



Solstice Backup™ 6.0 Administrator's Guide

Sun Microsystems, Inc.
901 San Antonio Road
Palo Alto, CA 94303
U.S.A. 650-960-1300

Part No. 806-5607-10
August 2000, [Revision A](#)

[Send comments about this document to: docfeedback@sun.com](mailto:docfeedback@sun.com)

Copyright 2000 Sun Microsystems, Inc., 901 San Antonio Road • Palo Alto, CA 94303-4900 USA. All rights reserved.

This product or document is protected by copyright and 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. For Netscape Communicator™, the following notice applies: Copyright 1995 Netscape Communications Corporation. All rights reserved.

Sun, Sun Microsystems, the Sun logo, AnswerBook2, docs.sun.com, Solstice Backup, Java, and Solaris 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 and Sun™ Graphical User Interface was developed by Sun Microsystems, Inc. for its users and licensees. Sun acknowledges the pioneering efforts of Xerox in researching and developing the concept of visual or graphical user interfaces for the computer industry. Sun holds a non-exclusive license from Xerox to the Xerox Graphical User Interface, which license also covers Sun's licensees who implement OPEN LOOK GUIs and otherwise comply with Sun's written license agreements.

RESTRICTED RIGHTS: Use, duplication, or disclosure by the U.S. Government is subject to restrictions of FAR 52.227-14(g)(2)(6/87) and FAR 52.227-19(6/87), or DFAR 252.227-7015(b)(6/95) and DFAR 227.7202-3(a).

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 2000 Sun Microsystems, Inc., 901 San Antonio Road • Palo Alto, CA 94303-4900 Etats-Unis. Tous droits réservés.

Ce produit ou document est protégé par un copyright et distribué avec des licences qui en restreignent l'utilisation, la copie, la distribution, et la décompilation. Aucune partie de ce produit ou document ne peut être reproduite sous aucune forme, par quelque moyen que ce soit, sans l'autorisation préalable et écrite de Sun et de ses bailleurs de licence, s'il y 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. La notice suivante est applicable à Netscape Communicator™: Copyright 1995 Netscape Communications Corporation. Tous droits réservés.

Sun, Sun Microsystems, the Sun logo, AnswerBook2, docs.sun.com, Solstice Backup, Java, et Solaris 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 et Sun™ a été développée par Sun Microsystems, Inc. pour ses utilisateurs et licenciés. Sun reconnaît les efforts de pionniers de Xerox pour la recherche et le développement du concept des interfaces d'utilisation visuelle ou graphique pour l'industrie de l'informatique. Sun détient une 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.

CETTE PUBLICATION EST FOURNIE "EN L'ETAT" ET AUCUNE GARANTIE, EXPRESSE OU IMPLICITE, N'EST ACCORDEE, Y COMPRIS DES GARANTIES CONCERNANT LA VALEUR MARCHANDE, L'APTITUDE DE LA PUBLICATION A REPENDRE A UNE UTILISATION PARTICULIERE, OU LE FAIT QU'ELLE NE SOIT PAS CONTREFAISANTE DE PRODUIT DE TIERS. CE DENI DE GARANTIE NE S'APPLIQUERAIT PAS, DANS LA MESURE OU IL SERAIT TENU JURIDIQUEMENT NUL ET NON AVENU.



