *not* enter a slash before the device name; this will cause the **jbconfig** program to complete with no errors, but NetApp will not recognize the tape device or autochanger, and you will not be able 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 12-6 on page 495](#) lists NetApp-specific variables that may be required. Enter these variables in the Application Information attribute of the Client resource, as described in [“To Configure an NDMP Client Resource” on page 474](#).

**TABLE 12-6** NetApp Application Information Variables (1 of 2)

| Variable       | Required?              | Values and Meaning                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|----------------|------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| DIRECT= y/n    | Yes for DAR operations | y/n. 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 <a href="#">“Performing DAR Backups” on page 480</a> .                                                                                                                                                                            |
| EXCLUDE=       | No                     | This string value determines which files are excluded from backup. The following rules apply: <ul style="list-style-type: none"><li>• The string is a filename.</li><li>• The asterisk (*) is accepted as a wildcard, but it must be the first or last character in the string (or both).</li><li>• To list multiple files, separate each name with a comma. A comma cannot appear as part of the filename.</li><li>• You can specify up to 32 strings.</li></ul> |
| EXTRACT_ACL= y | Yes to recover ACLs    | y/n. To recover access control lists (ACLs) when using DAR with a NetApp host, you must set the value to y. The default is n.                                                                                                                                                                                                                                                                                                                                     |

**TABLE 12-6** NetApp Application Information Variables (2 of 2)

| Variable | Required?                    | Values and Meaning                                                                                                                                                                                                                                                                                                                                                                    |
|----------|------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| UTF8=y/n | Yes for UTF-8 formatted data | y/n. Specifies whether the data on the filer is in UTF-8 format. To support non-ASCII characters using NDMP, the UTF-8 settings on the filer and the Sun StorEdge EBS server must match.<br>If the data from the filer is not in UTF-8 format, set the value to <i>n</i> .<br>If the data from the filer is in UTF-8 format, set the value to <i>y</i> .<br>The default is <i>n</i> . |

## EMC Celerra

The following 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 *Legato Software Compatibility Guide*.

### EMC Celerra Considerations with NDMP

The following sections discuss NDMP issues that are specific to EMC Celerra filers.

#### *Using an Earlier Version of NDMP on EMC Celerra Data Movers*

Generally, 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 functionality 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.

#### ▼ *To Lower the NDMP Version on EMC Celerra*

To set the NDMP version on EMC Celerra data movers so that the data mover only negotiates at the specified NDMP version or an earlier version:

---

**Note** – The NDMP version must be set individually on each data mover that requires the modification.

---

1. Log in as root to the EMC Celerra Control Station.
2. At the command prompt, enter the following:

```
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 entering the following command:

```
vi netd
```

4. At the end of the NDMP string in the *netd* file, add the following:

```
proto=ndmp_version_#
```

For example, the resulting NDMP string looks like the following when setting the NDMP version to 3:

```
ndmp port=10000 proto=3
```

5. Save the file.
6. Reboot the data mover by entering the following command at the prompt:

```
server_cpu server_name -reboot -m now
```

### ▼ *To Set the ntape Parameter*

To help avoid tape drive issues, EMC recommends setting 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. Using 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 the following line:

```
[NDMP ntape=n]
```

**Note:** The value of *n* cannot be greater than 8.

3. In order for the changes to take effect, reboot the data mover.

## ▼ To Find the Device Pathname for an EMC Celerra Host

To configure the autochanger for NDMP backups and recoveries, you must obtain the pathname of the media device. Use the pathname to configure the autochanger with the `jbconfig` program, as described in [“Configuring Autochanger Support for NDMP” on page 467](#).



---

**Caution** – 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 the following:

- Re-create all devices and view all nondisks by entering the following 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, enter:

**`http://EMCcontrolhostIP:8000/top_level.htm`.**

- b. At the prompt, enter your user name 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 12-7 on page 499](#) lists Celerra-specific variables that may be required. Enter these variables in the Application Information attribute of the Client resource, as described in [“To Configure an NDMP Client Resource” on page 474](#).

**TABLE 12-7** EMC Celerra Variables

| Variable    | Required?               | Values and Meaning                                                                                                                                                                                                                       |
|-------------|-------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| DIRECT= y/n | Yes for DAR operations  | y/n. 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 <a href="#">“Performing DAR Backups” on page 480</a> .    |
| OPTIONS=NT  | Yes for ACL restoration | This value must be set to make sure all ACLs are restored when NT data is contained in the backup path.<br><b>Note:</b> In addition to setting this variable, <b>axttrp</b> must exist in the <code>/nas/server/slot_#/netd</code> file. |

## EMC IP4700

The following 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 *Legato Software Compatibility Guide*.

### ▼ To Find the Device Pathname for an EMC IP4700 Host

To configure the autochanger for NDMP backups and recoveries, you must obtain the pathname of the media device. 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 Autochanger Support for NDMP” on page 467](#).

To obtain the autochanger and tape drive device pathnames:

1. In the browser window, access the EMC IP4700 Administrator program at:  
**http://IP4700name/**
2. Click the Tape Drives button.
3. Enter 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 12-8 on page 500](#) lists IP4700-specific variables that may be required. Enter these variables in the Application Information attribute of the Client resource, as described in [“To Configure an NDMP Client Resource” on page 474](#).

**TABLE 12-8** EMC IP4700 Variables

| Variable            | Required | Values and Meaning                                                                                                                                                                                                                                     |
|---------------------|----------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| FAST_INCREMENTAL= y | No       | y/n. Performs a fast incremental backup. This function works only if fast incremental is enabled on the volume. You must have the SnapView /IP option to perform a fast incremental backup. The default is y.                                          |
| SNAPSHOT= y         | No       | y/n. Backs up a snap copy of the filesystem. 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. |

## Auspex

The following 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 *Legato Software Compatibility Guide*.

## Auspex-Specific Considerations with NDMP

The following are NDMP issues that are specific to Auspex filers.

## *Auspex NS3000 Requirements*

In order to support Auspex NS3000, you must perform the following:

- 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.

### ▼ *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 the following:  
`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 the following 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*

The following limitations apply to Auspex backups and recoveries:

- Root partition backups are not supported. Do not back up the root partition.
- It is recommended that file-by-file recoveries not exceed 10,000 files per recovery.

To recover a larger number of files, perform *one* of the following:

- Set up multiple save sets for a given client to break the filesystem into directory-level save sets rather than whole volumes.
- Perform a save set recovery using the Sun StorEdge EBS Administrator program or the `nsrndmp_recover` command from the command prompt.

## ▼ To Find the Device Pathname for an Auspex Host

To configure the autochanger for NDMP backups and recoveries, you must obtain the pathname of the media device. 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 Autochanger Support for NDMP” on page 467](#).

**Note:** 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 12-9 on page 503](#) lists Auspex-specific variables that may be required. Enter these variables in the Application Information attribute of the Client resource, as described in [“To Configure an NDMP Client Resource” on page 474](#).

**TABLE 12-9** Auspex Variables

| Variable      | Required?   | Values and Meaning                                                                                                                                                                                                                                         |
|---------------|-------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| DIRECT= y / n | Yes for DAR | y/n. 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.<br>For additional information about DAR, see <a href="#">“Performing DAR Backups” on page 480</a> . |
| SNAPSHOT= y   | No          | y/n. Enables or disables snapshots. The default is n.                                                                                                                                                                                                      |

## Procom NetFORCE

The following sections provide information about autochanger configuration for the Procom host. For a complete list of supported versions, refer to the *Legato Software Compatibility Guide*.

### ▼ To Find the Device Pathname for a Procom NetFORCE Host

To determine the device pathname, enter the following command at the prompt:

```
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 Autochanger Support for NDMP” on page 467](#).

---

**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.

---

## Post-Configuration Requirements for Procom NetFORCE Host

After using the `jbconfig` program to configure an autochanger for a Procom NetFORCE host, perform the following:

1. Reset the autochanger by entering the following command:

```
/usr/sbin/nsrjb -HE
```

2. If you want to read previously labeled Sun StorEdge EBS tapes back into the database, you must inventory the autochanger:

```
/usr/sbin/nsrjb -IE
```

## Mirapoint

The following sections provide information about requirements and autochanger configuration for the Mirapoint host. For a complete list of supported versions, refer to the *Legato Software Compatibility Guide*.

### Requirements

The following 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 restore.
- 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 using the date of the image in order to perform the selection.

### Device Pathname for a Mirapoint Host

The device pathname is required to successfully run the `jbconfig` program. The device pathname is `/dev/nrst<n>`, where *n* starts at 0 and increases one number for each tape drive. This value is constant. Use the pathname to configure the autochanger with the `jbconfig` program, as described in [“Configuring Autochanger Support for NDMP” on page 467](#).

Once the jukebox is configured, certain attributes must be modified in the Jukebox resource window. Set the Load Sleep and Eject Sleep attributes to the following values:

- Load Sleep: 90
- Eject Sleep: 30

---

**Note** – If the NDMP tape server is running version 2 or does not support NDMP\_CONFIG interfaces, you must enter 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 the following in mind:

- When performing NDMP restore operations, a `$__NDMP__` directory may be created at the root level of the filesystem to which the data is being restored. This directory holds the file list that the Sun StorEdge EBS server passes for an index restore, if that file list contains more than 1,024 filenames. Do not change the directory and its contents during an active restore operation. If no restore 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 restore operations, a message similar to the following may appear:

`NDMP session: Unknown environment variable name ignored.`

You can ignore this message, as it does not impact the operation.

For complete information on configuring a BlueArc system for NDMP operations, refer to "Backing Up Data with NDMP" in the *BlueArc System Administrator's Guide*.

## Filers That Act as Tape Servers Only

The following sections provide information about the following filers that act as tape servers only.

### DinoStor TapeServer

The following 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. Keep the following in mind 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 a jukebox 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.

#### ▼ To Specify Device Handles for a DinoStor TapeServer

When configuring the jukebox with the `jbconfig` utility, the DinoStor TapeServer has one hostname with multiple device handles. Use the hostname to configure the autochanger with the `jbconfig` program, as described in [“Configuring Autochanger Support for NDMP” on page 467](#).

To specify the device handles:

1. Access the DinoStor TapeServer interface.
2. Click the Configure page.
3. Click the SCSI tab.

The device handles are listed on the left side. When running `jbconfig`, specify the DinoStor hostname as the device that controls the robotics and the device handles as the additional devices.

## ▼ 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 [“To Configure an NDMP Client Resource” on page 474](#).

## SnapImage Module

The following sections discuss the SnapImage 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 `jbconfig` program, as described in [“Configuring Autochanger Support for NDMP” on page 467](#).

## ▼ To Identify the Device Pathname for a Solaris Computer

To identify the device pathname, determine the media device pathname by entering the following command at the prompt:

```
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 entering the following command at the prompt:

---

```
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.

## ▼ To Identify the Device Pathname for an HP-UX Computer

To display the device pathname, use *one* of the following methods:

- Use the `ioscan` command by entering the following command at the prompt:

```
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.

---

## ▼ To Perform a Remote Backup with the SnapImage Software

To perform a remote backup:

1. Set up the servers as shown in either [FIGURE 12-3 on page 461](#) or [FIGURE 12-4 on page 462](#).
2. Install the Sun StorEdge EBS software on the server designated as the Sun StorEdge EBS server.
3. Install the SnapImage Module 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 software on that Sun StorEdge EBS server.
  - If you are backing up the data to a different NDMP server, install the SnapImage Module software on the host to which the data is being backed up.

For instructions on installing the SnapImage Module software, refer to the *Legato NetWorker SnapImage Module Installation and Administrator's Guide, UNIX Version*. Do not use the configuration instructions in the SnapImage Module guide. Instead, use the configuration instructions provided in this chapter.

4. To allow other servers to access the SnapImage server, run the `ndmp_passwd` command, located in `/usr/ndmphone/bin`.

- a. Use the following syntax:

```
ndmp_passwd -u user
```

where *user* is the root or administrator name of the SnapImage server.

- b. When prompted, enter the password.

5. 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 [“To Configure an NDMP Client Resource” on page 474](#).



## SNMP Module

---

The Sun StorEdge Enterprise Backup Software Simple Network Management Protocol (SNMP) Module allows Sun StorEdge EBS servers to send notification messages to SNMP management agents.

SNMP is a protocol for network transactions that specifies the transfer of structured management information between SNMP managers and agents. This chapter describes the Sun StorEdge EBS server implementation of SNMP.

For detailed explanations of SNMP management operation, including how to configure your management software, refer to your network management software documentation.

The main sections in this chapter are:

- [“SNMP Traps” on page 511](#)
- [“Configuring Sun StorEdge EBS SNMP Notifications” on page 512](#)
- [“Configuring SNMP Management Software” on page 515](#)

---

## SNMP Traps

The Sun 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 (in this case, 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 Sun 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 Sun StorEdge EBS SNMP notification, see [“Configuring Sun StorEdge EBS SNMP Notifications” on page 512](#).

Additionally, SNMP-enabled network management software must be configured to accept traps from the Sun StorEdge EBS server. Depending on the configuration of the network management software, traps received from the Sun StorEdge EBS server may appear in the management software's graphical interface, may be logged into a file, or processed in another manner.

---

# Configuring Sun StorEdge EBS SNMP Notifications

Sun StorEdge EBS software provides notifications to a variety of resources about Sun StorEdge EBS server events. The Sun StorEdge EBS SNMP Module is one of those resources. The module then forwards the notification to the SNMP management software 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 Sun StorEdge EBS SNMP notification, you must first enable the Sun StorEdge EBS SNMP module. For information about enabling the Sun StorEdge EBS SNMP module, see the Sun StorEdge EBS Roadmap.

## Command Line Options for nsrtrap

The Sun StorEdge EBS SNMP Module uses the `nsrtrap` program to communicate SNMP traps from the Sun StorEdge EBS server to the SNMP management software. [TABLE 13-1 on page 512](#) lists the command line options that can be included in the Action attribute when the SNMP notification is configured.

**TABLE 13-1** Command Line Options for nsrtrap (1 of 2)

| Option              | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|---------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| -C <i>community</i> | This option allows you to specify 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, meaning that the public community can receive traps from the Sun StorEdge EBS server.<br>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 using this option when you configure the SNMP notification. |

**TABLE 13-1** Command Line Options for nsrtrap (2 of 2)

| Option                  | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|-------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| -t <i>trap_type</i>     | This option is used to set the type of trap the Sun 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 (error messages, for example), 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.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| -s <i>specific_type</i> | <p>This option is 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 coming from the Sun StorEdge EBS server.</p> <p>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:</p> <pre><b>/usr/sbin/nsr/nsrtrap -s 1 host</b></pre> <p>With this setting, the SNMP management software can be configured to recognize that Sun StorEdge EBS traps with the specific trap type of <b>1</b> 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.</p> |
| -v                      | This option is used to set the output mode to verbose. In verbose mode, nsrtrap echoes the community, trap type, specific trap type, and the hostname or IP address to the command line.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |

## ▼ To Modify the Preconfigured Sun 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 Sun StorEdge EBS Administrator program, select Notifications from the Customize menu to open the Notifications window.
2. From the View menu, select Details to display the event and priorities for this notification.
3. Highlight the SNMP Notification.

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 512](#).

The events and priorities associated with the preconfigured SNMP notification cannot be modified. To set different events and priorities for the SNMP notification, see [“To Create Sun StorEdge EBS SNMP Notifications” on page 514](#).

## ▼ To Create Sun StorEdge EBS SNMP Notifications

To create additional Sun StorEdge EBS SNMP notifications:

1. In the Sun StorEdge EBS Administrator program, select Notifications from the Customize menu to open the Notifications window.
2. From the View menu, select Details to display the event and priorities for this notification.
3. Click Create.
4. For the Name attribute, enter "SNMP notification request."
5. (Optional) For the Comment attribute, enter a description of the notification.
6. For the Action attribute, enter the following:

```
/usr/sbin/nsr/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 command line options for `nsrtrap`, such as the SNMP community, in this attribute if necessary. For more information about command line options, see [“Command Line Options for nsrtrap” on page 512](#).

For example, the Action attribute might read:

```
/usr/sbin/nsr/nsrtrap SNMPhost -c ebs
```

7. Select the events and priorities for the notification to communicate.

---

**Note** – The events and priorities cannot be modified after the notification is created.

---

8. Click Apply.

---

## 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. Instructions for configuring SNMP management software to accept specific traps vary with the type of management software you are using.

For specific instructions on configuring the types of traps that the software will accept, 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 that you want it to accept. This is done by using the Structure of Management Information (SMI) Network Management Private Enterprise Code that applies to the specific network application that will be sending traps to your management 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, icons representing each Sun StorEdge EBS server appear on the network management console. You can then configure the management software to determine what it should do when it receives a trap from a Sun StorEdge EBS server. Using the network management software, you may want to:

- Configure the way in which the management software indicates 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 setting 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 512](#).

# DiskXtender Data Manager Filesystem Support

---

This chapter discusses Sun StorEdge EBS support for backing up and recovering data that resides on DiskXtender Data Manager (DXDM) filesystems.

Sun StorEdge EBS backup and recover provides DXDM filesystem support through an engineered awareness of Data Management API (DMAPI) metadata. This added support is automatic. When backup or recovery of a DXDM filesystem is requested, the Sun StorEdge EBS software adjusts to the additional requirements without administrative intervention.

The following sections describe this added support:

- [“Supported Configurations” on page 517](#)
- [“DiskXtender Data Manager Filesystem Overview” on page 519](#)
- [“Backup of DiskXtender Data Manager Filesystems” on page 522](#)
- [“Recovery of DiskXtender Data Manager Filesystems” on page 525](#)

---

## Supported Configurations

The requirements for using Sun StorEdge EBS backup and recover with DXDM filesystems are described in the following sections:

- [“Operating Systems” on page 518](#)
- [“Filesystems” on page 518](#)
- [“Sun StorEdge EBS Configurations” on page 518](#)
- [“Path Information” on page 518](#)
- [“Permissions” on page 519](#)

# Operating Systems

The following operating systems are supported:

- AIX version 4.3.3
- HP-UX 11.11 (11i version 1.0) on a HP-UX PA-RISC server in 64-bit mode
- Solaris release 6 for SPARC
- Solaris releases 7, and 8 for 64-bit SPARC

# Filesystems

The following filesystem types are supported:

- UNIX filesystem (UFS)
- VERITAS filesystem (VxFS) 3.4 on Solaris and HP-UX
- Journaled filesystem (JFS) on AIX

# Sun StorEdge EBS Configurations

For Sun StorEdge EBS backup and recover support of DXDM filesystems the following configurations are permissible:

- DXDM software and Sun StorEdge EBS server, client, or storage node software installed on the same computer.
- DXDM software installed on a computer that is a client of a computer running Sun StorEdge EBS server, client, or storage node software.
- DXDM software installed on a computer that has Sun StorEdge EBS HSM software installed and enabled.

# Path Information

The file `/etc/dxuldm.path` is required by Sun StorEdge EBS backup and recover when supporting a DXDM filesystem. It is automatically created during DXDM installation.



## Permissions

The Sun StorEdge EBS daemon `nsrexecd` controls automated backup and recovery of DXDM filesystems. It is configured to run with set user ID (suid) root permissions. Root permissions are also required to perform manual Sun StorEdge EBS backup and recover operations with DXDM filesystems.

---

## DiskXtender Data Manager Filesystem Overview

DXDM filesystems use the following enhancements not found in standard filesystems:

- The creation and retention of DMAPI information for each file.
- The ability to migrate files to a storage target.
- The ability to purge file data from the filesystem after migration.
- The retention of a data stub for each purged file.

Through these enhancements DXDM provides filesystem access to large numbers of files while storing the bulk of the file data on one or more storage target systems.

## File Data in a DXDM Filesystem

When a file is placed into a DXDM filesystem it is available to list, view, access and change the same as in any standard UNIX filesystem. The difference is that as soon as a file is placed into a DXDM filesystem, DMAPI metadata is created for the file. This extended metadata permits the transparent archiving of the file's data outside of the filesystem.

Some of the DMAPI information that is tracked on each file includes:

- The file's migration status.
- The file's purge status.
- The file's data stub size.

After a period of time the file's data is migrated to a storage target. As shown in [FIGURE 14-1 on page 521](#) the storage target can be a DiskXtender Storage Manager (DXSM) system or an EMC Centera Storage System (Centera). After a file's data is migrated it exists on both the DXDM filesystem and the storage target.

After a period of time without changes the file's data is purged from the filesystem. The filesystem retains the file's DMAPi metadata and data stub. The data stub consists of a user-configured number of bytes from the beginning of the file.

## User Access to File Data

[FIGURE 14-1 on page 521](#) depicts eight DXDM filesystems on one host system. Four of the filesystems use a DXSM storage target and four use a Centera. The diagram shows that the data from each DXDM filesystem exists in its own namespace on the storage target. On a DXSM storage target relative pathnames and capabilities are used to identify files. On a Centera storage target C-Clip data is created to identify each file.