Contents

Preface xxxv

1. Introduction 1

Solstice Backup Storage Management 1

 Cross-Platform Management 2

 Performance 3

 Ease of Use 3

 Scalability 4

How the Solstice Backup Software Works 4

 Solstice Backup Resources and Attributes 5

 Solstice Backup Daemons and Programs 8

How the Solstice Backup Software Backs Up Data 10

How the Solstice Backup Software Recovers Data 13

 Solstice Backup Operations with Storage Nodes 14

Solstice Backup License Modes 16

 Evaluation Mode 17

 Enabled Mode 17

 Authorized Mode 17

 Solstice Backup Power Edition Evaluation Enablers 18

| | |
|--|----|
| Authorization of Update Enablers | 18 |
| Rehosting a Solstice Backup Server | 18 |
| Solstice Backup Product Family | 19 |
| Server Software Products | 19 |
| Client Software Products | 19 |
| Other Products | 20 |
| Solstice Backup Documentation | 20 |
| Documentation Suite CD-ROM | 20 |
| Web-Accessible Documentation | 21 |
| Printed Solstice Backup and Legato NetWorker Product Manuals | 21 |
| Administrator's Guide | 21 |
| Installation Guide | 22 |
| Disaster Recovery Guide | 22 |
| Release Supplement | 22 |
| Online Help | 23 |
| Online Manual Pages | 23 |
| Error Message Guide | 24 |
| Command Reference Guide | 24 |

2. Getting Started 25

| | |
|---|----|
| Operating Solstice Backup Software for the First Time | 25 |
| ▼ To Start the Solstice Backup Administrator Program | 26 |
| ▼ To Label and Mount a Volume | 26 |
| Backing Up Your Data | 26 |
| ▼ To Configure a Scheduled Backup | 26 |
| ▼ To Start a Backup Immediately | 27 |
| ▼ To Monitor and Verify Data During a Backup | 27 |
| ▼ To Add Clients | 28 |

| | |
|---|-----------|
| Customizing and Configuring the Solstice Backup Software | 29 |
| Administrator Privileges | 29 |
| Sequence of Solstice Backup Configuration Tasks | 29 |
| Installation | 29 |
| Server Configuration | 29 |
| Client Configuration | 30 |
| Devices Configuration | 30 |
| Licensing | 30 |
| Solstice Backup Resources | 30 |
| Solstice Backup Administrator Interfaces | 31 |
| Solstice Backup Administrator Program (nwadmin) | 31 |
| Hidden Attributes | 33 |
| Character-Based Solstice Backup Interface (nsradmin) | 33 |
| UNIX Command Line | 33 |
| Solstice Backup Manual Backup, Recovery, and Archiving Interfaces | 34 |
| nwbackup | 34 |
| nwrecover | 35 |
| nwarchive | 36 |
| nwretrieve | 38 |
| Licensing | 39 |
| Overview | 39 |
| ▼ To Obtain Enabler Codes | 39 |
| ▼ To Enter an Enabler Code | 40 |
| ▼ To Remove an Enabler Code | 40 |
| ▼ To Authorize the Solstice Backup Software for Permanent Use | 41 |
| License Manager | 41 |
| 3. Server Management | 43 |

| | |
|--|----|
| Solstice Backup Server Connectivity | 43 |
| ▼ To Change Servers | 43 |
| ▼ To Terminate Your Server Connection | 44 |
| Client/Server Communication Configuration | 44 |
| Server Configuration | 44 |
| ▼ To Set Up the Server | 44 |
| Performance Factor Controls | 45 |
| Parallelism | 46 |
| Target Sessions | 47 |
| Server Access Management | 48 |
| Administrator Permissions | 48 |
| ▼ To Add Administrators | 49 |
| ▼ To Remove Administrators | 49 |
| Notifications | 50 |
| Preconfigured Notifications | 50 |
| Customizing Notifications | 51 |
| Events | 52 |
| Actions | 52 |
| Priorities | 53 |
| Printing Notifications | 53 |
| ▼ To Create Notifications | 54 |
| ▼ To Edit Notifications | 55 |
| ▼ To Delete Notifications | 55 |
| Reports | 55 |
| Routine Data Movement Operations Reports | 55 |
| Storage Management Application Reports | 56 |
| Solstice Backup Server Statistics and Diagnostic Reports | 57 |
| Message Log Files | 57 |

| | |
|--|----|
| Firewall Support | 57 |
| Service Ports | 57 |
| Connection Ports | 57 |
| ▼ To Configure Firewall Support | 58 |
| ▼ To Configure Firewall Support Using Default Settings | 58 |
| ▼ To Restrict the Port Range Solstice Backup Uses | 59 |
| Task 1: Determine Minimum Port Range Required | 59 |
| Server Port Range Requirements | 59 |
| Example: Setting the Server's Service Port Range | 60 |
| Example: Setting the Server's Connection Port Range | 61 |
| Client Port Range Requirements | 61 |
| Example: Setting the Client's Connection Port Range | 61 |
| Storage Node Port Range Requirements | 62 |
| Example: Setting the Storage Node's Service Port Range | 62 |
| Example: Setting the Storage Node's Connection Port Range | 62 |
| Task 2: Change Solstice Backup Default Settings | 63 |
| ▼ To Configure Service and Connection Ports using nswadmin | 63 |
| ▼ To Configure Service and Connection Ports using nsrports | 63 |
| Task 3: Completing Configuration | 64 |
| Example: Sample rpcinfo Output | 64 |
| Troubleshooting Port Configuration Problems | 65 |
| Storage Nodes | 66 |
| Prerequisites | 66 |
| Licensing | 67 |
| Storage Node Configuration | 67 |
| Installing Storage Node Software | 68 |
| Installing on the Server and Creating a New Storage Node | 68 |
| Installing on the Server Only | 68 |

| | |
|--|----|
| Installing on the Server and Existing Storage Node | 69 |
| Bootstrap Backup on a Storage Node | 69 |
| Storage Node Affinity | 69 |
| Storage Nodes and Cloning | 69 |
| Online Index Structures and Management | 70 |
| A Note on Upgrading from Solstice Backup 5.x Software to This Release | 71 |
| Understanding the Online Indexes | 71 |
| Automated Solstice Backup Index Actions | 72 |
| Client File Index Location Issues | 72 |
| ▼ To Designate a New Client's File Index Location | 72 |
| ▼ To Move an Existing Client File Index's Location | 73 |
| Accessing Online Index Information | 74 |
| ▼ To View Index Information | 74 |
| ▼ To View Save Set Information | 74 |
| ▼ To View Details of Each Instance of a Save Set Backup | 74 |
| Managing the Size of the Online Indexes | 75 |
| ▼ To Reduce Client File Index Size | 75 |
| ▼ To Reduce Media Database Size | 76 |
| ▼ To Remove the Oldest Save Set Cycles | 76 |
| ▼ To Remove Volume-Based Entries from the Online Indexes | 77 |
| ▼ To Compress the Media Database | 78 |
| Index Conversion Issues | 78 |
| Changes in Disk Space Requirement from Previous Solstice Backup Releases | 79 |
| Enhancing Client File Index Conversion | 79 |
| Converting a Specific Client File Index | 79 |
| Converting Multiple Client File Indexes Simultaneously | 80 |

| | |
|--|-----------|
| CPU Usage for nsrck Instances | 80 |
| Server Maintenance Tasks | 81 |
| Message Log Management | 81 |
| ▼ To Set the Startup Script to Trim Log Files | 81 |
| ▼ To Use To Solaris Services to Trim Log Files | 82 |
| ▼ To Move Your Solstice Backup Server Software to a Different Computer | 83 |
| ▼ To Define the Index Entry | 84 |
| ▼ To Complete the Move | 85 |
| ▼ To Set Up an Old Solstice Backup Server as a Client | 86 |
| ▼ To Recover Bootstrap Data from Another Server's Tapes Using mmrecov | 86 |
| Device Parameter Environmental Variables | 86 |
| NSR_DEV_BLOCK_SIZE_ | 87 |
| NSR_DEV_TAPE_FILE_SIZE_ | 87 |
| NSR_DEV_DEFAULT_CAPACITY_ | 87 |
| NSR_DEV_LOAD_TIME_ | 87 |
| 4. Client Operations | 89 |
| Solstice Backup Clients | 89 |
| Client Configuration | 89 |
| ▼ To Create Clients | 90 |
| ▼ To Edit Clients | 92 |
| ▼ To Delete Clients | 92 |
| ▼ To Access the Client ID | 93 |
| ▼ To Change a Client's Name | 93 |
| Example: Changing a Client's Name | 94 |
| ▼ To Restore a Deleted Client | 94 |
| Remote Access Rights for Other Clients | 94 |
| Remote User Attribute | 95 |