This storage target namespace information is not required by DXDM filesystem users for access to the files. Users need only know the full pathname of a file on the DXDM host system or, when access is provided over NFS, its relative pathname from the filesystem's NFS mount point.

When a file on a DXDM filesystem is changed it is marked as not migrated. Any previously migrated data for the changed file is retained on the storage target but renamed with a version label.

When a file is deleted from a DXDM filesystem it is fully removed from the filesystem. Any previously migrated data is retained on the storage target and renamed with a version label and a deleted label.




**FIGURE 14-1** Prototypical DXDM Installation

---

# Backup of DiskXtender Data Manager Filesystems

A DXDM filesystem consists of files and metadata. The files can be in a variety of migration and purge states. The metadata is contained in the filesystems's DMAPI attributes file. [TABLE 14-1](#) shows the types of files and metadata that can exist in a DXDM filesystem and notes the types that are included in a Sun StorEdge EBS backup.

**TABLE 14-1** DXDM Filesystem Data Types Included in a Sun StorEdge EBS Backup

| Data Type                                                       | Included?        |
|-----------------------------------------------------------------|------------------|
| New files which have not been migrated.                         | Yes              |
| Changed files which have not been migrated since being changed. | Yes              |
| Files which have been migrated but not purged.                  | Yes              |
| Data stubs for files which have been migrated and purged.       | Yes              |
| File data which exists only on a storage target system.         | No               |
| DMAPI metadata.                                                 | Yes <sup>a</sup> |

<sup>a</sup> The DMAPI metadata is included by default. It can safely be excluded. For more information, see [“Excluding the DMAPI Attributes File” on page 523](#).

File data which has not changed since being migrated is considered to be protected by the storage target. Sun StorEdge EBS backup automatically excludes this data from its client/save sets when the data has been purged. This results in the following benefits:

- Smaller client/save sets.  
Purged data which is protected on the storage target is excluded from the client/save set. This is normally the bulk of the data in a DXDM filesystem.
- Much faster backups.  
In addition to the time savings derived from the smaller client/save set, Sun StorEdge EBS backup prevents the lengthy process of retrieving purged file data from the storage target.

---

**Note** – It is recommended that DXDM filesystem client/save sets exclude the DMAPI metadata. For more information, see [“Excluding the DMAPI Attributes File” on page 523](#).

---

## ▼ To Backup a DXDM Filesystem

Backups of a DXDM filesystem are performed in the same manner as backups of a standard filesystem. Backups can be performed automatically and manually. Both methods require root permissions. For information on scheduling automatic backups, see [Chapter 6](#). For information on performing manual backups, see [Chapter 7](#).

## Excluding the DMAPI Attributes File

Every DXDM filesystem has a DMAPI attributes file which stores volatile DMAPI metadata for the filesystem. This file does not need to be backed up because it is completely re-created during the recovery process. Since these files can become quite large, a global Directive resource normally should be created to simplify the exclusion process.

## ▼ To Exclude a DMAPI Attributes File

To exclude a DMAPI attributes file:

1. Create a global Directive resource which excludes `.DMATTR` files.  
DXDM stores DMAPI attribute information in files located at the top level of each filesystem. The filename for these files is `.DMATTR`. When creating the global Directive resource, use the Application Specific Module named Skip to exclude `.DMATTR` files. For more information on directives, see [“Directives” on page 144](#).
2. Apply the Directive resource created in [Step 1](#) when creating client/save sets for the filesystem.

## Aborted Backups

Sun StorEdge EBS initiates a DMAPI process whenever it starts a backup of a DXDM filesystem. DXDM assigns a DMAPI session label of DXULDMLIB to this type of process. At the conclusion of a successful backup the process is removed.

If a backup is aborted for *any* reason the DXULDMLIB process continues to exist as a defunct process. For information on how to determine whether a defunct DXULDMLIB process exists, see [“To View Existing DMAPI Processes” on page 524](#).

Defunct DXULDMLIB processes exist solely as process table entries and use no system resources. These processes normally do not cause any problems and are removed when the system is rebooted. They can be removed manually. For information on this, see [“To Remove a Defunct DXULDMLIB Process” on page 524](#).

## ▼ To View Existing DMAPI Processes

To view all DMAPI processes on a DXDM host system:

1. Log in as root on the DXDM host system.
2. Create the DXDM environment.

Use the correct command and file for the current shell:

- C shell (csh):

```
source /opt/dxuldm/etc/dxuldm.login
```

where `/opt/dxuldm` represents the full path to the DXDM installation directory.

- Korn shell (ksh) or Bourne shell (sh):

```
. /opt/dxuldm/etc/dxuldm.profile
```

where `/opt/dxuldm` represents the full path to the DXDM installation directory.

3. Enter the `prtdmsession` command:

```
prtdmsession
```

The `prtdmsession` command displays a listing for each existing DMAPI session.



---

**Caution** – “[To Remove a Defunct DXULDMLIB Process](#)” on page 524, describes how to use the `delmdsession` command to remove defunct DXULDMLIB sessions. Caution must be exercised with this command to avoid the accidental removal of active DMAPI sessions. If active sessions are removed, DXDM processes terminate and filesystem activity is blocked until those processes are restarted.

---

## ▼ To Remove a Defunct DXULDMLIB Process

To remove a defunct DXULDMLIB process:

1. Log in as root on the DXDM host system.
2. Create the DXDM environment, see [Step 2](#) in “[To View Existing DMAPI Processes](#)” on page 524.
3. Enter the `prtdmsession` command.

Find the listing which has DXULDMLIB in the Session Info field and note the integer in the listing's Session ID field.

4. Enter the `deldmsession` command:

```
deldmsession -s DXULDMLIB -i session-id -n
```

where *session-id* is the integer from the defunct DXULDMLIB process's Session ID field.

## Sun StorEdge EBS Archive Feature

DXDM filesystems are not supported by the Sun StorEdge EBS Archive feature. For information on using this feature with standard filesystems, see [Chapter 8](#).

---

# Recovery of DiskXtender Data Manager Filesystems

Sun StorEdge EBS recovery of a DXDM filesystem consists of restoring the data from a client/save set and synchronizing the filesystem's DMAPI information with the data on the storage target. These topics are described in the following sections:

- [“Initiating a Recovery” on page 525](#)
- [“Filesystem Synchronization” on page 527](#)

## Initiating a Recovery

Recovery of backed-up data from a DXDM filesystem is performed using the same procedures that are used for a standard filesystem. For information on these procedures, see [“Data Recovery” on page 287](#). A DXDM filesystem recovery can consist of individual files, directories or an entire filesystem.



---

**Caution** – With DXDM filesystems, file data must be recovered into the same filesystem from which it is backed up. Attempting to recover backed up data into a new filesystem causes:

- The unavailability of all data on the storage target.
- The DXDM processes to stop responding.

This requirement is based on the nature of the DMAPI metadata that DXDM filesystems use to locate and access data on the storage target. The DMAPI metadata uses full pathnames. If the filesystem mount point changes because of a recovery into a new filesystem, the pathname information in the metadata becomes inaccurate.

---

## Recovered Files List

As part of the recovery process each restored file's pathname is listed in a file on the DXDM host system. DXDM uses this information to synchronize its metadata, see [“Filesystem Synchronization” on page 527](#).

When a recovery is complete all data that is actually in the filesystem is available without waiting for the synchronization process. Data on the storage target becomes available after synchronization. The following types of data require synchronization:

- Migrated and purged file data.
- Purged and deleted file data, see [“Restoring Deleted Files and Previous File Versions” on page 526](#).

## Restoring Deleted Files and Previous File Versions

DXDM filesystems retain file versions. This allows administrators to restore specific versions of files. In a DXDM filesystem a file is migrated when it is first placed into the filesystem and again after each change. Each of these migrations causes a new version of the file to be created on the storage target.

### ▼ To Restore a File Version

To restore a previous version of a file, whether or not the file still exists in the filesystem:

1. Recover the file from the client/save set that backed up the desired version. For information on this procedure, see [“Data Recovery” on page 287](#).
2. Complete the synchronization process, see [“Filesystem Synchronization” on page 527](#).

To expedite the synchronization of a particular file complete the procedure described in [“Manual Synchronization” on page 528](#).



The synchronization process restores the data for the version of the file that is recovered. It will also restore data for a file which has been inadvertently deleted. At the conclusion of the synchronization process the recovered file version or deleted file is available.

## Filesystem Synchronization

As part of a recovery of one or more files into a DXDM filesystem Sun StorEdge EBS recover creates a file containing a list of each recovered file. This list file has a pathname with the following format:

```
/opt/dxuldm/adm/recdir/rec.date.pid
```

where:

- `/opt/dxuldm` is the installation directory for DXDM.
- `date` is the date of the recovery.
- `pid` is the process ID of the Sun StorEdge EBS client process.

DXDM uses the list file to rebuild the filesystem's DMAPi attributes file. This process synchronizes the filesystem's metadata with the data on the storage target.



---

**Caution** – Do not remove a list file created by Sun StorEdge EBS recover. Automatic file synchronization will not occur if a list file is removed.

Manual synchronization can be conducted without a list file. For more information, see [“To Manually Synchronize a File” on page 528](#).

---

Synchronization is an automatic process which does not require administrative intervention, see [“Automatic Synchronization” on page 528](#). To make a file's data available more quickly, see [“Manual Synchronization” on page 528](#).



---

**Caution** – If Sun StorEdge EBS recover cannot write to the `recdir` directory, for any reason, each recovered file is synchronized as it is recovered. This file-by-file synchronization significantly slows down the recovery process.

The `recdir` directory is created during the installation of DXDM. Sun StorEdge EBS recover will be unable to write to it if the directory is removed or if the partition on which it is mounted becomes full.

---

## Automatic Synchronization

DXDM uses the list file created by Sun StorEdge EBS recover to determine which files require synchronization. Each file on the list is synchronized. If synchronization is not successfully completed, for any reason, DXDM repeatedly retries until all files on the list have been successfully synchronized. After successful synchronization of all files, the list file is removed from the `recdir` directory.

Automatic synchronization is performed by the script `dxuldmcronscript` which is invoked by root's crontab every 10 minutes. The cron job to invoke `dxuldmcronscript` is placed in root's crontab when DXDM is installed.

## Manual Synchronization

To make a file available before automatic synchronization is complete use manual synchronization. This uses the DXDM command line utility `dxuldmrecover`.

### ▼ *To Manually Synchronize a File*

To manually synchronize a file:

1. Log in as root on the DXDM host system.
2. Create the DXDM environment, see [Step 2](#) in “[To View Existing DMAPI Processes](#)” on page 524.
3. Run the `dxuldmrecover` utility:

```
dxuldmrecover -p filepath
```

where *filepath* is the full path to the file that is being synchronized.

The `dxuldmrecover` utility takes the full path of a file as its argument. It can be invoked from a script to manually synchronize more than one file.

When synchronization is complete the following message appears:

```
Successfully recovered file filepath.
```

---

**Note** – An entire DXDM filesystem can be manually synchronized using the `dxuldmrecoverfs` utility. This utility examines and synchronizes every file in the filesystem. For more information on this utility, refer to the *DiskXtender Data Manager, Release 2.6, Installation and Administrator's Guide*.

---



## Notes and Troubleshooting Tips

---

This appendix contains troubleshooting tips and workarounds to help you use Sun StorEdge Enterprise Backup Software software.

The main sections in this appendix are:

- [“Before You Call Technical Support” on page 531](#)
- [“Sun StorEdge EBS Backup and Recovery” on page 534](#)
- [“Client/Server Communications” on page 550](#)
- [“Sun StorEdge EBS Archive and Retrieve” on page 559](#)
- [“Storage Nodes” on page 561](#)
- [“Devices and Autochangers” on page 562](#)
- [“Diagnostic Tools” on page 575](#)
- [“Sun StorEdge EBS Locale and Code Set Support” on page 577](#)
- [“Resource Database Notes” on page 578](#)

---

## Before You Call Technical Support

If the solutions in this appendix do not solve the problem, be prepared to provide the following information when you call Sun Technical Support:

- The software release of the Sun StorEdge EBS component.
- The version of operating system that you are running. You can determine this with the **uname -a** command.
- The hardware configuration.
- Information on your devices and other SCSI IDs. Use the **/usr/sbin/nsr/inquire** command as root to obtain the required information.

- 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. You can determine this from the output of **pkginfo -x SUNWebsd**.
- You should also be able to relate the following:
  - How to reproduce the problem
  - The exact error messages
  - How many 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

You can use either the Sun StorEdge EBS Administrator program or the character-based `nsradmin` interface to determine the version of Sun StorEdge EBS software running on clients.

---

**Note** – The version of Sun StorEdge EBS software running on the server appears in the Server Information pane of the main window of the Sun StorEdge EBS Administrator program.

---

### ▼ To Determine the Client Software Version Using the Sun StorEdge EBS Administrator Program

To determine the client software version using the Sun StorEdge EBS Administrator program:

1. In the Sun StorEdge EBS Administrator program, select Client Setup from the Clients menu.
2. Select the appropriate client.
3. Scroll to the Sun StorEdge EBS Version attribute. This is a read-only attribute that lists the version of the Sun StorEdge EBS software running on the client.

### ▼ To Determine the Client Software Version Using `nsradmin`

To determine the client software version using `nsradmin`:

1. At the system prompt enter:

```
nsradmin
```

2. At the nsradmin prompt, enter the following:

```
nsradmin> show name; NetWorker version

nsradmin> print type: NSR client
```

This displays the version of Sun StorEdge EBS software running on each client.

## The nsr\_support Program

The `nsr_support` program gathers system information, including Sun StorEdge EBS software specific details and general information regarding the computer environment, and generates a user-friendly HTML format report that can be used for troubleshooting Sun StorEdge EBS software configurations and system analysis.

If you contact Technical Support for assistance, you might be asked to run the `nsr_support` program on your system. Depending on the nature of the problem, you might also be asked to supply specific command line parameters when running this program.

For information about the command line parameter options available for the `nsr_support` command, see the `nsr_support` man page, or see the *Legato Command Reference Guide*.

To run the `nsr_support` program in interactive mode:

1. Navigate to the directory where the Sun StorEdge EBS software binary files are installed. For example:
  - `/opt/networker/bin`
  - `/usr/sbin`
2. Enter the following at the command prompt to run the `nsr_support` program and generate an HTML format report:

```
nsrsup --html
```

3. Follow the prompts to finish generating the report.

---

**Note** – It could take some time for the `nsr_support` command to gather all required information.

---

The reports generated by the `nsr_support` program will be written to the client's filesystem. For example:

- `/tmp/nsrsup.pid.out.html`
- `/nsr/tmp/nsrsup.pid.out.html`

---

**Note** – Future generated html format reports will be written to different file names, as identified by the process id for each log.

---

---

## Sun StorEdge EBS Backup and Recovery

This section provides tips, notes, and workarounds for issues regarding Sun StorEdge EBS 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 spawned for each active save stream.

### Checking the Sun StorEdge EBS Daemons

If you have trouble starting Sun StorEdge EBS programs, the daemons might not be running properly. To determine whether the required daemons are running, 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
12235 ? S 0:00 /usr/sbin/nsr/nsrmmd -n 3
12236 ? S 0:00 /usr/sbin/nsr/nsrmmd -n 4
12410 pts/8 S 0:00 grep nsr
```

If the response indicates that the daemons are not present, start the Sun StorEdge EBS daemons with the following command:

```
/etc/init.d/networker start
```

## 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 that invokes the `savegrp` program:

```
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 to within one minute of each other.

## 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. Determine the group that is currently being backed up by looking at the messages displayed in the Sun StorEdge EBS Administrator program.
2. Determine the process ID for the currently running **save** process using *one* of the following commands:

```
ps -ef | grep nsr

ps -ax | grep nsr
```

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 [“To Stop a Group Immediately” on page 258](#).
- Shut down and restart the Sun StorEdge EBS daemons. To do this:
  1. Shut down the Sun StorEdge EBS daemons using the `nsr_shutdown -a` command.
  2. Restart the Sun StorEdge EBS daemons using the `nsrd` command.
- Stop the **save** process for each process ID, using the `kill -9 pid` command.

## Save Set Cutoffs for 5.x Clients

The Save Set Cutoff option is no longer supported in Sun StorEdge EBS 6.0 or later releases. When Sun StorEdge EBS clients earlier than 6.0 back up to a Sun StorEdge EBS 6.0 or later server, the server respects the save set cutoff value. If both the server and the client are Sun StorEdge EBS 6.0 or later releases, then the save set cutoff feature is not present.

## Memory Usage When Browsing Large Save Sets

Browsing or recovering from a large save set, such as a save set with 1 million or more files, may consume all of the host's memory. The recommended workaround if this problem occurs is to perform a saveset recovery instead. For information on save set recover, see [“Save Set Recovery” on page 303](#).

## Licensing-Based Errors with Release Clients Earlier Than 5.0

Clients running Sun StorEdge EBS software earlier than release 5.0 do not have “self-ID” capabilities and are unable to report their OSs to the server at backup time. These clients are all listed under client types earlier than release 5.0 when `nsrlic -v` is run. Since the server does not know what OS the client is running, it uses the `workstation license, even if the client is on a Windows NT or UNIX server. However, such non-ID clients have no effect on which Legato NetWorker ClientPak licenses are required to allow Sun StorEdge EBS clients to function; the appropriate licenses are still required.

Sun StorEdge EBS backups are rejected in the following cases because of licensing enforcement, and an error message is generated:

- If the client type is not allowed by the enabler code, or if a ClientPak enabler is not present. For example, the following error message might be sent for a Solaris server with an unlicensed Windows NT client:

```
RAP error, ClientPak for NT does not exist in the system.
```

- If no more licenses are available for that type of client. For example, you might receive an error message similar to this:

```
RAP error, Too many clients. Maximum is 2
```

## 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” on page 312](#).

## 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 [“Online Indexes” on page 88](#). For more information, refer to the `nsr ls`, `nsrck`, and `nsrim` man pages, or to the *Legato Command Reference Guide*.

## Media Position Errors Encountered When Auto Media Verify Is Enabled

When you enable Auto Media Verify for a pool, the Sun StorEdge EBS server verifies the data written to volumes from the pool during the save. This is done by reading a record of data written to the media and comparing it to the original record. Media is verified after the server finishes writing to the volume, which might occur when a volume becomes full or when the server no longer needs the volume for saving data.

To verify media, `nsrmmd` must reposition the volume to read previously written data. It does not always succeed in the first attempt. These warning messages appear in the message window of the Sun StorEdge EBS Administrator program:

```
media warning: /dev/rmt2.1 moving: fsr 15: I/O error

media emergency: could not position jupiter.007 to file 44, record
16
```

No action is required. The Sun StorEdge EBS server continues to attempt to find the proper position. If the server can find the correct position, media verification succeeds and a successful completion message appears.

```
media info: verification of volume "jupiter.007" valid 30052
succeeded.
```

In this case, ignore the earlier messages because they only indicate that the server had problems finding the desired position on the media.

If media verification fails, perform the following:

- 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.

Next, restart the Sun StorEdge EBS server manually using the following command:

```
/etc/init.d/networker start
```

## 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
```

## Index Recovery to a Different Location Fails

If you attempt to recover indexes to a different directory 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.
```

Do not attempt to recover the indexes to a different directory. After the indexes are recovered to their original location, you can move them to another directory. To move the indexes, log in as root and invoke the following command from within the `/nsr/index` directory:

```
/usr/lib/nsr/uasm -s -i client_index_directory_name | (cd target_directory;
/usr/lib/nsr/uasm -r)
```

## Potential Cause for Client Alias Problems

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.
- It appears that automatic index management, as set up in the browse and retention policies, does not occur. An indication of this is when the filesystem that contains the indexes continuously increases in size.

- 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 for the following situations:

- Machines that have two or more network interfaces
- Sites that mix short and fully qualified hostnames for the same machines; for example, `mars` and `mars.jupiter.com`
- Sites using both (Network Information Services (NIS) and DNS

Edit the Client resource for clients with this problem. Add all network names for this host to the Aliases attribute.



---

**Caution** – Do not include aliases that are shared by other hosts in the Aliases attribute.

---

## Illegal Characters to Avoid in Configurations

When naming label templates, directives, groups, policies, and schedules, the following characters are not allowed:

`/ \ * [ ] ( ) $ ! ^ ' ; ' ~ < > & | { _`

In addition, the following characters cannot be used in label templates because they are reserved for use as field separators:

- Colon (:)
- Hyphen (-)
- Period (.)
- Underscore (\_)

## 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, and you receive the following message:

`please enter record size for this volume ('q' to quit) [xx]`

where the number in the bracket [xx] is the entry from the last query.

The `scanner` command always rewinds the tape and reads the volume label to determine the block size. If the volume label is corrupted or unreadable, a message prompts you to enter the block size (in kilobytes).

Type in the block size; it must be an integer equal to or greater than 32. If you enter an integer that is less than 32, the following message appears:

```
illegal record size (must be an integer >=32)
```

## 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 on a system, the following error message appears:

```
nwrecover: Program not found.
```

To save disk space, Sun StorEdge EBS software delays the creation of the client file index until the first backup is completed. The `nwrecover` program cannot recover data until the client file index has entries for browsing. To avoid the problem, perform a Sun StorEdge EBS backup on the client.