| | |
|---|-----|
| Solstice Backup Clients of Different Platforms | 95 |
| UNIX Clients | 96 |
| Windows NT Clients | 96 |
| Windows 95/98 Clients | 98 |
| NetWare Clients | 99 |
| Multiple Clients from the Same Computer | 100 |
| Using Unique Client/Save Set Combinations | 100 |
| Client/Save Set Combination | 101 |
| Save Set Status Values | 102 |
| ▼ To Change the Status of a Save Set | 103 |
| Browse and Retention Policies | 104 |
| How the Solstice Backup Server Uses Browse and Retention Policies | 104 |
| How the Browse Policy Works | 105 |
| Example: One Week Browse Policy | 105 |
| Example: Two-Week Browse Policy | 106 |
| How the Retention Policy Works | 107 |
| Example: Three-Week Retention Policy | 108 |
| Managing the Data Life Cycle | 110 |
| ▼ To Assign Multiple Policies to a Single Client | 112 |
| Example: Assigning Different Policies for an Identical Client | 112 |
| Preconfigured Policies | 112 |
| ▼ To Create Policies | 113 |
| ▼ To Edit Policies | 113 |
| ▼ To Delete Policies | 114 |
| Browse and Retention Policies for Manual Backups | 114 |
| Browse and Retention Save Set Policy Modification | 115 |
| Example: Changing Browse and Retention Policies | 116 |
| ▼ To Modify Manual Backup Policies | 116 |

| | |
|--|------------|
| Reports on Browse and Retention Policies for Save Sets | 117 |
| Granularity | 117 |
| ▼ To Override Save Set-Based Browse and Retention Policies | 118 |
| Backup Command Customization | 118 |
| Using the save Command in a Customized Backup Script | 119 |
| Using savepnp With a Customized Backup Program | 120 |
| ▼ To Back Up a ClearCase [®] VOB | 124 |
| Directives | 127 |
| ▼ To Write Directives | 127 |
| Preconfigured Directives | 130 |
| ▼ To Create Directives | 131 |
| ▼ To Edit Directives | 132 |
| ▼ To Delete Directives | 132 |
| Client Priority | 132 |
| Solstice Backup Client Licensing | 133 |
| The nsrlic Program | 134 |
| 5. Device and Media Management | 135 |
| Pools | 135 |
| How the Solstice Backup Server Uses Pools | 135 |
| Example: Using Pools Configurations to Sort Data | 136 |
| Solstice Backup Pool Types | 137 |
| Preconfigured Pools | 137 |
| Default Pool | 138 |
| Expression Matching of Save Sets to Pools | 138 |
| How the Solstice Backup Server Uses Pool Criteria to Sort Data | 139 |
| Directing Client Indexes and Bootstrap to a Separate Pool | 139 |
| Meeting the Criteria for More Than One Pool Configuration | 140 |

| | |
|---|-----|
| Not Meeting the Criteria for Any Pool | 141 |
| Configuring Pools | 142 |
| Incremental Backups | 142 |
| Manual Backups | 143 |
| Clone Data | 143 |
| Archive Data | 144 |
| Migration Data | 144 |
| Using Storage Devices and Pool Configuration to Sort Data | 144 |
| Using Different Media | 144 |
| Using a Specific Device for Backup Data | 145 |
| ▼ To Create Pools | 145 |
| Auto Media Verification | 146 |
| ▼ To Edit Pools | 147 |
| ▼ To Delete Pools | 147 |
| ▼ To Create Archive Pools | 147 |
| Volume Labeling | 148 |
| How the Solstice Backup Server Uses Label Templates | 148 |
| Preconfigured Label Templates | 149 |
| Key Label Templates Attributes | 151 |
| Label Template Naming Strategies | 153 |
| ▼ To Create Label Templates | 153 |
| ▼ To Edit Label Templates | 154 |
| ▼ To Delete Label Template | 154 |
| Devices | 155 |
| Device Support by Solstice Backup Edition | 155 |
| Supported Devices and Media Types | 156 |
| Nonrewinding Tape Device Usage | 156 |
| ▼ To Add Devices | 157 |

- ▼ To Edit Devices 158
- ▼ To Delete Devices 158
- Filesystem Devices 159
 - ▼ To Configure Filesystem Devices 159
 - Restrictions 160
 - Use with Staging 160
- Logical Volume Backup 160
- Media Handling Error Resolution 161
 - ▼ To Limit Repeated Device Errors 161
- Auto Media Management 162
- Storage Nodes 164
 - ▼ To Configure a Storage Node 164
 - Administrator Privileges 165
 - Installing the Software on a New Storage Node 165
 - Installing an Update Only on the Solstice Backup Server 165
 - Installing an Update on an Existing Storage Node 165
 - ▼ To Configure Autochangers and Silos as Storage Node Remote Devices 166
 - ▼ To Set Timeouts for Storage Node Remote Devices 166
- Storage Management Operations 167
 - How the Solstice Backup Server Uses Volume Labeling 167
 - How the Solstice Backup Server Uses Pools 167
 - How the Solstice Backup Server Selects and Mounts a Volume During a Backup 168
 - How the Solstice Backup Server Selects a Volume for Relabeling 169
- Volume Status 169
 - Read-Only Mode 171
- Solstice Backup Volume Operations 171
 - ▼ To View Volume Information 171

- ▼ To View Save Set Information 172
- ▼ To Mount Volumes 173
- ▼ To Label and Mount Volumes in One Operation 173
- ▼ To Label Without Mounting 175
 - Labeling Tips 175
- ▼ To Unmount Volumes 175
- ▼ To Find a Volume Name 176
- ▼ To Change the Mode of a Volume 176
- ▼ To Remove Bad Volumes from the Online Indexes 177

Save Set Staging 177

Staging Data from a Filesystem Device 178

- ▼ To Create Staging Policies 178
- ▼ To Edit Staging Policies 179
- ▼ To Delete Staging Policies 179

6. Backup Scheduling Tools 181

Overview of Solstice Backup Scheduling 181

Backup Groups 182

How the Solstice Backup Server Uses Groups 182

Example: Using Groups to Back Up Multiple Clients 183

Preconfigured Groups 184

Key Group Attributes 185

Abandoned Group Backups 185

- ▼ To Create Groups 185
- ▼ To Set Group Backup Time Intervals 186
- ▼ To Force Incremental Backups 187
- ▼ To Edit Groups 188
- ▼ To Delete Groups 188

| | |
|--|-----|
| Group Backup Management | 188 |
| ▼ To Start Groups Immediately | 189 |
| ▼ To Stop Groups | 189 |
| ▼ To Restart Groups | 190 |
| ▼ To View Group Control Details | 190 |
| ▼ To Preview a Group Backup | 191 |
| Save Group Completion Message | 191 |
| Bootstrap Generation and Printout | 191 |
| ▼ To Save the Bootstrap to a File | 192 |
| ▼ To Send the Bootstrap File Via E-mail | 192 |
| System Console Log | 192 |
| How the Solstice Backup Server Backs Up Open Files | 193 |
| Schedules | 193 |
| How the Solstice Backup Server Uses Schedules | 193 |
| Preconfigured Solstice Backup Schedules | 194 |
| Example: Weekly Backup Cycle | 195 |
| Example: Default Schedule with Multiple Clients | 196 |
| Example: Staggered Weekly Schedules for Multiple Clients | 197 |
| Schedule Planning Considerations | 198 |
| Scheduling Large Client Filesystems | 198 |
| Key Schedule Attributes | 199 |
| Resource Configuration Order for Schedules | 200 |
| ▼ To Create Schedules | 200 |
| ▼ To Edit Schedules | 201 |
| ▼ To Delete Schedules | 201 |
| ▼ To Override a Client's Regular Backup Schedule | 202 |
| Backup Levels | 202 |