## Client File Index Is Missing Message

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 re-create
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
```

## Recovering Files from an Interrupted Backup

If you terminate a backup by stopping the Sun StorEdge EBS daemons, you cannot recover the files because the media database is not updated when the daemons stop. Consequently, the Sun StorEdge EBS server does not know on which volumes the requested files reside.

## Backup of a New Client Defaults to a 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. If the save set had been backed up, the save set would appear in the media database.

Before you can perform an incremental or level backup on a save set, you must 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.

## 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 [“Directed Recovery” on page 294](#).

## The `savegrp` Command Line Program Cannot Be Stopped from the Sun StorEdge EBS Administrator Program

If you start the `savegrp` command at the command prompt and attempt to stop the backup from the Sun StorEdge EBS Administrator program, the following message appears:

```
Only automatically started groups that are currently running can
be stopped
```

To stop the command, you must first determine the process ID of the `savegrp` process, and then use the UNIX `kill` command to stop the process.

### ▼ To Stop the `savegrp` Command Line Program

To stop the `savegrp` command line program:

1. Determine the `savegrp` process ID using the `ps` command:

```
ps -ef | grep savegrp
```

The process ID appears in the following format:

```
root process_id 4769 0 13:53:10 pts/2 0:00 savegrp myGroup
```

2. Stop the `savegrp` process using the `kill` command:

```
kill process_id
```

## Disk Label Errors

If you have incorrectly configured a nonoptical device as an optical device within the Sun StorEdge EBS server, 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, you might need to enter the printer's name as a hidden attribute in the Group resource.

1. Select Details from the View menu to display hidden attributes.
2. In the Printer attribute of the Group resource, enter the name of the printer where you want the bootstraps to print.

## Server Index Not Forced

If the Sun StorEdge EBS server belongs to a group that is not enabled or if it does not belong to any group, the `savegrp` program does not back it up, because it is not necessary to force a server index backup.

The information to recover server indexes is stored in the media database.



---

**Caution** – If the datazone’s planning has depended on `savegrp` forcing a server index backup that has never been explicitly backed up, note that `savegrp` no longer performs a server index backup for servers not in an active group.

---

## Copy Violation

If you installed Sun StorEdge EBS software on more than one server and used the same Sun StorEdge EBS enabler code for them all, messages similar to the following appear in the save group completion mail:

```
--- 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 daemons on the server where you want the backups to go.

## Xview Errors

If you receive the following error message when you attempt to start the graphical administration interface with the Sun StorEdge EBS Administrator program from a client machine, 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 Sun StorEdge EBS graphical interfaces:

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 at the prompt:

```
setenv DISPLAY client_name:0.0
```

For command shells other than `CSH` enter:

```
DISPLAY=client_name:0.0

export DISPLAY
```

## Limited Frame Buffer Capabilities May Cause `nwadmin` to Display Incorrectly

If you are using a system with limited frame buffer capabilities (for example, a 8-bit frame buffer), the Sun StorEdge EBS Administrator program may sometimes start with both a white foreground and a white background. This might happen if another application is open and is using most of the available colors.

To resolve this problem:

1. Log in as root:

```
% su root
```

2. Change to the `/usr/lib/X11/app-defaults` directory:

```
cd /usr/lib/X11/app-defaults
```

3. Change the permissions of the read-only `Networker` file so that root has write access:

```
chmod 644 Networker
```

4. Open the `Networker` file in a text editor

5. Search for all lines in this file that contain the text `*background`, for example:

```
Networker*background #5a757b
```

6. Change the color mapping to black, for example:

```
Networker*background black
```

Repeat for all `*background` entries in the file.

7. Save the changes to the `Networker` file and exit.

8. Set the `$XAPPLRESDIR` environment variable to `/usr/lib/X11/app-defaults`; for example:

- ksh:

```
export XAPPLRESDIR=/usr/lib/X11/app-defaults
```

- csh:

```
setenv XAPPLRESDIR /usr/lib/X11/app-defaults
```

This will change the background color of the Sun StorEdge EBS Administrator to black.

## Converting Sparse Files to Fully Allocated Files

The Sun StorEdge EBS server determines that files are sparse or holey by comparing the allocated blocks with the byte size. Use the `ls -ls` command to compare blocks. If the allocated blocks do not account for the size of the file, the file is considered to be sparse and is saved using an algorithm that replaces long strings of zeroes with “holes” in the recovered file.

Some files that were not sparse when saved might be recovered as sparse. Oracle databases are particularly susceptible to this problem because they are zero-filled, fully allocated files and are not sparse.

The workaround for this situation is to 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** – If you choose this workaround, you must have enough free disk space to accommodate a duplicate of each sparse file you choose to copy.

---

## Cannot Manage Sun StorEdge EBS User Groups from Sun StorEdge EBS Release 6.x or Earlier

If you use a version of the Sun StorEdge EBS Administrator program from Sun StorEdge EBS release 6.x or earlier to connect to a Sun StorEdge EBS server running release 7.0 or later, you will be unable to manage the User Groups resource.

## 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 referenced by the `-N` option must match the case of the save set name as it was 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:\'
```

---

## Client / Server Communications

Many of the problems that Sun StorEdge EBS users report when they set up and configure Sun StorEdge EBS software concern network communications. This section contains a procedure for testing the communications in a network.

The information provided in this section is for UNIX platforms only. If your Sun StorEdge EBS setup includes other platforms, refer to *Legato Technical Bulletin 299: IP Naming in Heterogeneous Environments (UNIX | NT | NetWare | Windows 95)*, available on the Legato web site at [www.legato.com](http://www.legato.com).

### ▼ To Troubleshoot IP Errors

If any of the following error conditions occurs, you probably have 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 or daemons failing to start
- Sun StorEdge EBS services or daemons 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 e-mail or fax the exact steps and error message text directly to Sun.
2. Set up a host table for each Sun StorEdge EBS client and Sun StorEdge EBS server. For more information, see [“To Set Up a Host Table” on page 551](#).
3. Disable other name servers to simplify testing. For more information, see [“Troubleshooting Hostname Resolution Problems” on page 552](#).
4. Use `ping` to establish basic connectivity. For more information, see [“To Use ping to Verify Network Connections” on page 553](#).
5. Use `rpcinfo` to verify that sessions can be established and that portmapping is correct. See [“To Use rpcinfo to Verify That Sessions Can Be Established” on page 554](#) for more information.

## ▼ To Set Up a Host Table

It is recommended that you 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, you can enable the named server you are using.

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 [“To Use ping to Verify Network Connections” on page 553](#) 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 unremarked 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.

## Troubleshooting Hostname Resolution Problems

To simplify the troubleshooting of hostname resolution problems, try disabling services like DNS, DHCP, and Windows Internet Naming Service (WINS). If you have hostname resolution problems, first configure the host table for each computer, and then test backups.

Common problems encountered with DNS, DHCP, and WINS services include:

- The DNS is not configured with a reverse lookup table.
- The clients are configured with the wrong IP addresses for DNS or WINS servers.
- The DHCP services do not properly update the WINS server with new addresses.

You do not need to disable DNS for the entire network, just for the initial setup of the Sun StorEdge EBS clients and the Sun StorEdge EBS server you want to test. Only disable the ability of a client to obtain IP naming information from a DNS server. Typically, you do not need to disable the DNS server itself.

To disable the DNS server on most UNIX platforms, rename the file `/etc/resolv.conf` and reboot.

Instead of renaming `resolv.conf`, you can set up the IP name search order so that the host table is searched before DNS.

## ▼ 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
```

## ▼ 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

## ▼ To Use `rpcinfo` to Verify That Sessions Can Be Established

If `ping` is successful and backup problems still exist, you can also test with `rpcinfo`. Because Sun StorEdge EBS software relies heavily on mapping of ports, use `rpcinfo` to test the operation of the portmapper. Using `ping` tests the connection up to the network layer in the Open System Interconnection (OSI) model; `rpcinfo` checks for communication up to the session layer.

Use the same tests with `rpcinfo` as with `ping`.

For `rpcinfo` to be used successfully, the machine whose hostname you enter on the command line must have a portmapper running. In most cases, Sun portmappers are compatible with fully functional portmappers from other vendors (this is called a third-party portmapper). If you are using a product that provides its own portmapper, we recommend not loading the third-party portmapper until you have verified that the Sun StorEdge EBS software works with the rest of the environment. This process lets you test portmapper compatibility without adding other unknowns.

The `rpcbind` daemon 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*, just like for `ping`.

View other `rpcinfo` command line options by entering `rpcinfo` at the line. Notes on the `rpcinfo` command and its error messages are available in the UNIX man page for `rpcinfo`. 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 as shown in the following example:

```
rpcinfo for mars

program vers proto port
100000 2 tcp 111 portmapper
100000 2 udp 111 portmapper
390103 2 tcp 760
390109 2 tcp 760
390110 1 tcp 760
390103 2 udp 764
390109 2 udp 764
390110 1 udp 764
390113 1 tcp 7937
390105 5 tcp 821
390107 4 tcp 819
390107 5 tcp 819
390104 105 tcp 822
```

## Replacement for Functional Use of portmapper and rpcbind

The `nsrexecd` process now combines the functionality of `portmapper` and `rpcbind`. The `nsrexecd` process spawns a `nsrexecd` subprocess that handles the `portmapper` functions. As a result, you see more than one instance of the

nsrexecd process. In the following example, the instance of nsrexecd running under the process ID (PID) 20415 is the main process and the instance of nsrexecd running under the PID 20416 is the subprocess:

```
jupiter# ps -ef | grep nsr

root 20429 204200 11:23:59 ? 0:00 /usr/sbin/nsr/nsrmmd -n 1
root 20415 1 0 11:23:49 ? 0:01 /usr/sbin/nsr/nsrexecd -s jupiter
root 20416 204150 11:23:49 ? 0:01 /usr/sbin/nsr/nsrexecd -s
jupiter
root 20428 204200 11:23:57 ? 0:00 /usr/sbin/nsr/nsrindexd
root 20420 1 0 11:23:53 ? 0:02 /usr/sbin/nsr/nsrd
root 20446 203490 11:27:22 pts/6 0:00 grep nsr
root 20427 204200 11:23:55 ? 0:00 /usr/sbin/nsr/nsrmmdbd
```

This change allows the Sun StorEdge EBS processes to function in a more secure manner that does not utilize ports or protocols that other programs commonly use.

## ▼ To Verify Firmware for Switches and Routers

If you are using switches or routers from any vendor, 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.

## Naming Requirements

The Sun StorEdge EBS client uses the `servers` file in the `/nsr/res` subdirectory to determine whether a Sun StorEdge EBS server is authorized to back up the client's data. If you do not have the `servers` file, you can create it in `/nsr/res` using any text editor.

Make sure the `servers` file on a client contains both the short name and the long name of the server you want to use to back up that client's data. For example, the `servers` file on a Sun StorEdge EBS client would 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.

## 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 inside of long-lived processes known as daemons.

For clients to find these daemons, register the daemons with a registration service. When the daemons start, they register themselves with the registration service provided by the portmapper.

Sun StorEdge EBS servers provide a backup and recovery service. The servers receive data from clients, store the data on backup media, and retrieve it on demand. If the Sun StorEdge EBS daemons 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 daemons `nsrd`, `nsrexecd`, `nsrindexd`, `nsrmmmd`, and `nsrmmdbd` might not be running.

To restart the daemons, log in as root and start the Sun StorEdge EBS daemons with the following command:

```
/etc/init.d/networker start
```

## Saving Remote Filesystems

The following error messages might appear in the save group completion mail when a backup for a remote client fails:

```
Host hostname cannot request command execution

hostname: Permission denied
```

The first message means that the `nsrexecd` daemon on the client is not configured to allow the server to back up its files. The second message means that the `nsrexecd` daemon is not currently running on the client.

To resolve these problems, make sure that the `nsrexecd` daemon is running on the client, and that the server's hostname is listed in the boot-time file. The boot-time file is automatically generated before the installation script is completed, and takes your responses to the query for the names of all the servers, in order of precedence, that can contact a client for backups. [TABLE A-1](#) lists the location for the boot-time file. For detailed information about the `nsrexecd` daemon, refer to the `nsrexecd(1m)` man page.

**TABLE A-1** Boot-Time File Locations

| Operating System | Boot-Time File                     |
|------------------|------------------------------------|
| Solaris          | <code>/etc/init.d/networker</code> |
| SunOS 4.1.x      | <code>/etc/rc.local</code>         |

## 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. You can add or remove user names 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, you must 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 only recover files for which they have read permission, relative to the file mode and ownership 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.



---

# Sun StorEdge EBS Archive and Retrieve

This section explains how to troubleshoot problems you might encounter during archiving and retrieval with the Sun StorEdge EBS Archive Module.

## Remote Archive Request from Server Fails

If you cannot perform a remote archive request of a computer from the Sun StorEdge EBS server, the archive client's user name (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. Granting administrator privileges creates a potential security issue, because 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.

If you want save sets to appear separately during retrieval, 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, you can use `mminfo` to locate the archive clone with the same save set ID (ssid) as the archive you want.

## Wrong Archive Pool Is Selected

If you create multiple archive pools, the one selected for archive is not the default archive pool. When you create multiple archive pools, 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. To avoid this problem, do not create two archive requests with the same name; the newer one is not executed.

## Command Line Archive Does Not Start Immediately

If you run `nsrarchive` from the command line, the archive does not start immediately after you enter the annotation and then [Ctrl]+[D] to start the archive. Wait a short time; there is a delay before 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 stat" error. Filenames that include carriage returns are not allowed.

---

# Storage Nodes

This section discusses errors and solutions associated with storage nodes.

## Storage Node Affinity Errors

You may have a storage node affinity problem 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 listed in the Storage Nodes attribute.
- The devices do not have volumes that match the pool required by the backup request.
- All devices are set to read-only.

You must fix the problem and restart the backup. To fix the problem, 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 enabled devices meeting 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. For more information, refer to the online help.

## 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 nsrmmmd control timeout attribute instead
```

If you receive this message:

1. Shut down the Sun StorEdge EBS daemons.
2. Remove the environment setting for NSR\_MMDCONTROL.
3. Restart the Sun StorEdge EBS daemons.
4. Start the Sun StorEdge EBS Administrator program.
5. Adjust the value of the Nsrmmmd Control Timeout attribute to the value previously assigned to the NSR\_MMDCONTROL variable, or one that best meets the current requirements.

---

## Devices and Autochangers

This section discusses features, problems, and solutions associated with using an autochanger with Sun StorEdge EBS software.

### Additional Attributes in the Jukebox Resource

The Jukebox resource now contains attributes that provide a detailed view of options that the nsrjb program uses. These are hidden attributes; select Details from the View menu in the Jukebox window to display them.

[TABLE A-2](#) lists additional Jukebox attributes.

**TABLE A-2** Additional Jukebox Attributes

| Attribute        | Description                                                                                                                                                                                                                                                                                                                                                   |
|------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Nsrjb Release    | Indicates the version of the nsrjb program that is installed on the computer that uses the selected Jukebox resource. The version displayed for Nsrjb Release might differ from the version displayed for the Sun StorEdge EBS software on the server, depending on whether you updated the storage nodes to the same Sun StorEdge EBS release as the server. |
| Jukebox Features | Displays the features that the autochanger supports for use by the nsrjb program. If you are installing a software update, the values that were once assigned to the Jukebox Options attribute migrate to the equivalent selections for the new Jukebox Features attribute.                                                                                   |

**TABLE A-2** Additional Jukebox Attributes

| Attribute       | Description                                                                                                                                                                     |
|-----------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Eject Sleep     | Defines the number of seconds an autochanger remains idle (“sleep”) after an eject operation is completed.                                                                      |
| Cleaning Delay  | Defines the number of seconds an autochanger remains idle before attempting to unload a cleaning cartridge.                                                                     |
| Unload Sleep    | Defines the number of seconds an autochanger remains idle after an unload operation is completed.                                                                               |
| Load Sleep      | Defines the number of seconds a jukebox remains idle after a load operation is completed.                                                                                       |
| Deposit Timeout | Defines the number of seconds a jukebox waits for a deposit in the mail slot before the <code>nsrjb</code> program abandons further attempts to complete the deposit operation. |



**Caution** – The timed settings are set to the optimal settings for optimal performance. Do not change these settings unless advised to do so by a Technical Support representative.

## Maintenance Commands

Sun StorEdge EBS device driver software provides the following maintenance commands for diagnosing problems on tape devices and autochangers. For more detailed information regarding these commands, refer to the *Legato Command Reference Guide* or the corresponding man pages.

[TABLE A-3](#) lists autochanger maintenance commands.

**TABLE A-3** Autochanger Maintenance Commands (1 of 2)

| Command               | Description                                        |
|-----------------------|----------------------------------------------------|
| <code>lusbinfo</code> | Prints out SCSI information                        |
| <code>lusdebug</code> | Sets the library debugging level                   |
| <code>lrescan</code>  | Rescans for devices                                |
| <code>lreset</code>   | Resets the SCSI bus                                |
| <code>changers</code> | Lists the SCSI autochangers attached to the system |

**TABLE A-3** Autochanger Maintenance Commands (2 of 2)

| Command  | Description                                                                                                                                                                                                                                                                                          |
|----------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| hpflip   | Controls whether an HP optical disk drive reports that it is an optical device or a disk (direct access) device. Many systems do not have drivers for optical devices; <b>hpflip</b> allows you to toggle the device between optical and direct access without having to install additional drivers. |
| ielem    | Initializes the element status                                                                                                                                                                                                                                                                       |
| inquire  | Lists the devices available                                                                                                                                                                                                                                                                          |
| ldunld   | Loads or unloads a tape device                                                                                                                                                                                                                                                                       |
| msense   | Retrieves mode sense data                                                                                                                                                                                                                                                                            |
| pmode    | Prints mode sense data                                                                                                                                                                                                                                                                               |
| relem    | Reads the element status                                                                                                                                                                                                                                                                             |
| tur      | Tests whether the unit is ready                                                                                                                                                                                                                                                                      |
| writebuf | Writes a device buffer                                                                                                                                                                                                                                                                               |
| sjielm   | Tests the standard jukebox interface (SJI) SJIELEM command                                                                                                                                                                                                                                           |
| sjiinq   | Tests the SJI SJINQ command                                                                                                                                                                                                                                                                          |
| sjirdp   | Tests the SJI SJIRDP command                                                                                                                                                                                                                                                                         |
| sjirdtag | Tests the SJI SJIRTAG command                                                                                                                                                                                                                                                                        |
| sjirelem | Tests the SJI SJIRELEM command                                                                                                                                                                                                                                                                       |
| sjirjc   | Tests the jukebox                                                                                                                                                                                                                                                                                    |

## ▼ To Test the Device Driver Installation

After you install the Sun StorEdge EBS device driver software, use the **lusdebug** program to verify the server connection and the **jbexercise** program to test the autochanger by using the following commands:

```
lusdebug control_port 0

jbexercise -c control_port -m model
```

where *control\_port* is the value of the control port assigned to the autochanger (for example, *scsidev@0.6.0*).

If these commands fail or if you receive error messages, see the following sections for information on the possible causes and solutions.

## The `lusdebug` Command Fails

If the `lusdebug` command fails, review these suggestions to identify the potential problems and their solutions:

- Issue the `sjinq` command as root, and provide the `control_port` as an argument. A message similar to the following should appear:

```
scsidev@0.6.0:<EXABYTE EXB-10i EXB-10i >
```

Verify that the information in the message is correct.

Incorrect vendor and model names indicates the wrong SCSI ID was supplied for the device ID during the driver installation. The installation script asks for the SCSI ID of the robot mechanism, not the tape drive.

Uninstall the device driver and reinstall it, and supply the correct address for the autochanger (robotic arm). Make sure that each device on the SCSI bus has a different SCSI ID address.

- Inspect the following items to verify that the autochanger is properly connected:
  - a. Make sure all the connectors on the SCSI bus are firmly connected.
  - b. Make sure none of the SCSI cables are defective.
  - c. Verify that the SCSI bus is properly terminated and is within the length specified by ANSI SCSI-II specifications (ANSI X3.131-1994).

Both ends of the SCSI bus must be terminated with the appropriate resistors to be properly terminated. Single-ended SCSI buses are 220 ohms to +5 VDC, 330 ohms to ground. Differential terminators have a 122-ohm characteristic impedance (-5 VDC to +5 VDC). The ends of the SCSI bus are considered to be the last SCSI device at either end of the bus, where both peripheral devices and systems are considered as peer SCSI devices.

Additional termination (termination placed on devices not at either end of the SCSI bus) is ill-advised. Additional termination causes the hardware bus drivers on each device on the bus to work harder (for example, out of the range of their nominal specification) to affect signal transitions. As a result, they might not be able to meet the timing requirements for some signal transitions.

- d. SCSI bus length limitations affect the quality of the signals, thus increasing the likelihood of transmission errors on the bus. For single-ended SCSI buses (the most prevalent), the length is 6 meters, unless Fast SCSI devices are attached and in use, in which case the length limit is 3 meters. This length includes the

length of the bus within a device, as well as the length of external cables. A reasonable rule of thumb for internal length is to assume 1 meter of internal bus length for the computer chassis and about 0.25 meters per device for external peripheral boxes.

Differential option SCSI buses can be much longer (due to the electrical differences from single-ended). Allow for a maximum of 25 meters. Never mix differential and single-ended devices.

- Check to see whether an old autochanger driver is still installed. This can be the autochanger driver shipped with earlier versions of Sun StorEdge EBS software, or release 1.1 or earlier of the Parity driver, which only supported SCSI bus 0.

Uninstall the driver according to the instructions shipped with the old driver, and then reinstall the latest version. Special instructions on how to remove the autochanger driver are available in *Legato Technical Bulletin 086: De-installing the AAP or Parity Driver (UNIX)*, available on the Legato web site.

- Check the SCSI IDs on all devices attached to the same bus; make sure that none are the same. If two devices have the same target ID, the following symptoms occur:
  - SCSI bus reset errors appear in system log files
  - The machine does not boot
  - The `probe-scsi` boot prompt command on SPARC systems hangs
- If the sensor that verifies whether the tape drive door is open is out of place, refer to the instructions provided with the autochanger hardware to determine the problem, or contact the hardware vendor.
- If the autochanger is in sequential mode, change the setting to random mode.

If none of these suggestions resolves the problem, contact Sun Technical Support. Provide the information described in [“Before You Call Technical Support” on page 531](#) and the captured output of the `jbexercise`, `sjiinq`, and `sjirjc` programs. For more information, refer to the *Legato Command Reference Guide* or the `jbexercise`, `sjiinq`, and `sjirjc` man pages.

## The `jbexercise` Command Fails

If the `jbexercise` command fails, review the following list of suggestions to identify potential problems and their solutions:

- The `jbexercise` program prompts you for a nonrewinding device name (for example, on Solaris, `/dev/rmt/0mbn`). Verify that you have supplied the correct device pathname for the tape drive. The device name must belong to a tape drive in the autochanger, not the autochanger itself.



If a nonrewinding device name was not entered, the following error message appears:

```
device not ready
```

- Make sure that the tape drive for which you enter the pathname works. Insert a volume into the drive and perform the following tests:
  - a. Use the `tar` command to copy a small file to the volume.
  - b. Verify more extensive operations by issuing the `tapeexercise` command.

If these tests fail, the tape drive is not functioning. For further information on how to configure the tape drive, contact the hardware vendor.

If none of these suggestions resolves the problem, contact Sun Technical Support. Provide the information described in [“Before You Call Technical Support” on page 531](#) and the captured output of the `jbexercise`, `sjinq`, and `sjirjc` programs. For more information, refer to the *Legato Command Reference Guide* or the `jbexercise`, `sjinq`, and `sjirjc` man pages.

## Autodetected SCSI Jukebox Option Causes Server to Hang

If you install an autodetected SCSI jukebox using `jbconfig`, and the server hangs:

1. Select the `jbconfig` option that installs an SJI jukebox. A list of jukeboxes appears.
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.
```

## X11 Libraries Required for jbconfig

To use the `jbconfig` program on UNIX, the X11 graphical libraries must be installed. If you attempt to run the `jbconfig` program on a computer that does not have the X11 graphical libraries installed, the following error message appears:

```
Symbol S..... in use libvgalaxy.a is not defined
Cannot load libvgalaxy.a
System error, can not run a file.
```

Before you can recover data using an autochanger, you must add and configure the autochanger using the `jbconfig` command.

On an HP-UX computer, ensure that the `/nsr/jukebox` directory and its contents are present before attempting to configure an autochanger. If the directory is not present or it is empty, copy the contents of `/opt/networker/lib` to `/nsr/jukebox` before attempting to configure an autochanger.

## Autochanger Inventory Problems

Any of the following situations can 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 software cannot use the autochanger.

To make the autochanger usable again:

1. Verify that the media cartridge is correctly installed in the autochanger and that the autochanger door is closed.
2. Log in as root on the Sun StorEdge EBS server.
3. Reset the autochanger by entering the following command at the prompt:

```
nsrjb -Hv
```

4. Perform an inventory by entering the following command at the prompt:

```
nsrjb -Iv
```

After the inventory operation is finished, the Sun StorEdge EBS server can once again use the autochanger.

For complete information on the use of the `nsrjb` command, refer to the `nsrjb(8)` man page or see [Chapter 10](#).

## Destination Component Full Messages

If a manual operation is performed on an autochanger, for example, physically unloading the tape drive by means of the buttons on the autochanger rather than using the Sun StorEdge EBS server to unmount the volume, the following error message may appear:

```
Destination component full
```

Manual autochanger operations can cause the server to lose track of the status of the media in the autochanger. To resolve the problem, use the Sun StorEdge EBS command `nsrjb -H` to reset the autochanger.

## Tapes Do Not Fill to Capacity

You might encounter situations where the Sun StorEdge EBS software does not fill tapes 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 enable the server to use the tape capacity to its fullest, select the highest density device driver appropriate for the device. When a tape is labeled, the server writes to it at the highest density the device supports.

There are several reasons why the server appears to fill tapes prematurely:

- Write errors occur during a backup.

Most tape drives try to read after a write operation to verify that the tape was written correctly, and retry if it was not. A write error indicates either an end-of-tape or read error. At any tape error, the Sun StorEdge EBS server marks the tape full.

To prevent tape write errors, clean the tape drive regularly and use only data-quality tapes. If cleaning the drive does not seem to help, make sure:

- The device driver is properly configured
- Any necessary switch settings on the tape drive are set to the manufacturer's specifications
- All cabling is secure
- Other potential SCSI problems have been addressed

■ Sun StorEdge EBS filemarks take up 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. On some drives, filemarks can consume several megabytes. The number of filemarks the server writes to tape is a function of how many save sets are on the tape. Many small save sets require more filemarks than a few larger ones.

■ Tape capacities vary from tape to tape.

Tape capacities are not constant from tape to tape. 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 kind 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.

Be sure to verify tape lengths. A 120-meter DAT tape holds more data than a 90-meter DAT tape, and without examining the printed information on the tape cassette carefully, the two tapes can appear identical.

For more detailed information, refer to *Legato Technical Bulletin 176: NetWorker and Tape Capacity Discrepancies (UNIX | NetWare)*, available on the Legato web site at [www.legato.com](http://www.legato.com).

For Solaris, if the tape devices are not directly supported by Sun Microsystems, you must re-create the entries in the *st.conf* file. If you need assistance with this, contact Sun Technical Support.

## Server Cannot Access Autochanger Control Port

The control port controls the autochanger loading mechanism. For instructions on how to verify whether the control port is properly connected, refer to the autochanger's hardware installation manual. If you cannot determine whether the control port is working, contact the autochanger vendor for assistance.

## Nonrewinding Device Requirement

You must 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. You will be unable to recover the previous backup data at a later date.

## ATL 2640 Libraries

If a computer connected to an ATL 2640 library is upgraded from a release of Sun StorEdge EBS software earlier than 5.5, perform the following procedure to allow the ATL 2640 library to be used:

1. Set the *polled\_import* and *stacker\_export* jukebox attributes using the Sun StorEdge EBS Administrator program.
  - a. Select Jukeboxes from the Media menu.
  - b. In the Jukeboxes dialog box, select Details from the View menu to display hidden attributes.
  - c. In the Jukeboxes dialog box, select the *Polled\_import* and *Stacker\_export* options for the Jukebox Features attribute.
  - d. Click Apply.
2. Configure the ATL library to enable User Op IOD (inport/outport device). This must be done to allow the IOD door to open to insert tapes.
  - a. Connect a DOS/Windows computer to the library DIAG (diagnostic) serial port.
  - b. Load the ATL diagnostic program *tapelib.exe* on the computer.
  - c. Select User Op IOD from the Config menu, and then select Enable.

In addition, verify that the library is configured with an IOD on the correct side of the device (it must match the physical configuration of the library). Also verify that the passthrough is enabled.

---

**Note** – Do not specify any port numbers on the `nsrjb` command line. In addition, starting with Sun StorEdge EBS release 5.5, the numbering scheme for the import/export ports changed from zero-based (0,1,2,3) to 1-based (1,2,3,4) to match the scheme used for all other items (slots and drives).

---

## Depositing Volumes to Slots

Instead of performing an inventory of the autochanger after depositing volumes, `nsrjb` enables you to specify a range of volume names already contained within the media database.

To deposit a labeled volume that is already contained within the media database into the autochanger, specify only the names of the volumes to be deposited and not the associated port or slot number. The Sun StorEdge EBS server deposits volumes starting from the first port to the first slot that is not empty, for example:

```
nsrjb volume_name1 volume_name2 volume_name3
```

If you do not want to inventory an autochanger after depositing labeled volumes that are already listed in the Sun StorEdge EBS media database, make sure to:

- Specify a range of empty slots for the `-S slots` option. The range specified should begin with the first volume in the range of ports.
- Specify a range of nonempty ports for the `-P ports` option, starting from the first port. If the range of volumes to be deposited does not start from the first port, then the range of ports must be specified using the `-P` option. For example:

```
nsrjb -S slots -P ports volume_names
```

## Depositing Volumes to Slots on a Solaris Server or Storage Node

When you use the `nsrjb -S` command to deposit a range of volumes to a range of slots in an autochanger attached to a Solaris server or storage node, the deposit does not occur if the volume names are also specified with the command. To avoid the problem, do not specify the volume names at the command line.

## Using the -o Command Option for nsrjb

If you use the `-o` option twice in the same `nsrjb` command line, the first `-o` option stated is ignored. For example, in the following command, the `-o notreadonly` task is not executed:

```
nsrjb -o notreadonly -o notfull -S 2
```

To perform multiple `nsrjb` tasks with the `-o` option, enter the commands separately. For example, the previous commands must be entered as follows:

```
nsrjb -Y -o notreadonly -S n

nsrjb -Y -o notfull -S n
```

## Scanner Command Behaves Differently with Adv\_file Type Device

The scanner command behaves differently when you are using this command with the advanced file type device.

When both primary and `_AF_readonly adv_file` type devices are unmounted, running the following command results in `_AF_readonly` device being mounted:

```
scanner -m -S ssid primary_device_name
```

This is the expected behavior of the scanner command with `adv_file` type device.

# Misleading Error Message Displayed If Server and Remote Storage Node Clocks Are Not Synchronized

The following error message might appear if you are cloning to a remote storage node device, despite a successful cloning process:

```
Starting cloning operation...from storage node <storage-node>
nsrclone: Failed to clone any save sets
```

To prevent this misleading error message from being displayed, ensure the remote storage node and server clocks are synchronized. You can verify that the clone worked using the `mminfo` command, for example:

```
mminfo -s server -c client -S
```

---

**Note** – To prevent misleading error messages from appearing during any Sun StorEdge EBS operations, ensure all clocks on your system configuration are synchronized.

---

## Sleep Times Required for TZ89 Drive Types

If you are unloading a TZ89 drive and receive the error, you must modify 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 retrying the unload attempt.



4. Use the following settings for the sleep time attributes:

- Eject Sleep: 18 secs
- Unload Sleep: 30 secs
- Load Sleep: 30 secs

For information about setting the sleep attributes, see [“Adjusting Sleep Periods for Autochanger Tasks” on page 375](#).

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 you enable the CDI feature while using an NDMP tape device or disk 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.

## Barcode Attribute for the Sun L280 Library

When configuring a Sun L280 library, the `jbconfig` program will incorrectly set the Barcode Reader attribute to Yes. You must edit the Jukebox resource to change the setting for this attribute to No. For more information, see [“To Edit Autochanger Attributes” on page 373](#).

---

## Diagnostic Tools

A variety of diagnostic tools are available as operating system services and as part of the Sun StorEdge EBS product.

To verify that communications sessions can be established, test with `ping` and `rpcinfo`, which are tools provided with the operating system software.

Because the Sun StorEdge EBS server relies heavily on mapping of ports, use `rpcinfo` to test the operation of the portmapper. Using `ping` tests the connection up to the network layer in the OSI model; `rpcinfo` checks for communication up to the session layer. For instructions on using `ping` and `rpcinfo`, see [“Client/Server Communications” on page 550](#).

For more tools on testing communications, contact Sun Technical Support.

## Vendor-Specific ASC and ASCQ Codes Supported in Sun StorEdge EBS Software

With release 7.0 of the Sun StorEdge EBS software, when you install the server or storage node, the Additional Sense Code (ASC), `ascd` code, and vendor-specific directories are automatically installed. This `ascd` code utility translates standard SCSI and many vendor-specific ASC and Additional Sense Code Qualifiers (ASCQ) into English error messages.

The ASC utility is located in the Sun StorEdge EBS driver package. Since the driver package is in the storage node, the utility is automatically installed with the server and storage node. If you only install the client package, the ASC utility is not installed.

The ASC utility translates vendor-specific error messages for the following vendors:

- ADIC
- ATL
- BHTi
- Benchmark
- Exabyte
- Fujitsu
- HP
- IBM
- Overland
- Plasmon
- Qualstar
- Quantum
- Seagate
- STK
- Tandberg

---

# Sun StorEdge EBS Locale and Code Set Support

Sun StorEdge EBS software does not support locales (defined by your operating system) or code sets that remap characters having special meaning for filesystems. Depending on the filesystem, 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).

The following sections discuss Sun StorEdge EBS support of locales and code sets:

- [“Locale Support” on page 577](#)
- [“Code Set Support” on page 578](#)

## Locale Support

Sun StorEdge EBS software supports *only* the following mixed locale environments:

- An English server and a supported localized client
- An 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

TABLE A-4 indicates the Solaris Language environment settings for each supported code set supported in Sun StorEdge EBS releases 6.0.1L and later.

**TABLE A-4** Solaris Language Environment Settings for Supported Code Sets

| Locale             | Solaris Language Environment Settings     |
|--------------------|-------------------------------------------|
| French             | fr<br>fr_FR.ISO8859-1<br>fr_FR.ISO8859-15 |
| Japanese EUC       | ja                                        |
| Simplified Chinese | zh<br>Chinese                             |

---

## Resource Database Notes

With Sun StorEdge EBS release 7.0, resource information resides in the following directory structure:

```
/nsr/res/nsrdb/00
.
.
.
/nsr/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 /nsr/res/nsr1a.res file.

---

## Viewing Resources

You can view all Sun StorEdge EBS resources (and administer a Sun StorEdge EBS release 7.0 server) using the Sun StorEdge EBS Administrator program from any release later than 6.0.



---

**Caution** – 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 script access to the resource database is through either of the following:

- **nsradmin -s** *server*
- **nsradmin -d** */nsr/res/nsrdb*

If you inadvertently specify the wrong path with the `nsradmin -d` command, empty resource directories are created. If this happens, manually delete the incorrect directories.

---

## Restoring Old Resource Files

There might be situations in which you want to restore resource files from a previous Sun StorEdge EBS release. For example, after you have updated to Sun StorEdge EBS release 7.0, you might want to recover a bootstrap that was created by the older Sun StorEdge EBS release.

To restore the resource files from a Sun StorEdge EBS release earlier than 7.0:

1. Use the Sun StorEdge EBS `mmrecov` command to restore the resource files. For further information, refer to the `mmrecov` man page.
2. Stop all Sun StorEdge EBS daemons.
3. Rename or remove the `/nsr/res/nsrdb` directory.
4. Move or copy `/nsr/res.R` (which was created by `mmrecov`) to `/nsr/res`.
5. Start the Sun StorEdge EBS daemons.

## Error Logged If Old and New Resource Files Exist

If you start the `nsrd` daemon while both the `nsrdb` database and either of the old resource files (`nsr.res` or `nsrjb.res`) is present, the older resource files are ignored and the following message is written to `daemon.log`:

```
nsrd: NSR configuration database (/nsr/res/nsrdb) already exists,
previous release version ignored
```

This situation could occur if you install (or update to) Sun StorEdge EBS release 7.0 and then revert to an earlier release, or if you perform a disaster recovery using a bootstrap that was created before the resource files were converted to the new format.

If you encounter this error:

1. Stop all Sun StorEdge EBS services.
2. Perform one of the following:
  - If you want to use the old resource configuration, remove the entire `/nsr/res/nsrdb` directory structure.
  - If you want to use the new resource configuration, remove the `nsr.res` and `nsrjb.res` files.
3. Restart the Sun StorEdge EBS services.

## Repairing Resource Database Corruption

Corruption of Sun StorEdge EBS resource database files can be caused by such things as a power outage, an operating system crash, or manual editing. 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 /nsr/res/nsrdb/01/00019803aa14713c89456b41

nsrd: Invalid resource saved at
/nsr/res/nsrdb/dbg/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 Sun StorEdge EBS Administrator program or the `nsradmin` command.

After you inspect a corrupt resource file, delete it.

If you do not know the cause of the resource file corruption, contact Sun Technical Support for assistance.

## UNIX and Linux Platform-Specific Notes

---

This appendix addresses Sun StorEdge Enterprise Backup Software software features and functionality specific to one or more UNIX platforms. The main sections in this chapter are:

- [“Solaris” on page 581](#)
- [“Linux” on page 582](#)
- [“HP-UX” on page 583](#)
- [“AIX” on page 595](#)
- [“IRIX” on page 596](#)

---

### Solaris

This section provides information specific to Sun StorEdge EBS software running 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.

The best solution is to set the hidden attribute Executable Path for a client that has this problem.

## ▼ To Set the Executable Path

To set the Executable Path:

1. In the Sun StorEdge EBS Administrator program, select Client Setup from the Client menu.
2. In the Clients dialog box, select Details from the View menu to display hidden attributes.
3. For the Executable Path attribute, enter the path of the executables, `/usr/sbin/nsr`.

Another solution is to modify the search path for root on the Sun StorEdge EBS server to include `/usr/sbin/nsr` even if it does not exist locally.

---

## Linux

This section provides information specific to Sun StorEdge EBS software running on the Linux platform.

### Configure Linux Operating System to Detect SCSI Devices

Proper configuration of the SCSI subsystem is required to get full usage of SCSI devices and allow the operating system to detect SCSI devices attached to the machine. If you have a device configured with multiple LUNs, then the kernel parameter "Probe all LUNs of each SCSI Device" must be set to Yes. For more information on configuring the Linux 2.4 SCSI subsystem, refer to *The Linux Documentation Project* web site. For information on your SCSI device, contact the manufacturer.



# Linux Journaled Filesystem Support

Sun StorEdge EBS software supports backup and recovery operations on the following Linux journaled filesystems:

- ext3
- reiserfs
- jfs
- xfs

However, Sun StorEdge EBS software does not support backup and recovery of extended attributes or ACLs on these systems.



---

**Caution** – If you have an ext3 filesystem with the journal set to visible, do not back up or recover the journal. Recovering the journal might cause the filesystem to become unstable. You can use a Sun StorEdge EBS directive to ensure that this filesystem is excluded from a backup. For information about using Sun StorEdge EBS directives, see [“Directives” on page 144](#).

---

## HP-UX

This section provides information specific to Sun StorEdge EBS storage node software running on the HP-UX platform.

## Autochanger Installation on an HP-UX System

The following sections explain how to install and configure the drivers supplied by Hewlett-Packard.

### Selecting SCSI Addresses for the Autochanger

Determine which SCSI address is assigned to each SCSI bus and select 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.

SCSI addresses range from 0 to 7. However, SCSI address 7 is reserved for the SCSI interface board in on a computer, so you cannot use it. The primary hard disk is usually on SCSI address 6.



---

**Caution** – Some devices, such as the HP Model 48AL autochanger, choose 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. Follow the example that matches your configuration.

## Installing the SCSI Pass-Through Driver

There are two procedures for installing a SCSI pass-through (spt) driver, each depending on the type of hardware you are installing:

- If you are installing a SCSI pass-through driver on an HP 9000 Series 800 containing an HP-PB interface card, see, [“To Install an HP-PB or NIO Pass-Through Driver” on page 584](#).
- If you are installing a SCSI pass-through driver on an HP 9000 Series 700 or 800 containing an HSC or PCI interface card, see, [“To Install a GSC, HSC, or PCI Pass-Through Driver” on page 587](#).

### ▼ *To Install an HP-PB or NIO Pass-Through Driver*

For further information regarding 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 you are using:

```
ioscan -f
```

A screen output similar to the following should appear:

| Class   | I  | H/W Path   | Driver | S/W State | H/W Type  | Description                  |
|---------|----|------------|--------|-----------|-----------|------------------------------|
| ext_bus | 7  | 4/8        | c720   | CLAIMED   | INTERFACE | GSC add-on<br>FAST/WIDE SCSI |
| target  | 4  | 4/8.1      | tgt    | CLAIMED   | DEVICE    |                              |
| tape    | 6  | 4/8.1.0    | stape  | CLAIMED   | DEVICE    | STK 4890                     |
| ext_bus | 8  | 8/4/4      | scsi3  | CLAIMED   | INTERFACE | HP 28696A-wide<br>SCSE ID=7  |
| target  | 5  | 8/4/4.0    | target | CLAIMED   | DEVICE    |                              |
| tape    | 7  | 8/4/4.0.0  | tape2  | CLAIMED   | DEVICE    | Quantum DLT4000              |
| ext_bus | 12 | 10/4/4     | scsi3  | CLAIMED   | INTERFACE | HP 28696A-Wide<br>SCSI ID=7  |
| target  | 6  | 10/4/4.6   | target | CLAIMED   | DEVICE    |                              |
| unknown | -1 | 10/4/4.6.0 |        | UNCLAIMED | UNKNOWN   | HP 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/sbin/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 desired device (from the `ioscan` command).

If you used the previous `ioscan` display with the library robotics connected to the hardware path `10/4/4.6.0`, the line 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` might prevent the system from recognizing tape drives added to that SCSI bus. You might have 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: 97/07/18 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:

```
ioscanner -f -C spt
```

Output from the `ioscanner` should resemble the following:

| Class | I | H/W Path   | Driver      | S/W State | H/W Type      | Description |
|-------|---|------------|-------------|-----------|---------------|-------------|
| spt   | 0 | 10/4/4.6.0 | SPT CLAIMED | device    | HP C6280-7000 |             |

14. Proceed to [“To Create a Device File” on page 589](#).

## ▼ 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 [“To Create a Device File” on page 589](#).  
**Note:** You can choose any name for the device as long as it is not something reserved by the system. For example, do not use 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 “Are you sure?” indicate Yes, and press [Return].
8. The Creating Kernel message appears (for approximately 1 to 2 minutes), followed by the Move Kernel Message. Select OK and press [Return]. The system reboots.
9. Proceed to [“To Create a Device File” on page 589](#).

## Installing the Autochanger Hardware

To install the autochanger hardware, follow the instructions in the vendor’s hardware installation manuals. Use the SCSI addresses that you selected in the section, [“Selecting SCSI Addresses for the Autochanger” on page 583](#), to set any required switches on the autochanger for the devices and robotics controller.

For example, using the previous Series 800 example, you set the robotics SCSI address on the autochanger to 3 and the device SCSI address to 4, according to the instructions in the hardware manual supplied by the vendor.

### ▼ *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 entering 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. If *time* is zero, **shutdown** gives users very little time to log out but it runs more quickly.

2. Turn the power to the server off.

## ▼ *To Start the System*

To start the system:

1. Turn the main power switch of the autochanger to the On position.
2. Turn the power to the server on.

This reboots the server, as required after installing new autochanger hardware.

## ▼ *To Create a Device File*

To create a device file:

1. Verify that the *spt* has been 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, you must verify that the installation has been completed properly.
3. If the device entry is already completed by the operating system, then use the OS-defined entry and proceed to verify the installation.
4. If the device entry is not defined, you must create a device file. Use the **mknod** command to create a device file to access the robotics controller. The command syntax is:

```
/user/sbin/mknod /dev/sji devfilename c majornum minornum
```

- *devfilename* is the user-defined name of the device file
- *majornum* is the character 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; proper *ext\_bus* entry)
- *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:

|            | Character | Block | Driver | Class |
|------------|-----------|-------|--------|-------|
| HP-PB      | 75        | -1    | spt    | spt   |
| HSC or PCI | 203       | -1    | sctl   | ctl   |

The value for *majornum* is the number in the Character column.

## Minor Number

To determine the value for *minornum*, use the `ioscan` command. The relevant lines in the `ioscan` output are those:

- Pertaining to the controller itself (which contains HP C6280-7000 in the Description column).



- Pertaining to 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 Path   | Driver | S/W State | H/W Type  | Description      |
|---------|----|------------|--------|-----------|-----------|------------------|
| ext_bus | 12 | 10/4/4     | scsi3  | CLAIMED   | INTERF    | ACEHP 28696-Wide |
|         |    |            |        |           | SCSI ID=7 |                  |
| target  | 6  | 10/4/4.0   | target | CLAIMED   | DEVICE    |                  |
| spt     | 0  | 10/4/4.6.0 | spt    | CLAIMED   | DEVICE    | HP C6280-7000    |

Given the information in the previous example, the correct `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 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 Path   | Driver | S/W State | H/W Type | Description   |
|-------|---|------------|--------|-----------|----------|---------------|
| spt   | 0 | 10/4/4.6.0 | schgr  | CLAIMED   | DEVICE   | HP C6280-7000 |

If the `schgr` driver is not configured on the system, no driver appears to be associated with the library. The `ioscan` output line resembles:

| Class   | I  | H/W Path   | Driver | S/W State | H/W Type | Description   |
|---------|----|------------|--------|-----------|----------|---------------|
| unknown | -1 | 10/4/4.6.0 | schgr  | UNCLAIMED | DEVICE   | HP C6280-7000 |

## ▼ To Test the Device Driver and Device File Installation

After you have installed the device driver and created the device file, 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
```

## ▼ 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 might 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 you encounter this message:

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 at the prompt:

```
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 and do not support the use of 90- or 120- meter tapes. To determine the type of tape an HP drive supports, refer to the drive's hardware manual.

- If you attempt to use unsupported media in an HP tape drive, you might encounter the following types of error messages in the specified situations: When you use the `nsrmm` or `nsrjb` command to label the tape:

```
nsrmm: error, label write, No more processes (5)
```

- When you attempt to use the `scanner -i` command:

```
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 you use the `nsrjb -u -S` command 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** – When you use the `nsrjb` command with the `-f devicename` attribute and the device name contains backslashes, you must use the backslash “escape” character with the backslash. Therefore, a device with the name `rd=sn1:\\.\\Tape0` would be entered at the command line 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 filesystem the symbolic link points to.

---

# AIX

This section provides information specific to Sun StorEdge EBS storage node software running on the AIX platform.

## ▼ To Activate Full Diagnostic Core Dumps on AIX

Earlier releases of Sun StorEdge EBS software on an AIX system do not automatically provide full diagnostic core dumps. To activate full diagnostic core dumps you need to set the `NSR_FULLDUMP` environment variable to 1.

To set the environment variable and the core file size:

1. Log in as root and set the environment variable with one of the following commands as appropriate for your shell tool:

```
export NSR_FULLDUMP=1

setenv NSR_FULLDUMP 1
```

2. Use the System Management Interface Tool (SMIT) or edit the `/etc/security/limits` file to set the core file size to -1 for root or default.
3. Set the Enable Full CORE Dump option to True (select Change or Show Characteristics of Operating System from the System Environments choices).

For the core file size change to take effect, you might need to reboot the system after you edit the limits file. You can check the file size limit with the `ulimits -a` command. For the size that you set in the procedure described, the `ulimits` command should return a value of “unlimited.”

## STK-9840 Drives Attached to AIX

If you attach an STK-9840 drive to an AIX server for Sun StorEdge EBS backups, 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 you use an IRIX Sun StorEdge EBS storage node, you might receive the following message 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. You need to ensure that the device is supported and configured for your machine by applying a kernel patch, available from the Silicon Graphics web site.

To resolve the problem:

1. Go to the web page [www.sgi.com/support/patch\\_intro.html](http://www.sgi.com/support/patch_intro.html).
2. Download the recommended patch.
3. Follow the instructions provided to apply the patch set.

SGI tests specific firmware revisions for each drive and robot with IRIX Sun StorEdge EBS software. Even if you use drives and robotics that SGI has qualified, you might need to apply patches to IRIX to make the combination work.

Visit the SGI Services and Support web site for information about and distribution of patches you might require. You can also obtain information regarding the current drives, robotics, and firmware revisions IRIX supports.

# Glossary

---

This glossary provides definitions for terms used in this guide.

|                             |                                                                                                                                                                                                                                                                                                                                                                                                          |
|-----------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>ACL</b>                  | <p>An abbreviation for access control list. This is a list that specifies the permissions assigned to a specific file or directory.</p> <p>To recover a file that has an associated ACL, you must either be logged into the system as root, as a Sun StorEdge EBS Administrator, or as the file's owner.</p>                                                                                             |
| <b>active group</b>         | <p>A Sun StorEdge EBS backup group that has its Autostart attribute enabled.</p>                                                                                                                                                                                                                                                                                                                         |
| <b>administrator</b>        | <p>The person normally responsible for installing, configuring, and maintaining Sun StorEdge EBS software.</p>                                                                                                                                                                                                                                                                                           |
| <b>Administrators group</b> | <p>A Microsoft Windows user group whose members have all the rights and abilities of users in other groups, plus the ability to create and manage all the users and groups in the domain. Only members of the Administrators group can modify Windows NT and Windows 2000 operating system files, maintain the built-in groups, and grant additional rights to groups.</p>                               |
| <b>annotation</b>           | <p>A comment that you associate with an archive save set to help identify that data later. Annotations are stored in the media database for ease of searching and are limited to 1,024 characters.</p>                                                                                                                                                                                                   |
| <b>archive</b>              | <p>The process by which Sun StorEdge EBS software backs up directories or files to an archive volume and then grooms them to free disk space. When data is archived, it is written to one or more storage volumes and then marked so that it is never subject to automatic recycling. You can delete the archived files from the client, thus freeing disk space. See also <a href="#">grooming</a>.</p> |
| <b>archive clone pool</b>   | <p>A pool composed exclusively of archive clone save sets.</p>                                                                                                                                                                                                                                                                                                                                           |

|                               |                                                                                                                                                                                                                                                                                                                                                                                                                               |
|-------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>archive pool</b>           | A volume <a href="#">pool</a> composed exclusively of archive save sets. Archived save sets are in a different format than regular backup save sets, and must be maintained on separate media.                                                                                                                                                                                                                                |
| <b>archive volume</b>         | A tape or other storage medium used to store Sun StorEdge EBS archive data, as opposed to a <a href="#">backup volume</a> .                                                                                                                                                                                                                                                                                                   |
| <b>ASM</b>                    | An abbreviation for Application Specific Module. An ASM is a program that, when used in a directive, specifies the way that a set of files or directories is to be backed up and recovered. For example, <b>compressasm</b> is a Sun StorEdge EBS directive used to compress files.                                                                                                                                           |
| <b>ASM specification</b>      | A directive that specifies how files or directories with a matching pattern are backed up. This specification appears in the format:<br>[+] <i>asm: argument</i><br>For more information, refer to the <b>nsr_5</b> man page or the <i>Legato Command Reference Guide</i> .                                                                                                                                                   |
| <b>attribute</b>              | A feature of a resource. It is a service or information that the <a href="#">resource</a> provides.                                                                                                                                                                                                                                                                                                                           |
| <b>authorization code</b>     | A code that is unique to a network that unlocks the Sun StorEdge EBS software for permanent use.                                                                                                                                                                                                                                                                                                                              |
| <b>autochanger</b>            | A mechanism that uses a robotic arm to move media among various components located in a device, including slots, media drives, media access ports, and transports. Autochangers automate media loading and mounting functions during backup and recovery. The term autochanger refers to a variety of robotic libraries, including autoloader, <a href="#">carousel</a> , datawheel, jukebox, library, and near-line storage. |
| <b>auto media management</b>  | A feature that enables the storage device controlled by the Sun StorEdge EBS server to automatically label, mount, and overwrite a volume it considers unlabeled. For stand-alone devices, volumes that are eligible for reuse are also automatically recycled.                                                                                                                                                               |
| <b>backup</b>                 | The writing of saved data to a volume.                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>backup cycle</b>           | The period of time from one level full backup to the next level full backup.                                                                                                                                                                                                                                                                                                                                                  |
| <b>backup group</b>           | See <a href="#">group</a> .                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>backup level</b>           | See <a href="#">level</a> .                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Backup Operators group</b> | A group of Microsoft Windows users who can log on 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.                                                                                                                                                                                                                                        |



|                                |                                                                                                                                                                                                                                                                                                                                                                                                                      |
|--------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>backup volume</b>           | A tape or other storage medium used to store Sun StorEdge EBS backup data, as opposed to an archive volume.                                                                                                                                                                                                                                                                                                          |
| <b>base enabler code</b>       | See <a href="#">enabler code</a> .                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>bootstrap</b>               | A save set that is essential for the Sun StorEdge EBS disaster recovery procedures. The bootstrap is composed of three components that reside on the Sun StorEdge EBS server: the <a href="#">media database</a> , the <a href="#">resource database</a> , and a <a href="#">server index</a> .                                                                                                                      |
| <b>browse policy</b>           | A policy that determines how long entries for backed up data remain in the client file index.                                                                                                                                                                                                                                                                                                                        |
| <b>carousel</b>                | A tray or tape cartridge that holds multiple backup volumes.                                                                                                                                                                                                                                                                                                                                                         |
| <b>client</b>                  | A computer that accesses the Sun StorEdge EBS server to back up or recover files. Clients may be workstations, computers, or file servers.                                                                                                                                                                                                                                                                           |
| <b>client file index</b>       | A database of information the Sun StorEdge EBS server maintains that tracks every database object, file, or <a href="#">filesystem</a> backed up. The Sun StorEdge EBS server maintains a single client index file for each client computer.                                                                                                                                                                         |
| <b>client-initiated backup</b> | See <a href="#">manual backup</a> .                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>clone</b>                   | The Sun StorEdge EBS process used to make an exact copy of saved data (save sets). You can clone individual save sets or the entire contents of a backup volume. Cloning is different from a simple copy operation carried out on an operating system or hardware device. Sun StorEdge EBS software indexes and tracks clones in both the <a href="#">client file index</a> and the <a href="#">media database</a> . |
| <b>clone pool</b>              | A pool of volumes composed exclusively of cloned data. There are three types of clone pools: backup clone, archive clone, and migration clone. Save sets of different types (for example, archive and migration) cannot be intermixed on the same clone volume.                                                                                                                                                      |
| <b>clone volume</b>            | A volume belonging to a clone pool.                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>component</b>               | A Sun StorEdge EBS-defined group of related data that must be treated as a single unit for backup and recovery. Within <a href="#">VSS</a> , a component is a subordinate unit of a <a href="#">writer</a> .                                                                                                                                                                                                         |
| <b>connection port</b>         | The port Sun StorEdge EBS processes use to perform backup and recovery sessions through a firewall.                                                                                                                                                                                                                                                                                                                  |

|                           |                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>consolidate</b>        | To create a complete backup of a save set by merging the most recent level 1 save set with its corresponding full level save set. For more information, refer to the <b>nsrssc</b> man page or the Legato <i>Command Reference Guide</i> .                                                                                                                                                                  |
| <b>continued save set</b> | Data associated with a save set that is continued from a previous volume. Continued save sets are created by the backup server when large save sets are being backed up, cloned, or archived to multiple volumes.                                                                                                                                                                                           |
| <b>daemon</b>             | A program that lies dormant waiting for a specified condition to occur.                                                                                                                                                                                                                                                                                                                                     |
| <b>datawheel</b>          | See <a href="#">autochanger</a> .                                                                                                                                                                                                                                                                                                                                                                           |
| <b>DDS</b>                | An abbreviation for dynamic drive sharing, which allows Sun StorEdge EBS software to recognize shared drives.                                                                                                                                                                                                                                                                                               |
| <b>device</b>             | <ol style="list-style-type: none"> <li>1. A storage unit that reads from and writes to storage volumes (see volume). A storage unit can be a tape device, optical drive, <a href="#">autochanger</a>, or file connected to the <a href="#">server</a> or <a href="#">storage node</a>.</li> <li>2. When dynamic drive sharing (DDS) is enabled, refers to the access path to the physical drive.</li> </ol> |
| <b>DFS</b>                | An abbreviation for Distributed File System. DFS is a Microsoft Windows add-on that allows you to create a logical directory of shared directories that span multiple machines across a network.                                                                                                                                                                                                            |
| <b>DFS component</b>      | <p>A DFS component is one of the following:</p> <ul style="list-style-type: none"> <li>• A namespace for files and DFS links, called a DFS root.</li> <li>• A connection to a shared file or folder, called a DFS child node</li> </ul>                                                                                                                                                                     |
| <b>directed recovery</b>  | A recovery method used to recover data that originated on one computer and re-create it on another computer.                                                                                                                                                                                                                                                                                                |
| <b>directive</b>          | An instruction that directs Sun StorEdge EBS software to take special actions on a given set of files for a specified client during a backup.                                                                                                                                                                                                                                                               |
| <b>domain controller</b>  | A computer that stores directory data and manages user interactions with a domain, including logon, authentication, directory searches, and access to other shared resources.                                                                                                                                                                                                                               |
| <b>drive</b>              | When dynamic drive sharing (DDS) is enabled, refers to the physical backup object, such as a tape drive, disk, or file. See also <a href="#">device</a> .                                                                                                                                                                                                                                                   |

|                           |                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|---------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>enabler code</b>       | A special code provided by Sun that activates the software. The enabler code that unlocks the base features for software you purchase is referred to as a base enabler. Enabler codes for additional features or products (for example, autochanger support) are referred to as add-on enablers.                                                                                                                                     |
| <b>exit code</b>          | An indicator that specifies whether a backup or recovery session succeeded. An exit code of zero (0) indicates the session completed successfully. A nonzero exit code indicates the session did not complete successfully.                                                                                                                                                                                                          |
| <b>expiration date</b>    | The date when the volume changes from read / write to read-only.                                                                                                                                                                                                                                                                                                                                                                     |
| <b>expired save set</b>   | A save set that has reached its browse time and can no longer be browsed. In addition, the save set has been removed from the client file index.                                                                                                                                                                                                                                                                                     |
| <b>file index</b>         | See <a href="#">client file index</a> .                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>filesystem</b>         | <ol style="list-style-type: none"> <li>1. A file tree located on a specific disk partition or other mount point.</li> <li>2. The entire set of all files.</li> <li>3. A method of storing files.</li> </ol>                                                                                                                                                                                                                          |
| <b>firewall</b>           | A system designed to prevent unauthorized access to or from a private network. All messages entering or leaving the intranet pass through the firewall, which examines each message and blocks those that do not meet the specified security criteria. There are several types of firewall techniques. The Sun StorEdge EBS software supports client backups from computers that are protected by <a href="#">packet filtering</a> . |
| <b>fork</b>               | A subprocess the Sun StorEdge EBS software creates to perform a requested operation. In instances where a command uses a parallelism value, Sun StorEdge EBS software creates multiple instances of that command. Each instance of the command is identical to the original command and is referred to as a subprocess. Once each subprocess is created, they are run simultaneously.                                                |
| <b>full backup</b>        | See <a href="#">level</a> .                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>grooming</b>           | The Sun StorEdge EBS process of removing the original files from a local disk after a successful archive operation.                                                                                                                                                                                                                                                                                                                  |
| <b>group</b>              | A client or group of clients configured to start backing up files to the Sun StorEdge EBS server at a designated time of day.                                                                                                                                                                                                                                                                                                        |
| <b>inactivity timeout</b> | An attribute that indicates the number of minutes Sun StorEdge EBS software waits before determining that a client is unavailable for backup.                                                                                                                                                                                                                                                                                        |

|                        |                                                                                                                                                                                                                                                                                                                                                                                                                   |
|------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>incremental</b>     | See <a href="#">level</a> .                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>jukebox</b>         | See <a href="#">autochanger</a> .                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>legacy method</b>   | The use of special-case Microsoft APIs to back up and recover operating system components, services, and applications. This is the method Sun StorEdge EBS software uses from release 5.7 to release 7.0. In release 7.1, Sun StorEdge EBS software also uses the legacy method in certain situations. For example, if a <a href="#">VSS</a> client license does not exist or if VSS is disabled.                 |
| <b>level</b>           | <p>A measurement that determines how much data Sun StorEdge EBS software saves during a scheduled or manual backup.</p> <p>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.</p> |
| <b>library</b>         | See <a href="#">autochanger</a> .                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>license enabler</b> | The enabler code that enables you to run a feature or product.                                                                                                                                                                                                                                                                                                                                                    |
| <b>local host</b>      | The node on which the client or server program is running.                                                                                                                                                                                                                                                                                                                                                        |
| <b>LUS</b>             | An abbreviation for Legato User SCSI. Sun software products use the LUS driver as a proprietary device driver that sends arbitrary SCSI commands to an autochanger.                                                                                                                                                                                                                                               |
| <b>manual backup</b>   | A backup that a user requests from the client's <b>save</b> program. The user specifies participating files, filesystems, and directories. A manual backup does not generate a <a href="#">bootstrap</a> save set.                                                                                                                                                                                                |
| <b>media</b>           | The physical storage medium to which backup data is written. Sun StorEdge EBS software supports tape, magnetic or optical disk, and filesystems as backup media. See also <a href="#">volume</a> .                                                                                                                                                                                                                |
| <b>media database</b>  | A database that contains indexed entries about the storage volume location and the life cycle status of all data and volumes the Sun StorEdge EBS server manages. See also <a href="#">volume</a> .                                                                                                                                                                                                               |
| <b>metadata</b>        | <a href="#">VSS</a> -defined information, stored in an XML file, that is passed from the <a href="#">writer</a> to the <a href="#">requester</a> . Metadata information includes the writer name, a list of <a href="#">components</a> to back up, a list of components to be excluded from the backup, and the methods to use for recovery.                                                                      |