| | |
|---|------------|
| How the Solstice Backup server Uses Backup Levels | 203 |
| How the Solstice Backup ServerBackup Levels Work | 205 |
| Example: Backup Levels Part 1 | 205 |
| Example: Backup Levels Part 2 | 206 |
| Example: Backup Levels Part 3 | 207 |
| Save Set Consolidation | 207 |
| How Save Set Consolidation Works | 207 |
| Requirements | 208 |
| Save Set Consolidation Usage | 208 |
| Limitations of Save Set Consolidation | 209 |
| ▼ To Start Save Set Consolidation | 210 |
| ▼ To Direct Data from a Consolidation Backup to a Specific Pool | 210 |
| Example: Consolidation of Level 1 and Full Backups | 211 |
| 7. Data Backup and Recovery | 213 |
| Manual Data Backups | 213 |
| ▼ To Manually Back Up Data | 213 |
| Data Recovery | 215 |
| ▼ To Recover Data | 216 |
| ▼ To Recover Data from a Previous Backup Time | 217 |
| ▼ To Relocate Recovered Data | 218 |
| ▼ To Resolve Recovered Data Conflicts | 219 |
| ▼ To Search for a Recoverable File | 220 |
| ▼ To View Details of Recoverable Files | 220 |
| ▼ To View the Version Details of Backed-Up Data | 220 |
| ▼ To View Volumes Required for Data Recovery | 221 |
| Disaster Recovery | 221 |
| Directed Recovery | 221 |

| | |
|--|-----|
| Uses of Directed Recovery | 222 |
| Access Requirements for Directed Recoveries of UNIX Clients | 222 |
| ▼ To Grant Administrative Rights for a Solstice Backup Server | 223 |
| ▼ To Grant Administrative Rights to a Remote Solstice Backup Server's Client | 223 |
| ▼ To Grant a Network Connection to a Destination Client | 225 |
| ▼ To Gain Remote Access to a Solstice Backup Client | 225 |
| Directed Recovery Between Different Client Platforms and Types | 226 |
| Directed Recovery Between Windows NT Clients | 226 |
| Directed Recovery Between Windows 95/98 and Windows NT Clients | 227 |
| Directed Recovery Between Windows NT and NetWare Clients | 227 |
| Directed Recovery Between UNIX Clients | 227 |
| ▼ To Perform a Directed Recovery Between UNIX Clients | 227 |
| recover Command in Directed Recoveries | 227 |
| Example: Directed Recovery Overwriting Existing Files | 228 |
| Directed Recovery of Specific Files and Directories | 228 |
| Save Set Recovery | 229 |
| How Save Set Recovery Works | 229 |
| ▼ To Recover Data from a Save Set | 230 |
| ▼ To Specify What Save Set Data Is Recovered | 231 |
| ▼ To View Volumes Required for Save Set Data Recovery | 232 |
| ▼ To Relocate Recovered Save Set Data | 232 |
| ▼ To Resolve Recovered Save Set Data Conflicts | 233 |
| ▼ To Restore a Save Set Entry in the Online Indexes | 233 |
| Client File Index Only | 233 |
| Client File Index and Media Database | 234 |
| ▼ To Recover an Entire Save Set to the Solstice Backup Server | 234 |

- ▼ To Recover One File Directly from a Volume 235
- Comparison of Recovery Methods 235
 - Advantages of Index-Based Recoveries 236
 - Disadvantage of Index-Based Recoveries 237
 - Costs of Save Set Recoveries 237
 - Deciding Between Index-Based and Save Set Recovery 238
- Recovering Client File Indexes 238
 - ▼ To Recover a Client's Current File Index 238
 - ▼ To Recover Client File Indexes for a Specific Time and Date 239
 - Example: Recovering a Client File Index for a Specific Time and Date 239
 - Relocating Recovered Online Indexes 240
 - Solstice Backup Release 6.0 Software 240
 - Solstice Backup Release 5.5 and Earlier Software 240
- Scanner Program 240