|                                      |                                                                                                                                                                                                                                                                                                                                             |
|--------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>migration</b>                     | The process of moving data from a local filesystem to storage media in the migration store to free up disk space on the local drive.                                                                                                                                                                                                        |
| <b>multiplexing</b>                  | A Sun StorEdge EBS feature that permits data from more than one save set to be simultaneously written to the same storage device.                                                                                                                                                                                                           |
| <b>NDMP</b>                          | An abbreviation for Network Data Management Protocol, which is a storage management client/server protocol for enterprise-wide backup of network-attached storage. Sun StorEdge EBS software uses NDMP to provide connections to computers with NDMP data modules for tape operations, allowing a significant reduction in network traffic. |
| <b>near-line storage</b>             | See <a href="#">autochanger</a> .                                                                                                                                                                                                                                                                                                           |
| <b>Sun StorEdge EBS client</b>       | See <a href="#">client</a> .                                                                                                                                                                                                                                                                                                                |
| <b>Sun StorEdge EBS server</b>       | See <a href="#">server</a> .                                                                                                                                                                                                                                                                                                                |
| <b>Sun StorEdge EBS storage node</b> | See <a href="#">storage node</a> .                                                                                                                                                                                                                                                                                                          |
| <b>NFS client</b>                    | A computer that can access files on a network filesystem (NFS) server.                                                                                                                                                                                                                                                                      |
| <b>NFS server</b>                    | A computer that contains exported filesystems that NFS clients can access.                                                                                                                                                                                                                                                                  |
| <b>nonclone pool</b>                 | Pools that contain data that has not been cloned.                                                                                                                                                                                                                                                                                           |
| <b>notification</b>                  | A message generated and sent to the Sun StorEdge EBS administrator about important Sun StorEdge EBS events.                                                                                                                                                                                                                                 |
| <b>online indexes</b>                | The databases located on the Sun StorEdge EBS server that contain all the information pertaining to the client backups ( <a href="#">client file index</a> ) and backup volumes ( <a href="#">media database</a> ).                                                                                                                         |
| <b>operator</b>                      | The person who monitors the server status, loads backup volumes into the server devices, and otherwise executes the day-to-day Sun StorEdge EBS tasks.                                                                                                                                                                                      |
| <b>override</b>                      | A Sun StorEdge EBS feature that allows you to configure a different backup level for a specific date listed in a Schedule resource.                                                                                                                                                                                                         |
| <b>packet filtering</b>              | A method of firewall protection that looks at each packet entering or leaving the network, and accepts or rejects it based on user-defined rules. See also <a href="#">firewall</a> .                                                                                                                                                       |

|                            |                                                                                                                                                                                                                                                                                               |
|----------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>parallelism</b>         | A Sun StorEdge EBS feature that enables the Sun StorEdge EBS server to either back up save sets from several clients or multiple save sets from one client at the same time. Parallelism is also available during recoveries.                                                                 |
| <b>pathname</b>            | A set of instructions to the operating system for accessing a file. An <i>absolute pathname</i> indicates how to find a file starting from the root directory and working down the directory tree. A <i>relative pathname</i> indicates how to find the file starting from where you are now. |
| <b>policy</b>              | A set of constraints that specify how long an entry can remain in a client file index. When a policy expires, the save sets associated with that policy are marked recyclable. Each Client resource uses two policies: a browse policy and a retention policy.                                |
| <b>pool</b>                | A feature that enables you to sort backup data to selected volumes. A pool contains a collection of backup volumes to which specific data has been backed up.                                                                                                                                 |
| <b>probe</b>               | The process Sun StorEdge EBS software uses to determine the directories or files to back up on each client.                                                                                                                                                                                   |
| <b>purging</b>             | The process of deleting all entries for files on the volume from the client file index, but allowing entries for the save sets to remain in the media database.                                                                                                                               |
| <b>recover</b>             | A recovery method that re-creates an image of the client filesystems and database on the Sun StorEdge EBS server.                                                                                                                                                                             |
| <b>recyclable save set</b> | A save set whose browse and retention policies have been reached; therefore, the save set has been removed from the media database.                                                                                                                                                           |
| <b>recyclable volume</b>   | A volume whose data has passed both its browse and retention policies and is now available for relabeling and use by a Sun StorEdge EBS server or storage node.                                                                                                                               |
| <b>Registry</b>            | A database of configuration information central to Microsoft Windows operations. It centralizes all Windows settings and provides security and control over system, security, and user account settings.                                                                                      |
| <b>remote device</b>       | A storage device that is attached to a Sun StorEdge EBS storage node.                                                                                                                                                                                                                         |
| <b>requester</b>           | An application that works within the <a href="#">VSS</a> infrastructure to initiate the creation and destruction of a <i>shadow copy</i> . Sun StorEdge EBS software is an example of a requester.                                                                                            |

|                               |                                                                                                                                                                                                                                                                                                                                                              |
|-------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>RPC</b>                    | An abbreviation for remote procedure call, which is the protocol the Sun StorEdge EBS server uses to perform client requests over a network.                                                                                                                                                                                                                 |
| <b>resource</b>               | A component of Sun StorEdge EBS software that describes the Sun StorEdge EBS server and its clients. Devices, schedules, clients, groups, and policies are examples of Sun StorEdge EBS resources. Each resource consists of a list of attributes that define the parameters to use for the specific Sun StorEdge EBS resource.                              |
| <b>resource database</b>      | A database that contains information about each of the configured backup server's resources.                                                                                                                                                                                                                                                                 |
| <b>retention policy</b>       | A policy that determines how long save set entries are retained in the Sun StorEdge EBS server's media database.                                                                                                                                                                                                                                             |
| <b>retrieve</b>               | The process of locating and copying back files and directories that Sun StorEdge EBS software has archived.                                                                                                                                                                                                                                                  |
| <b>retry mechanism</b>        | 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. By using this mechanism, a previously failed operation might be more successful. Another common situation whereby a retry mechanism might succeed is when the client is in a reboot cycle. |
| <b>root</b>                   | <ol style="list-style-type: none"> <li>1. (UNIX only) The UNIX superuser account (with user name "root" and user ID). By extension, the privileged system-maintenance login on any operating system.</li> <li>2. (Microsoft Windows and UNIX) The top node of the system directory structure; the home directory of the root user.</li> </ol>                |
| <b>save set</b>               | A group of files or a filesystem from a single client computer backed up onto storage media.                                                                                                                                                                                                                                                                 |
| <b>save set consolidation</b> | The process that merges a level 1 backup with the last full backup of a save set to create a new full backup. See also <a href="#">level</a> .                                                                                                                                                                                                               |
| <b>save set ID</b>            | An internal identification number that Sun StorEdge EBS software assigns to a save set.                                                                                                                                                                                                                                                                      |
| <b>save set recover</b>       | The recovery of specified save sets to the Sun StorEdge EBS server.                                                                                                                                                                                                                                                                                          |
| <b>save set status</b>        | The Sun StorEdge EBS attribute that indicates whether a given save set is browsable, recoverable, or recyclable. The save set status also indicates whether the save set has been successfully backed up.                                                                                                                                                    |

|                           |                                                                                                                                                                                                                                                                                                                                  |
|---------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>save stream</b>        | The data and save set information being written to a storage volume during a backup. A save stream originates from a single save set.                                                                                                                                                                                            |
| <b>server</b>             | The computer on a network that runs the Sun StorEdge EBS server software, contains the online indexes, and provides backup and recovery services to the clients and storage nodes on the same network.                                                                                                                           |
| <b>server index</b>       | A database containing information about the server's files that have been backed up during scheduled backups. Also known as the server's <i>client file index</i> .                                                                                                                                                              |
| <b>service port</b>       | The port used by a server or storage node to listen for backup and recovery requests from clients through a firewall.                                                                                                                                                                                                            |
| shadow copy               | A temporary, point-in-time copy of a volume created using VSS technology. See also <i>Volume Shadow Copy Service</i> .                                                                                                                                                                                                           |
| <b>shell prompt</b>       | A cue in a shell window where you enter a command.                                                                                                                                                                                                                                                                               |
| <b>silo</b>               | A repository for holding hundreds or thousands of volumes. Silo volumes are identified by bar codes, not by slot numbers.                                                                                                                                                                                                        |
| <b>skip</b>               | A backup level in which designated files are not backed up. See also <i>level</i> .                                                                                                                                                                                                                                              |
| <b>SMS</b>                | An abbreviation for system management software, which is a Microsoft-based software installation system that allows the wide-scale, automatic installation of software products on clients from a single remote server.                                                                                                          |
| <b>SNMP</b>               | An abbreviation for Simple Network Management Protocol, which is a protocol that defines the communication between a manager (sometimes called a monitor or management station) and an object (the item being managed). Sun StorEdge EBS software uses SNMP to send messages to the administrator about Sun StorEdge EBS events. |
| <b>ssid</b>               | See <i>save set ID</i> .                                                                                                                                                                                                                                                                                                         |
| <b>staging</b>            | The process of moving data from one storage medium to another, less costly medium, and later removing the data from its original location.                                                                                                                                                                                       |
| <b>stand-alone device</b> | A <i>storage device</i> that contains a single drive for backing up data.                                                                                                                                                                                                                                                        |
| <b>storage device</b>     | The hardware that reads and writes data during backup, recovery, or other Sun StorEdge EBS operations.                                                                                                                                                                                                                           |
| <b>storage node</b>       | A storage device physically attached to another computer whose backup operations are administered from the controlling Sun StorEdge EBS server.                                                                                                                                                                                  |



|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|----------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>user groups</b>         | In the Sun StorEdge EBS software, refers to configuration resources that are used to assign users to access control groups and to configure the privileges associated with those groups.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>versions</b>            | The date-stamped collection of available backups for any single file.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>volume</b>              | A unit of storage media, such as a magnetic tape, an optical disk, or a file. A storage device reads from and writes to volumes, which can be physical units (for example, a labeled tape cartridge) or logical units (for example, optical media can store multiple volumes on a single physical platter).                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>volume ID</b>           | The internal identification Sun StorEdge EBS software assigns to a backup volume.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>volume name</b>         | The name you assign to a backup volume when it is labeled.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>volume pool</b>         | See <a href="#">pool</a> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| Volume Shadow Copy Service | <p>A Microsoft technology that acts as a coordinator among all the components that create, archive, modify, back up, or restore data. Volume Shadow Copy Service (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 filesystem, Sun StorEdge EBS software backs up data from the snapshot. Backing up data from a snapshot ensures that:</p> <ul style="list-style-type: none"> <li>• Applications can continue to write data to the volume during a backup.</li> <li>• Open files are no longer omitted during a backup.</li> </ul> <p>Backups can be performed at any time, without locking out users.</p> |
| VSS                        | See <a href="#">Volume Shadow Copy Service</a> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| writer                     | <p>A database, system service, or application code that works with VSS to provide <a href="#">metadata</a> information about what data to back up, and specific methods for properly handling <a href="#">components</a> and applications during backup and restore. Writers also identify the type of application or service that is being backed up.</p>                                                                                                                                                                                                                                                                                                                                                 |



# Index

---

## A

actions, notifications, [67](#)

administrators

adding, [60](#)

privileges, [39](#), [59](#)

storage nodes, [85](#)

Sun StorEdge EBS, [59](#)

Advanced File Type Device

concurrent operations, [211](#) to [212](#)

notification, [208](#)

retention policy, [208](#)

advanced file type device, [207](#)

affinity

clone storage node, [357](#)

storage node, [86](#)

AIX

core dumps, [595](#)

Alias attribute, [557](#)

aliases, host sharing restriction, [541](#)

ansrd program, [9](#)

application information field

Auspex, [503](#)

EMC, [499](#), [500](#)

NetApp, [495](#)

application specific module (ASM), [145](#)

archiving, [321](#)

archive pools

configuration, [174](#)

creating, [180](#)

errors, [560](#)

archive requests

creating, [331](#)

defined, [331](#)

deleting, [333](#)

disabling, [336](#)

editing, [332](#)

archive services

enabling, [328](#)

archive volume pool, [323](#)

cloning, [341](#)

display errors, [559](#)

compared to backups, [322](#)

defined, [19](#)

enabling, [326](#)

grooming files, [321](#)

indexed, [324](#)

licensing, [321](#)

managing, [335](#)

manual, [328](#)

manually archiving data, [329](#)

naming archive requests, errors, [560](#)

NDMP, [464](#)

nonindexed, [325](#)

nsrchive program considerations, [560](#)

permissions, [325](#) to [326](#)

pools

types of pools, [326](#)

remote requests, failure, [559](#)

requirements, [322](#)

retrieving from client machine, [333](#)

save sets, [322](#)

multiple, [559](#)

retrieving, [334](#)

scheduled, [331](#)

scheduling, [336](#)

starting, [336](#)

- stopping, 336
- tracking entries, 323
- troubleshooting, 559
- viewing details, 337
- volumes, cloning, 356
- ASM (application specific module), 145
- attributes
  - defined, 7
  - hidden, 42
  - See also* specific attribute name
- audience, for documentation, xlvii
- Auto Clean attribute, 371
- auto media management
  - autochangers, 399
  - cloning, considerations, 341
  - defined, 226
  - recyclable volumes, 125
- Auto Media Management attribute, 371
- auto media verification, 178
  - media position errors, 538
- Auto Media Verify attribute, 178
- Autochanger Module, 19
- autochangers, 188
  - adding, 361, 367, 378
  - AIX considerations, 595
  - ATL 2640 libraries, 571
  - attributes, 371, 395
    - new, 562
  - Auspex support, 502
  - auto media management, 399
  - Autochanger resource, 8
  - autodetection, 567
  - barcode labeling, 404, 406
  - barcodes labeling, 405
  - cartridges, 393
  - cleaning, 393
    - cartridges, 397
    - non-default slot, 395
  - control ports access, 571
  - defined, 6, 360
  - deleting, 374
  - devices
    - adding, 371 to 373
    - calibrating, 413
    - cleaning, 393, 395
    - drivers, 564
    - remote, 375
  - editing, 373
  - element status, disabling, 413
  - EMC Celerra support, 498
  - EMC IP4700 support, 499
  - emulation, 360
  - HP-UX considerations, 583
  - import/export function
    - depositing, 410
    - withdrawing, 410
  - installation, HP-UX considerations, 583
  - inventories, 407, 568
    - slots, 572
  - jbconfig program, 361, 367, 378
  - licensing, 359, 384
  - maintenance commands, 563
  - NDMP support, 468, 469
  - NetApp support, 493
  - notifications, 411
  - operations, 391
  - optical, 414
  - pools, 412
  - remote management, 374
  - requirements, 360
  - resetting, 413
  - sharing, 384
    - block size, 390
    - cloning, 388
    - configuration, 385
    - defined, 385
    - devices, 386
    - host device selection, 386
    - licensing, 384
    - NDMP clients, 471
    - recoveries, 386
    - Sun StorEdge EBS Editions, consideration for, 384
    - volume management, 389
    - volumes, 388, 389, 390
  - sharing with NDMP support, 471
  - sleep function, setting, 375
  - slots
    - available, 392
    - changing, 392
    - cleaning, 395
    - depositing volumes, 572
  - STK-9840, 595
  - storage nodes, 230, 374, 375
  - TapeAlert, 397
  - testing, 367, 369

- tips, [411](#)
- troubleshooting, [562](#)
- viewing, [360](#)
- volumes
  - depositing, [410](#)
  - exporting, [409](#)
  - importing, [409](#)
  - inventorying, [407](#)
  - labeling, [402](#), [404](#)
  - mounting, [400](#), [401](#)
  - recycling, [412](#)
  - selecting, [400](#)
  - unmounting, [402](#)
  - withdrawing, [410](#)
- X11 libraries, [568](#)
- autoloader. *See* [autochanger](#)
- Autorestart attribute, [433](#)
- Autostart attribute, [253](#)
- Available Slots attribute, [371](#)

## B

- backup groups. *See* [groups](#)
- backup levels
  - 1-9, [273](#)
  - described, [273](#)
  - options, [273](#)
  - overriding, [272](#)
  - planning, [274](#) to [276](#)
  - types
    - consolidated, [267](#), [273](#), [277](#), [278](#), [279](#)
    - full, [273](#), [276](#), [277](#), [543](#)
    - incremental, [273](#)
    - level, [277](#)
    - skip, [273](#)
  - usage, [276](#), [277](#)
- backup schedules. *See* [schedules](#)
- backups
  - balancing resources, [251](#), [267](#)
  - bootstrap, [87](#), [174](#), [261](#)
  - client
    - licenses, [157](#)
  - client file index entries, [120](#)
  - clusters
    - requirements, [456](#)
  - commands, [136](#)
  - savepnpc program, [140](#)

- completed, [261](#)
- consolidated, [267](#)
- customization scripts, [137](#)
- customizing, [30](#)
- cycle, using levels, [266](#)
- daemons, [11](#)
- data movement, [6](#)
- directives, [144](#)
- enforcement of licenses during, [158](#)
- filesystems, [118](#), [269](#)
- force incremental, [256](#)
- groups. *See* [groups](#)
- incremental, [173](#)
- large client filesystems, [117](#)
- large filesystems, [269](#)
- levels, [277](#), [282](#)
- managing, [257](#)
- manual, [285](#), [286](#)
  - media database, [174](#)
  - policies, [132](#)
  - pools, [173](#)
- NDMP, [464](#), [477](#)
  - consideration with [Auspex](#), [501](#)
- NDMP (Network Data Management Protocol), [458](#)
- non-scheduled, [286](#)
- nonscheduled, [285](#)
- nsrndmp\_save command, using, [478](#)
- online indexes, [87](#), [277](#)
- open files, [262](#)
- operations
  - details, viewing, [259](#)
  - monitoring, [32](#)
  - starting, [258](#)
  - stopping, [258](#), [536](#)
  - verifying, [32](#)
- overriding scheduled, [30](#)
- performance when cloning, [340](#)
- pools, [172](#)
- previewing, [260](#)
- programs, [11](#)
- recoveries, [534](#)
- rejection, [537](#)
- restarting, [259](#)
- save sets. *See* [save sets](#)
- sequence, [271](#)
- server, forced, [545](#)
- sessions, [58](#)
- storage node, [86](#)

- time intervals, setting, [255](#)
- troubleshooting, [534](#)
- types of backups. *See* backup levels
- unattended, [172](#), [238](#)
- volumes
  - clones, [350](#)
  - selection criteria, [232](#), [233](#)
- Bar Code Reader attribute, [371](#)
- barcodes
  - labels, [404](#), [405](#), [406](#)
  - identical, [406](#)
  - troubleshooting, [404](#)
  - volumes
    - labeling, [404](#)
- block size configuration, [390](#)
- BlueArc
  - requirements, [505](#)
- bootstrap, [13](#)
  - backups, storage node, [87](#)
  - clone, [87](#)
  - defined, [70](#)
  - file, [261](#)
  - NDMP limitation, [463](#)
  - pools, [170](#), [175](#)
  - printing, [261](#)
    - failure, [545](#)
  - storage nodes, [87](#)
- boot-time file, [558](#)
- browsable save set status, [119](#)
- browse policies
  - about, [120](#), [122](#)
  - client file index growth, [89](#)
  - clones, storage nodes, [357](#)
  - creating, [130](#)
  - defined, [121](#)
  - editing, [131](#)
  - granularity, [135](#)
  - save set recoveries, [314](#), [315](#)
  - save set status, [119](#)
  - usage, [121](#)
- Business Edition, [18](#)
- button *See* specific button name, [41](#)

## C

- CAP (cartridge access port), [409](#)

- carousel. *See* autchanger
- cartridges, cleaning, [397](#)
- Change Server button, [41](#)
- changers program, [563](#)
- characters, double-byte, [43](#)
- CHKDSK, running, [287](#)
- cleaning cartridges, [397](#)
- Cleaning Delay attribute, [563](#)
- Cleaning Slots attribute, [371](#)
- client alias, changing, [541](#)
- client file index
  - browse policies, [97](#), [121](#), [122](#)
  - checking, [538](#)
  - clones and storage nodes, [357](#)
  - defined, [6](#), [88](#)
  - entries
    - adding, [90](#)
    - removing, [90](#), [96](#), [98](#), [99](#), [342](#)
  - growth, [89](#)
  - location, designating, [94](#)
  - managing size, [96](#)
  - messages, missing, [543](#)
  - operations
    - backups, [277](#)
    - checking, [90](#)
    - moving, [95](#)
    - recovery, [309](#), [312](#), [313](#)
  - policies, [96](#)
  - pools, [170](#), [175](#)
  - restoration, [314](#)
  - save sets
    - cycles, removing, [98](#)
    - entries, [125](#)
    - removing, [96](#)
  - size, [89](#)
    - management, [96](#)
    - notification, [538](#)
- Client ID, [113](#)
- Client resource, [8](#), [251](#)
  - NDMP, [474](#)
- Client Retries attribute, [253](#)
- client type not allowed, troubleshooting, [158](#)
- ClientPak, [18](#)
- clients
  - aliases, problems, [540](#)
  - backup commands, [136](#)
  - client ID

- creating new client, 318
  - defined, 113
- client/server communication errors, 550
- cloning
  - examples, 351, 352, 353
  - storage nodes, 87
- communication, configuration, 55
- configuration, 104, 110
- daemons, 10
- defined, 6, 104
- editing, 112
- groups, 250
- large filesystems, 269
- licensing, 157, 537
- manual backups, 285
- multiple, 116
- NDMP considerations, 474
- operations
  - adding, 32
  - archive retrieve, 333
  - backing up renamed clients, 544
  - backups, 117
  - creating, 110
  - deleting, 113
  - editing, 112
  - indexes, moving, 95
  - installation, 110
  - manual backups, 285
  - recovery, 287
  - recovery, failure, 544
- policies, multiple, 129
- port requirements, 78
- priority, 155
- products, 18
- programs, 10
- resource, 37
- save sets, 117, 118
- Solaris binary location, 581
- storage node affinity, 86
- unsynchronized client clock, 535

Clone Storage Nodes attribute, 87, 357

Clone Volumes dialog box, 343

cloning
 

- archives, 356
  - display errors, 559
- autochangers, shared, 388
- automatic, 350
- defined, 339
- destination volume, defined, 340
- device requirements, 340
- examples, 351, 352, 353, 354
- manual, 347
- NDMP, 481
- online indexes, storage nodes, 357
- pools
  - configuration, 174
- recovery
  - save sets, 355
  - volumes, 343
- recycling volumes, 342
- requirements, 347
- save sets, 345
  - automated, 350
  - manually, 351
  - methods, 347
  - performing, 347
- source volume, defined, 340
- storage nodes, 87
  - affinity, 357
  - online indexes, 357
- volume duplication, comparison, 341
- volumes, 339, 350
  - automatic, 350
  - creating, 343
  - deleting, 343
  - details, viewing, 345
  - process, 339
  - purging, 342
  - requirements, 340

clusters
 

- backup requirements, 456
- backups, 433, 435
- clients, 433
- configuration options, 440 to 445
- data corruption, 442
- defined, 430
- devices, 440
- jbconfig program considerations, 442
- licensing, 454, 455
- nodes, 430
- recovery types, 435
- requirements, 455
- tape libraries, 440

code set support, 578

command line Sun StorEdge EBS utilities, 3

commands. *See* specific command names

- Common Device Interface, 197
  - cdidisable file, 197
  - SCSI command, 197
  - supported and not supported platforms, 197
  - using with tape, 198
- configuration
  - NDMP Device resource, 473
  - order for resources, 30
  - planning, 35 to 38
- connection ports, 77, 78, 79, 81
- consolidated backups, 277, 278, 279
- consolidating save sets, 278
  - nrssc command, 281
- cross-platform
  - capability, 3

## D

- daemon.log, 34, 101
- DAR backups, 480
- DAS (Dual Attach Station), 418
- dasadmin program, 418
- data
  - See also* cloning
  - backing up, 29
  - life cycle
    - cloned data, 342
    - managing, 127
  - management, cross-platform, 3
  - movement during backup, 6
  - parallelism, 57
- data compression, 570
- datawheel. *See* autochanger
- DDS (dynamic drive sharing)
  - about, 198, 378
  - clustering, 446
    - robotics control, 452
  - configuring, 199
  - device, creating, 204
  - Devices dialog box, 203
  - drive sharing, termination, 204
  - enabling, 378
  - idle devices, 383
  - licensing, 206, 378
  - NDMP, 201
  - SCSI, limitation, 199

- sharing, termination, 382
- default backup schedules, 265
- Default Cleanings attribute, 371
- Deposit Timeout attribute, 563
- Device dialog box, 58
- Device resource, 8
- devices
  - See also* autochangers
  - auto media management, 226
  - cloning, requirements, 347
  - concurrent, 188
  - DDS, 199
  - defined, 187
  - disk label errors, 545
  - drivers, testing, 564
  - environment variables, 194
  - errors, limiting, 226
  - filesystem
    - configuration, 223
    - staging, 245
  - Idle Device Timeout attribute, 204
  - labeling errors, 545
  - licensing, 84
  - maintenance commands, 563
  - media
    - nonrewinding, 192
    - pools, 175
  - media errors, 393
  - multiple, 188
  - multiplexing, 57
  - nonrewinding, 238, 571
  - operations
    - adding, 192
    - creating, 192
    - deleting, 194
    - editing, 193
  - parallelism, 57
  - pools, 175
  - prerequisites, 84
  - remote, 228, 230
    - autochangers, 230
    - naming, 375
    - privileges, 229
    - See also* storage nodes
    - storage, 83
    - timeout settings, 230
  - sessions, 234
  - silos, 415, 420



- stand-alone, 6
- storage nodes, 83
  - autochangers, 85
  - configuration, 85, 228
  - silos, 85
- support
  - Sun StorEdge EBS edition type, 187
- supported devices, 188
  - maximum number of, 188
- taking offline, 393
- testing, 369
- troubleshooting, 562
- volume loading problems, 204
- volumes, clones, 340, 350

Devices attribute, 371

DHCP (dynamic host configuration protocol), 39, 552

diagnostic tools
 

- AIX core dumps, 595
- communications tests, 575
- Sun StorEdge EBS reports, 74

DinoStor TapeServer
 

- configuration, 506

directed recoveries, 302
 

- administration, 295

directed recovery
 

- access, 297
- advantages, 309
- defined, 294
- NDMP, 486
- performing, 304
- recover program, 301
- use of, 297

Directive resource, 8

directives
 

- ASM (application specific modules), 145
- creating, 153
- customizing, 153
- defined, 144
- deleting, 153
- editing, 153
- naming restrictions, 541
- preconfigured, 151

disaster recovery, 294

*Disaster Recovery Guide*, 23

DNS (Domain Name System), 552
 

- hosts file, 39

- documentation
  - Sun StorEdge EBS, 20
- double-byte characters, displaying, 43
- drag and drop
  - mounting, 401
- drives
  - See devices, 199
- dynamic drive sharing (DDS). See DDS
- dynamic host configuration protocol (DHCP), 552

## E

ECB (Event Control Block), 539

editing
 

- clients, 112
- notifications, 69

editions of Sun StorEdge EBS, 17, 18

Eject Sleep attribute, 563

element status, disabling, 413

EMASS silos, 418

EMC, application information values, 499

enabler codes, 157

environment variables
 

- ACI\_MEDIA\_TYPE, 418
- DAS\_CLIENT, 418
- DAS\_SERVER, 418
- devices, 194
- format, 101
- NSR\_DEV\_BLOCK\_SIZE\_MEDIA\_TYPE, 195, 391
- NSR\_DEV\_LOAD\_POLL\_INTERVAL\_MEDIA\_TYPE, 196
- NSR\_DEV\_LOAD\_TIME\_MEDIA\_TYPE, 196
- NSR\_DEV\_LOAD\_TRY\_TIMEOUT\_MEDIA\_TYPE, 196
- NSR\_DEV\_TAPE\_FILE\_SIZE\_MEDIA\_TYPE, 195
- NSR\_DEV-DEFAULT\_CAPACITY\_MEDIA\_TYPE, 196
- NSR\_MMDCONTROL, 86

error messages
 

- administrator's list, 59
- barcode not present, 424
- ClientPak, does not exist, 537
- clients, too many, 537
- copy violation, 546

- disk label, [545](#)
- environmental variables, [561](#)
- hardware address not found, [592](#)
- hostname permission denied, [557](#)
- illegal record size, [542](#)
- matching devices, [87](#)
- media verification, [538](#)
- media waiting, [168](#)
- no client resource, [540](#)
- no space left in tape, [596](#)
- nwrecover program, [542](#)
- RPC errors, [557](#)
- save sets, [543](#)
- scanner program, [543](#)
- server, unavailable, [557](#)
- xview, [547](#)

evaluation mode

- clusters, [455](#)

Event Control Block (ECB), [539](#)

events, notifications, [67](#)

## F

failover, save group requirements, [433](#)

file conversion, sparse to fully-allocated, [549](#)

file type device, [206](#)

filemarks, [570](#)

files

- relocating during recovery, [292](#)
- servers, [154](#)

filesystem devices, [245](#)

filesystems

- backups, large, [117](#)
- remote, saving, [557](#)

firewall support, [601](#)

- configuration, [76](#), [81](#)
- packet filtering, [75](#)

firmware, verification, [556](#)

Force Incremental attribute, [270](#)

force incremental attribute, [256](#)

foreign languages, [43](#)

full backups, [276](#)

## G

Grau silos, [418](#)

Grooming attribute, [330](#)

grooming files, [321](#)

Group Control button, [42](#)

Group Control Details dialog box, [33](#), [260](#)

Group Detail dialog box, [33](#)

Group resource, [8](#)

groups, [11](#), [251](#)

- attributes, [253](#)
- backups
  - management, [257](#)
  - previewing, [260](#)
- bootstrap, [261](#)
- client policies, multiple, [129](#)
- completed, [261](#)
- defined, [250](#)
- details, viewing, [259](#)
- filesystems, large, [269](#)
- force incremental, setting, [256](#)
- naming restrictions, [541](#)
- operations
  - creating, [254](#)
  - deleting, [257](#)
  - editing, [257](#)
  - previewing, [260](#)
  - restarting, [259](#)
  - starting, [258](#)
  - stopping, [258](#)
  - time intervals, setting, [255](#)
- overview, [249](#)
- pools, [251](#)
- types, [252](#)

## H

Hardware ID attribute, [203](#)

HBA (host bus adapter), [390](#)

help, online, [22](#)

hidden attributes, [42](#)

host bus adapter (HBA), [390](#)

host transfer affidavit, [51](#)

hosts file, [39](#)

hosts table, setting up, [551](#)

hpflip program, [564](#)

## HP-UX

- creating device files, [589](#)
- installing autochangers, [583](#), [588](#)
- pass-through drivers, [584](#) to [588](#)
- system actions, [589](#)
- system shutdown, [588](#)
- testing autochangers, [592](#)

## I

- IBM, silos, [417](#)
- Idle Device Timeout attribute, [204](#), [383](#)
- ielem program, [564](#)
- Inactivity Timeout attribute, [253](#)
- incremental backup
  - NDMP limitation, [465](#)
- Index Save Sets dialog box, [92](#)
- indexed archives, [324](#)
- Indexes button, [41](#)
- Indexes dialog box, [91](#)
- inquire program, [372](#), [468](#), [564](#)
- Instances Details window, [93](#)
- interface
  - customizing Sun StorEdge EBS Administrator program, [42](#)
  - foreign languages, [43](#)
  - Sun StorEdge EBS Administrator program, [39](#)
- Interval attribute, [253](#)
- inventory, volume
  - silos, [426](#)
- inventorying, autochangers. *See* autochangers
- ioscan program, [583](#)
- IP
  - errors
    - hosts table, [551](#)
    - troubleshooting, [550](#)
  - name search order, setting, [553](#)

## J

- jbconfig program, [230](#)
  - adding autochangers, [367](#)
  - autochangers, [385](#), [386](#)
  - clusters, [442](#)
  - device naming, [386](#)

- hanging, [567](#)
- silos, [417](#)
- storage nodes, [85](#), [374](#)
  - administrator requirements, [229](#)
  - installation, [229](#)
- X11 libraries, [568](#)
- jbexercise program, [367](#)
  - device drivers, testing, [564](#)
  - failure, [566](#)
  - NDMP, not supported, [464](#)
- jbverify program, [368](#)
- Jukebox Features attribute, [562](#)
- Jukebox resource, [371](#), [375](#)
- jukebox. *See* autochanger

## L

- Label button, [41](#)
- Label Template resource, [8](#)
- label templates, [180](#), [181](#)
  - attributes, [182](#) to [184](#)
  - components, [185](#)
  - creating, [186](#)
  - deleting, [187](#)
  - editing, [186](#)
  - naming restrictions, [541](#)
  - naming strategies, [184](#)
  - number sequences, [184](#)
  - types
    - default, [180](#)
    - usage, [180](#), [232](#)
- labeling
  - autochanger considerations, [400](#)
  - autochangers, [402](#)
  - barcodes, [405](#), [406](#)
    - identical barcodes, [405](#)
    - matching names, [406](#)
    - types, [406](#)
  - consistently with templates, [403](#)
  - matching names, [403](#)
  - volumes, [180](#), [231](#)
    - barcodes, [404](#)
    - how to label a volume, [28](#)
    - silos, [423](#)
    - tips, [185](#)
- languages, foreign, [43](#)

- ldunld program, 564
- library. *See* autochanger
- licensing
  - archiving, 321
  - autochangers, 359, 384
    - sharing, 384
  - clients
    - backup rejection, 537
    - errors with, 537
  - clusters, 454
  - copy violation, 546
  - storage nodes, 84
    - Sun StorEdge EBS clients, 157
- list command, 498
- Load Sleep attribute, 376, 563
- local backups with NDMP
  - introduction, 459
- locale support, 577
- log files, 74
  - cloning information, 340
  - daemon.log, 34
  - log directory, 27
  - management, 100
  - messages file, 34
  - system log, 262
- logical backup volumes, 156
- lrescan program, 563
- lreset program, 563
- lsdev program, 590
- lusbinfo program, 563
- lusdebug program, 563, 564, 565

## M

- Mac OS X, 159
- man pages, 22
- manual backups. *See* backups, manual, 285
- Manual Recycle attribute, 240
- Manual Restart option, 259, 433
- manuals, 20
- Match Bar Code Labels attribute, 371
- media
  - auto media management, 226
  - handling errors, 225
  - nonrewinding, 192
  - silos, 421
  - supported devices, 188
- media database
  - backups, 277
  - cloned data, 342
  - clones and storage nodes, 357
  - compression, 97, 100
  - defined, 6
  - entries, removing, 98, 99
    - clones, 342
  - managing size, 88, 97
  - restoration, 315
  - retention policies, 121
  - save sets, entries, 125
  - silo inventory, 426
- media position errors, 538
- memory
  - NDMP requirements, 465
- message logs
  - files, 74
  - management, 100
  - reducing size, 101, 102
- message windows, 40
- messages file, 34
- Mirapoint
  - requirements, 504
- mminfo program, 71
  - program-readable reports, 71
  - reports, 135
- mmlocate program, 71
- mmrecov program, 87
- modules
  - NDMP, 476
- Monitor RAP attribute, 74
- Mount button, 41
- mounting volumes, 28
- moving StorEdge Backup Server to a different
  - computer, 51
- msense program, 564
- multiplexing
  - NDMP, 464
  - performance issues, 4
  - save sets, 57

## N

### name

- labeling, [403](#)

- name resolution, [552](#)

- name servers, disabling, [552](#)

- naming requirements, [556](#)

- naming restrictions, [541](#)

### NDMP (Network Data Management Protocol)

- Force Incremental setting, [476](#)

- group Parallelism setting, [476](#)

- silo support, [473](#)

### NDMP (Network Data Management Protocol), [473](#)

- application information variables, [475](#)

- archives, [464](#)

- attributes, [474](#)

#### Auspex

- autochanger handle, [502](#)

#### autochangers

- configuration, [469](#)

- list command, [498](#)

- sysconfig program, [493](#)

- backups, [477](#)

- command line options, [478](#)

- parallel, [464](#)

- bootstrap limitation, [463](#)

- clients, [474](#)

- definitions, [464](#)

- cloning, [481](#)

#### configuration

- options, [459](#)

- Sun StorEdge EBS server, [466](#)

- connection software, [458](#)

- DAR backup, [480](#)

- DDS considerations, [201](#)

- device pathname, [468](#)

- directed recoveries, [486](#)

#### EMC Celerra

- autochanger handle, [498](#)

#### EMC IP4700

- autochanger handle, [499](#)

- incremental backup limitation, [465](#)

- inquire program, [468](#)

- installation, [458](#)

- multiplexing, [464](#)

#### NetApp

- autochanger handle, [493](#)

- nsr resource attributes, [465](#)

- nsrndmp\_recover command

- syntax, [483](#), [488](#)

#### operations

- browsing, [464](#)

- recoveries, [482](#)

- recoveries, directory level, [484](#)

- recoveries, save set, [482](#)

- overview, [457](#)

- pool requirements, [464](#)

- prerequisites, [477](#)

- requirements, [463](#)

- save set operations

- consolidation, [464](#)

- recovery, [465](#)

- scanner program, [464](#)

- sharing autochangers, [471](#)

#### Sun StorEdge EBS

- changes with NDMP, [464](#)

- configuring, [466](#)

- support, [458](#)

### NDMP (Network Data Management Protocol)

- memory requirements, [465](#)

- protocol limitations, [465](#)

- sharing autochangers, [471](#)

- near-line storage. *See* autochanger

- Network Edition, [18](#)

- parallelism, [58](#)

### nonindexed

- archives, [325](#)

- Notification resource, [8](#)

### notifications

- autochangers, [411](#)

- defined, [65](#)

- events, [67](#)

#### operations

- creating, [69](#)

- customizing, [66](#)

- deleting, [70](#)

- editing, [69](#)

- preconfigured, [65](#) to [66](#)

- priorities, [68](#)

- programs, [67](#)

#### SNMP

- configuring, [512](#)

- creating, [514](#)

- modifying, [513](#)

- nsrtrap, [512](#)

- nsr\_getdate program, [133](#), [134](#)

- nsr\_shutdown program, [102](#), [299](#), [539](#)

- nsradmin program, [38, 44](#)
  - editing the nsrla.res database, [115](#)
- nsrarchive program, [323](#)
- nsrck program, [9, 313, 543](#)
- nsrd daemon, [9, 11, 12, 15, 102](#)
  - tasks, [10](#)
- nsrexec daemon
  - multiple processes, [534](#)
  - troubleshooting, [534](#)
- nsrexec program, [9](#)
- nsrexecd daemon, [10, 12, 102](#)
  - functionality, [555](#)
- nsrim daemon, [13](#)
- nsrim program, [9](#)
- nsrindexd daemon, [10, 15](#)
- nsrinfo
  - program-readable reports, [71](#)
- nsrjb program, [388, 389](#)
  - autochangers, shared, [388, 389](#)
  - cleaning cartridges, [397](#)
  - inventories, [572](#)
  - o option, errors, [573](#)
  - port number restriction, [572](#)
  - storage nodes, [374](#)
  - troubleshooting, HP-UX, [594](#)
- Nsrjb Release attribute, [562](#)
- nsrla.res database, [115](#)
- nsrlic program, [157](#)
- nsrmm program, [97, 133, 315, 539](#)
  - defined, [71](#)
- nsrmmmd daemon, [10, 13, 14, 15, 16, 196](#)
  - NDMP
    - unsupported options, [465](#)
  - Nsrmmmd Control Timeout attribute, [85](#)
- nsrmmdbd daemon, [10, 14, 15](#)
- nsrndmp\_recover program, [482, 484](#)
- nsrndmp\_save program, [477](#)
  - backing up data, [478](#)
  - command line options, [478](#)
- nsrports program, [81](#)
- nsrssc program, [281](#)
- nsrtrap, [512](#)
  - command line options, [512](#)
  - verbose mode, [513](#)
- nsrwatch program, [41](#)

- nwadmin program, starting the Sun StorEdge EBS
  - Administrator program, [28](#)
- nwarchive program, [48, 329](#)
- nwbackup program, [45, 286, 287](#)
- nwrecover program, [46, 287](#)
  - browse policy, [121](#)
- nwretrieve program, [49, 334, 335, 559](#)

## O

- online help, described, [22](#)
- online indexes
  - defined, [6](#)
  - entries
    - checking, [90, 538](#)
    - removing, [90, 98, 346](#)
  - information
    - viewing, [90](#)
  - management
    - manual, [88](#)
    - size, [96](#)
  - moving, [95](#)
  - recovery
    - location, [540](#)
  - restoration, [314](#)
  - save sets, viewing, [92](#)
  - size considerations, [89](#)
  - volumes
    - removing, [98, 99, 242](#)
- open files, backing up, [262](#)
- open files, backing up with VSS, [263](#)

## P

- packet filtering, [75](#)
- parallelism
  - defined, [57](#)
  - maximum, [57](#)
  - performance, [4](#)
- performance
  - cloning, [340](#)
  - features, [4](#)
  - multiplex, [57](#)
  - parallelism, [57](#)
  - target sessions, [58](#)
- permissions

- Archive feature, 325 to 326
- group, 59
- management of, 59
- storage nodes, 85
- user, 59
- pmode program, 564
- policies
  - backups, manual, 132
  - browse
    - about, 120, 122
    - data life cycle, 127
    - defined, 121
    - modifying, 133
    - usage, 121
  - clients, 129
  - creating, 130
  - data life cycle, 127
  - deleting, 131
  - editing, 131
  - granularity, 135
  - manual backups, modifying, 134
  - multiple, 129
  - naming restrictions, 541
  - overriding, 135
  - planning, 268
  - retention, 124, 125
    - about, 120, 121
    - data life cycle, 127
    - defined, 121
    - modifying, 133
    - usage, 121
    - volume relabeling, 121
- Policy resource, 8
- polling interval, 42
- Pool resource, 8
- Pool Type attribute, with archive pools, 327
- pools
  - archive, 174
    - creating, 180
    - errors, 560
  - auto media verification, 178
  - autochangers, 412
  - bootstrap, 170, 175
  - client file index, 170, 175
  - clones, 174
  - configuration, 173
    - archive, 174
    - clone, 174
    - criteria, 171, 172
    - devices, 175
    - incremental backups, 173
    - label templates, 181
    - manual backups, 173
  - conflicts, 171
  - default, 168, 174, 180
    - clone pool, 174
  - defined, 166
  - devices, 175
    - media, 175
  - expression matching, 170, 177
  - groups, 251
  - label templates, 181
  - NDMP requirement, 464
  - operations
    - creating, 176
    - deleting, 179
    - editing, 179
    - volumes, labeling, 180
  - preconfigured, 167
  - restrictions, 170, 177
  - save set consolidation, 282
  - save sets, 166
  - types, 167
  - usage, 169, 232
- portmappers
  - nsrexecd daemon, 555
  - registration, 82
  - verifying, 554
- ports
  - clients, 78
  - configuration, 81
  - configuring, 81
  - connection
    - configuration, 81
    - requirements, 78, 79
  - default settings, changing, 80
  - restricting range, 76
  - server requirements, 77, 78, 79
  - service, configuration, 81
  - storage node requirements, 79
  - troubleshooting, 83
- Power Edition, 18, 455
  - parallelism, 58
- priorities, notifications, 68
- Procom NetFORCE
  - requirements, 504

program-readable reports, [71](#)  
programs. *See* specific program name, [45](#)

## R

read-only volume mode, [235](#)

recover program, [10](#)  
    clusters, using with, [436](#)  
    media database, [125](#)  
    nsrmmd interaction, [16](#)  
    options, [301](#)  
    retention policy, [125](#)  
    suspect, [119](#)

recoverable save set status, [119](#)

recoveries, [287](#), [288](#)  
    archives, [333](#)  
    autochangers, sharing, [386](#)  
    client file index, [313](#)  
    clients, renamed, [544](#)  
    clone volumes, [343](#)  
    cluster, [435](#)  
    components interaction, [13](#)  
    conflicts, [288](#) to [291](#)  
    daemons, [14](#)  
    directed, [302](#)  
        access, [297](#)  
        administration, [295](#)  
        advantages, [309](#)  
        defined, [294](#)  
        recover program, [301](#)  
        usage, [297](#)  
    disaster-related, [294](#)  
    failure, [542](#)  
    files, finding for recovery, [293](#)  
    index-based  
        advantages, [309](#)  
    interrupted backups, from, [543](#)  
    methods, comparing, [309](#)  
    NDMP, [482](#), [484](#)  
        save set recoveries, [482](#)  
    planning, [268](#)  
    program, [14](#)  
    relocation, [292](#)  
    save set recovery, [303](#)  
    save sets  
        client file index, [314](#)  
        media database, [315](#)

    type, determination, [311](#)  
    volumes  
        required, [294](#)

recyclable save set status, [119](#)

recycling. *See* volumes

refresh, server information, [42](#)

Registration resource, [8](#)

re-hosting, [51](#)

relem program, [564](#)

remote access, [288](#)  
    recoveries, [558](#)

Remote Access list, [558](#)

remote archives, failure, [559](#)

remote filesystems, saving, [557](#)

Remote NDMP backups with DinoStor TapeServer  
    introduction, [462](#)  
    performing, [507](#)

remote procedure call (RPC), [9](#)

Remove Oldest Cycle button, [92](#)

reports, [70](#)  
    clones, [350](#)  
    daemon.log, [101](#)  
    data movement operations, [70](#)  
    message logs  
        files, [74](#)  
        management, [100](#)  
        reducing size, [101](#)  
    monitoring changes to Sun StorEdge EBS  
        resources, [74](#)  
    program-readable, [71](#)  
    rap.log file, [74](#)  
    save set policies, [135](#)  
    savegrp completion, [13](#)  
    server statistics and diagnostics, [74](#)  
    storage management application, [71](#)  
    system console log, [262](#)

resources  
    Autochanger, [8](#)  
    Client, [8](#)  
    configuration order, [30](#)  
    configuring for NDMP, [466](#)  
    defined, [7](#) to [9](#)  
    Device, [8](#)  
    dialog boxes, [44](#)  
    Directive, [8](#)  
    Group, [8](#)  
    Label Template, [8](#)



- Notification, 8
- Policy, 8
- Pool, 8
- Registration, 8
- Schedule, 8
- Server, 8
- Staging, 8
- retention policies, 125
  - about, 120, 121, 124
  - cloned volumes, 346
  - clones, storage nodes, 357
  - creating, 130
  - defined, 121
  - editing, 131
  - granularity, 135
  - read-only volume mode, 235
  - save set status, 119
  - usage, 121
  - volumes, relabeling, 121
- retrieval
  - troubleshooting, 559
- Retrieve Status window, 334
- retrieving archives, 333
- retrieving save sets, 334
- RPC (remote procedure call), 9
- rpcbind program, 555
- rpcinfo program, 81, 554

## S

- SAN (storage area network), 384
- save program, 11, 134, 137, 436
- save set
  - All* save set, 107
- Save Set Clone dialog box, 348
- save set consolidation, 278
- Save Set Cutoff option, unsupported, 536
- save sets
  - archives, retrieving, 334
  - backup commands, 136
  - backups, 117, 118
    - filesystem, 118
    - status, 260
  - client combination, 117
  - client file index, entries, 125
  - client priority, 155

- clones, 347
- cloning
  - automatically, 352
  - described, 345
  - manually, 351
- consolidation, 278
  - files, large, 279
  - limitations, 280
  - NDMP, 464
  - nsrssc command, 281
  - pools, 282
  - usage, 278, 279
- defined, 5
- defining, 118
- dependencies, 342
- details, viewing, 93
- indexes, viewing, 92
- information, viewing, 92, 237
- load balancing, 117
- media database, entries, 125
- multiplexing, 57, 59
- policies
  - modification, 133
  - reports, 135
- pools, 166
- predefined, 106
- recoveries, 303, 305, 306
  - browse period, 304
  - client file index, 314
  - conflicts, 307
  - entire, 308
  - files, 309
  - media database, 315
  - NDMP, 465
  - omitting, 120
  - online indexes, 314
  - relocating, 306
  - requirements, 304
  - usage, 303
  - volumes, 306
- recovery
  - compared to index-based recovery, 311
- staging. *See* staging
- status
  - completed, 260
  - failed, 260
  - manually changing, 119
  - pending, 260
  - retention policy, 124

- values, 118
  - suspected, 355
- Save Sets dialog box, 92
- savefs program, 11, 12
- Savegroup Completion Report, 70, 261
- savegrp program, 11, 12
  - backup limitation, 546
  - clients, 12
  - defined, 10
  - reports, 70
  - save set consolidation from command line, 281
  - stopping, 544
- savenpc program
  - message logging by, 143
  - using with customized backup program, 140
- scalability, Sun StorEdge EBS, 5
- SCANDISK, running, 287
- scanner program, 229, 308, 309, 315
  - deleted clients, 113
  - missing client file index, 543
  - NDMP, 464
  - options, 316
  - record size, 541
  - recovering clone volumes, 343
  - recovering save sets from volumes, 127
  - recreating online indexes, 306
  - recreating volume entries, 306
  - retention policy, 125
  - scann save set status, 119
  - storage nodes, 85
  - volumes, read-only, 539
- Schedule resource, 8, 250
- scheduled backups
  - Manual Restart option, 259
  - performing, 29
- schedules, 263
  - attributes, 269
  - backup cycle, 266
  - balancing, 267
  - creating, 271
  - default, 265, 266
  - deleting, 272
  - editing, 272
  - large filesystems, 269
  - load balancing, 117
  - naming restrictions, 541
  - overriding, 272
  - overview, 249
  - planning, 268
  - preconfigured, 265
  - sequence, 271
  - staggered, 267
  - usage, 264
- Schedules window, 270
- SCSI
  - troubleshooting, 565, 566
- SCSI address selection for HP-UX, 583
- server port requirements, 77
- servers
  - adding servers for a client, 154
  - administrators
    - adding, 60
    - defined, 59
    - privileges, 59
  - autochangers
    - viewing, 360
  - avoiding copy violation errors, 51
  - changing, 54
  - changing servers for a client, 154
  - client/server communication errors, 550
  - clients
    - adding, 32
    - license enforcement, 157
    - tasking rights, 154
  - cloning and storage nodes, 87
  - configuration, 55
  - connectivity, 53, 54
  - defined, 5
  - devices
    - volume management, 389
  - errors, binding to, 557
  - file, 154
  - index
    - backup, failure, 545
    - management, 89
  - maintenance, 100
  - notifications
    - creating, 69
    - deleting, 70
    - editing, 69
    - priorities, 68
  - parallelism, 57
  - ping command, testing, 553
  - port requirements, 77
  - refresh rate, 42

- Server resource, 8
- servers file, 33, 154
- setting up, 55
- storage nodes
  - affinity, 86
  - converting, 229
- target sessions, 58
- testing, 553
- verifying, 33
- service mode, 393
- service ports
  - configuration, 81
  - requirements, 78
  - storage nodes, 79
- session management, 4
- Set Up Server dialog box, 57
- Shared Devices attribute, 203
- Silo Software Module, 19
- silos, 230
  - defined, 415
  - devices, 415
    - cleaning, 420
    - defining, 419
    - releasing, 419
  - EMASS/Grau, 418
  - IBM 3494, 417
  - installation, 415
  - inventory, 426
  - jbconfig program, 417
  - media management, 421
  - naming conventions, 417 to 418
  - NDMP support, 473
  - slot numbers, 422, 426
  - storage nodes, 230
  - StorageTek, 417
  - tape library interface (STLI), 415
  - usage, 415
  - volumes, 415, 422, 423, 424, 425
- Simple Network Management Protocol. *See* SNMP
- sjiilm program, 564
- sjiinqm program, 564
- sjirdp program, 564
- sjirdtag program, 564
- sjirelem program, 564
- sjirjc program, 564
- SnapImage
  - autochanger handle, 507
- SNMP (Simple Network Management Protocol), 19, 511
  - configuring, 512, 515
  - defined, 511
  - notifications
    - configuring, 512
    - creating, 514
    - modifying, 513
  - nsrtrap, 512
  - traps, 511
- Solaris
  - troubleshooting, 581
- sparse files, converting, 549
- staging
  - described, 244
  - filesystem devices, 245
  - policies
    - creating, 246
    - deleting, 248
    - editing, 248
- Staging resource, 8
- stand-alone devices, defined, 6
- Start Time attribute, 253
- status values, scanned-in, 317
- STL (silos tape library), 415
- STLI. *See* silos
- storage area network (SAN), 384
- storage nodes
  - administrators, 85
  - affinity
    - clones, 357
    - defined, 86
    - problems, 561
  - autochangers, 85, 230, 374
  - backup clients to, 86
  - bootstrap backup, 87
  - clones, 357
  - cloning, 87
  - configuration, 85, 228
  - daemons, 15, 16
  - defined, 83, 228
  - devices, 85
    - defining, 374
    - volume management, 389
  - installation, 229
  - licensing, 84

- operations, [15, 16](#)
- parallelism, [57](#)
- port requirements, [79](#)
- prerequisites, [84](#)
- privileges, [229](#)
- programs, [15, 16](#)
- servers, converting from, [85, 229](#)
- silos, [85, 230](#)
- timeout settings, [230](#)
- timeouts, [561](#)
- troubleshooting, [87, 561](#)
- volumes, clones, [350](#)

StorageTek silos, [417](#)

Store Index Entries attribute, with archive pools, [327](#)

Sun cluster support, [432](#)

Sun StorEdge EBS

- administration interfaces, [38](#)
- archiving. *See* archiving
- backups. *See* backups
- Business Edition, [18](#)
  - autochanger sharing, [384](#)
- client license
  - enforcement, [157](#)
- clients
  - license enforcement, [157](#)
  - products, [18](#)
  - See also* clients
- command line usage, [3](#)
- configuration planning, [35](#)
- cross-platform capability, [3](#)
- daemon processes, graphics, [11 to 16](#)
- daemons, [9 to 10](#)
- documentation, [24](#)
- editions, [17, 18, 58](#)
  - device support, [187](#)
  - parallelism, [58](#)
- firewall support, [75](#)
- functionality, [5, 9](#)
- getting started, [25](#)
- interfaces, [45](#)
  - hidden attributes, [42](#)
  - nwarchive program, [48](#)
  - nwbackup program, [45](#)
  - nwrecover program, [46](#)
  - nwretrieve program, [49](#)
- introduction to, [2](#)
- NDMP
  - configuring, [466](#)
  - effects on, [464](#)
- Network Edition, [18](#)
  - autochanger sharing, [384](#)
  - defined, [18](#)
  - parallelism, [58](#)
- optimizing performance, [4](#)
- performance, [57](#)
- Power Edition, [18](#)
  - autochanger sharing, [384](#)
  - defined, [18](#)
  - parallelism, [58](#)
- preparing to use, [25](#)
- privileges, [39](#)
- products, [17, 19](#)
- programs, [9 to 10](#)
- remote access, [288](#)
- reports. *See* reports
- scalability, [5](#)
- servers
  - defined, [5](#)
- settings, preconfigured, [35](#)
- startup commands, [535](#)
- storage nodes. *See* storage nodes
- target sessions, [58](#)
- use
  - planning, [35](#)
  - preparation, [25](#)
- verifying processes, [33](#)
- volumes. *See* volumes
- WorkGroup Edition, [17](#)
  - autochanger sharing, [384](#)
- Workgroup edition
  - parallelism, [58](#)

Sun StorEdge EBS Administrator program, [38, 53](#)

- customizing interface, [42](#)
- getting started, [39](#)
- interface
  - described, [39](#)
  - foreign languages, [43](#)
- introduction, [3](#)
- message windows, [40](#)
- privileges, [39](#)
- server information area, [40](#)
- starting, [28](#)
- toolbar, [41](#)
- xview, errors, [547](#)

Sun StorEdge EBS command line utilities, [3](#)

- Sun StorEdge EBS DiskBackup, 206
  - advanced file type concurrent operations, 211, 212
  - advanced file type device, 206
  - converting file type devices to Sun StorEdge EBS 7.0, 207
  - converting to Sun StorEdge EBS 7.0 advanced file type system, 207
  - file size limits, 207, 208
  - file type and advanced file type differences, 209, 209 to 210
  - File Type Device, 206
  - file type device, 206
  - notification, 208
  - retention policy, 208
  - supported configurations, 216
- Sun StorEdge Enterprise Backup Software 7.0 Administrator's Guide*, 20
- Sun StorEdge Enterprise Backup Software 7.0 Installation Guide*, 21
- suspected save sets, 355
- swap space
  - NDMP requirements, 465
- sysconfig command, 493

## T

- TapeAlert, 397
- tapeexercise program, 567
- target sessions, 58
- TCP/IP
  - network requirements, 39
- technical support troubleshooting
  - information, 531
- The, 316
- three-party NDMP backups
  - introduction, 459
  - performing, 479
- tips. *See* troubleshooting
- toolbars
  - Sun StorEdge EBS Administrator program, 41
- tracking
  - cloned data, 342
  - online index information, 89
- traps
  - categories, 516

- SNMP, 511
- troubleshooting
  - AIX
    - core dump, 595
    - STK-9840, 595
  - archive pools, 560
  - archive requests
    - naming, 560
  - archives, 559
    - display errors, 559
    - multiple save sets, 559
    - nsrarchive program, 560
    - remote request failure, 559
  - auto media verification, 538
  - autochangers, 562, 565
    - AIX considerations, 595
    - ATL 2640 libraries, 571
    - attributes, 562
    - autodetected scsi errors, 567
    - control port access, 571
    - drivers, 566
    - HP-UX considerations, 589, 592, 593
    - inventorying, 568
    - maintenance, 563
    - X11 libraries, 568
  - backups, 534
  - backups levels, 543
  - backups, stopping, 536
  - barcodes, 404
  - bootstrap printing, failure, 545
  - client file index
    - messages, missing, 543
    - size growth, 538
  - clients
    - alias, 540
    - licensing, 537
    - Solaris, location, 581
    - unsynchronized errors, 535
  - clones, display errors, 559
  - daemons, 534
  - devices, 562
    - drivers, testing, 564
    - maintenance, 563
    - nonrewinding, 571
  - disk label errors, 545
  - display errors, 559
  - ECB counter, 539
  - file conversion, 549
  - firmware, 556

- hosts table, [551, 553](#)
- HP-UX
  - error messages, [592](#)
  - unloading drives, [594](#)
  - unsupported media, [593](#)
- illegal characters, [541](#)
- IP errors, [550, 551](#)
- IRIX, [596](#)
- jbexercise program, failure, [566](#)
- licensing, copy violation, [546](#)
- lusdebug program, [565](#)
- name resolution, [552](#)
- name servers, disabling, [552](#)
- naming requirements, [556](#)
- nsrexec processes, [534](#)
- nsrexecd program, [555](#)
- nsrjb program, [573](#)
- online indexes, [540](#)
- packet receive buffer, [539](#)
- portmappers, verifying, [554](#)
- ports, [83, 555](#)
- recoveries, [534](#)
  - interrupted backups, [543](#)
  - new installation, [542](#)
  - online indexes, [540](#)
  - remote access, [558](#)
- remote filesystems, saving, [557](#)
- renamed client backups, [544](#)
- retrievals, [559](#)
- routers, [556](#)
- rpcinfo command, [554](#)
- scanner program, [539, 541](#)
- SCSI considerations, [565, 566](#)
- server errors, binding to, [557](#)
- server index, [545](#)
- Solaris, [581](#)
- storage nodes, [87, 561](#)
- tapeexercise program, [567](#)
- volumes, capacity issues, [569](#)
- X11 libraries, [568](#)
- xview errors, [547](#)
- tur program, [564](#)

## U

- uasasm program, [308, 540](#)
- uninventoried volumes, mounting on
  - autochanger, [401](#)

- Unload Sleep attribute, [563](#)
- Unmount button, [41](#)
- unsynchronized clock errors, [535](#)
- user groups
  - customizing, [62](#)
  - preconfigured, [61](#)
- user ID, sending bootstrap to multiple, [262](#)

## V

- vendor specific ASC and ASCQ codes, [576](#)
- verification, auto media, [178](#)
- viewing autochangers, [360](#)
- virtual servers. *See* clusters
- volume pool
  - See also* pools
  - archive, [323](#)
- volumes
  - archive, [356](#)
  - auto media management, [226](#)
  - auto media verification, [178](#)
  - autochanger
    - barcodes, [406](#)
  - autochangers. *See* autochangers
  - bad, removal, [242](#)
  - barcodes, labeling, [404, 406](#)
  - capacity issues, [569](#)
  - client file index
    - removing, [98, 99](#)
  - cloning, [339](#)
    - archive data, [356](#)
    - creating, [343](#)
    - performance, [340](#)
    - recovery, [343](#)
  - duplicating, problems with, [341](#)
  - Idle Device Timeout attribute, [204](#)
  - information
    - viewing, [236](#)
  - labeling, [180, 231, 232, 239, 240](#)
    - barcodes, [404, 405, 406](#)
    - how to, [28](#)
    - maximum size, [185](#)
    - time, [234](#)
    - tips, [185](#)
    - unmatching labels, [406](#)
  - loading problems, [204](#)
  - logical backups, [156](#)

- management
  - server devices, [389](#)
  - storage nodes, [389](#)
- marking as full, [242](#)
- modes, [234](#)
  - changing, [241](#)
  - impact on cloned data, [342](#)
  - types, [235](#)
- mounting, [28](#), [238](#), [239](#), [400](#), [401](#)
  - Idle Device Timeout attribute, [204](#)
- names
  - finding, [241](#)
  - reuse, [405](#)
- nonrewinding, [571](#)
- operations, [236](#)
- recoveries, required for, [294](#)
- recycling, [97](#)
  - autochangers, [400](#)
  - comparison, [412](#)
- relabeling, [228](#), [233](#)
- removing, [97](#)
- save sets
  - information, [237](#)
  - recoveries, [306](#)
- selection criteria, [232](#)
- silos
  - inventory, [426](#)
- silos. *See* silos
- slots, depositing, [572](#)
- status, [234](#)
- unlabeled, [227](#)
- unmounting, [240](#)
- Volumes button, [42](#)
- Volumes window, [237](#)

## W

- WorkGroup Edition
  - parallelism, [58](#)
- writebuf program, [564](#)

## X

- xview, errors, [547](#)