8. Data Archiving Methods 243

- Overview 243
 - Archives Versus Backups 244
 - Requirements 244
- Licensing 244
- Archive Application Permissions 245
 - Server Permissions 245
 - Client Permissions 245
 - Restricting Archive Access 246
 - Enabling Public Archive Access 246
- How the Solstice Backup Server Archives Data 246
 - Archiving of Non-Indexed Data 247
 - Archive Pool 248

| | |
|--|-----|
| Archiving of Indexed Data | 248 |
| PC Archive Pool | 248 |
| Volume Tracking | 249 |
| The nwarehouse Program | 249 |
| Data Archiving Options | 249 |
| Compression | 249 |
| Cloning | 250 |
| Verification | 250 |
| Grooming | 250 |
| ▼ To Create an Archive Client | 250 |
| Enabling Archive Services for an Existing Client | 251 |
| Manually Archiving Data | 251 |
| ▼ To Manually Archive Client Data | 251 |
| Scheduled Archives | 253 |
| ▼ To Set Client Archive Privileges | 254 |
| Key Archive Request Attributes | 254 |
| ▼ To Create an Archive Request | 255 |
| ▼ To Edit an Archive Request | 256 |
| ▼ To Delete an Archive Request | 256 |
| Scheduled Archive Management | 257 |
| ▼ To Start a Scheduled Archive Immediately | 257 |
| ▼ To Stop a Scheduled Archive While In Progress | 257 |
| ▼ To Disable a Scheduled Archive | 258 |
| ▼ To Reschedule a Scheduled Archive | 258 |
| ▼ To View Details of a Scheduled Archive | 259 |
| Archived Data Retrieval Operations | 259 |
| The nwretrieve Program | 260 |
| Permissions | 260 |

| | |
|--|------------|
| Retrieving Archived Data to the Client | 260 |
| Retrieving Archived Data to Another Client | 261 |
| Non-Indexed Archive | 261 |
| Indexed PC Archive | 261 |
| ▼ To Retrieve Archived Save Sets | 261 |
| 9. Cloning | 263 |
| Volume Cloning | 263 |
| How Volume Cloning Works | 263 |
| Automatic versus Manual Cloning | 264 |
| Volume Cloning Versus Volume Duplication | 265 |
| Volume Cloning Requirements | 265 |
| Volume Cloning and Data Tracking Information | 265 |
| Volume Cloning Performance | 266 |
| Volume Cloning and Recovery | 267 |
| ▼ To Clone a Volume | 267 |
| ▼ To View Clone Volume Details | 268 |
| Save Set Cloning | 268 |
| Overview | 268 |
| Retention Policies | 269 |
| Performance Levels | 269 |
| Online Index Entries | 269 |
| Volumes | 269 |
| Devices | 269 |
| Save Sets Cloning Requirements | 270 |
| Save Sets Cloning Methods | 270 |
| Manual Save Set Cloning | 270 |
| Example: Preparing to Manually Clone Save Sets | 270 |

- ▼ To Manually Clone Save Sets 271
- ▼ To View Details of Save Sets Prior to Cloning 272
- Automated Save Set Cloning 273
 - Example: Automatically Cloning One Client 273
- ▼ To Automatically Clone a Single Client 273
 - Example: Automatically Cloning Several Clients 274
- ▼ To Automatically Clone Multiple Clients 276
- Cloning Archived Data 277
- Recovering Cloned Save Sets 277
- ▼ To Change the Original Save Set to Suspect 278
- Clone Storage Node Affinity 278

10. Autochangers and Silos 281

- Autochanger Installation and Configuration 281
 - ▼ To Add Autochangers 282
 - ▼ To Test the Autochanger Connection 287
 - Key Autochanger Attributes 288
 - ▼ To Add Devices to an Autochanger 288
 - ▼ To Edit an Autochanger 290
 - ▼ To Delete an Autochanger 291
 - Remote Autochanger Management 291
 - ▼ To Adjust Sleep Periods for Autochanger Tasks 292
- Sharing Autochangers Between Solstice Backup Hosts 293
 - Autochanger Sharing Licensing and Availability 293
 - How Autochanger Sharing Works 294
 - Autochanger Sharing Configuration 295
 - Example: jconfig Command Additions 295
 - Device Naming During Configuration 295

| | |
|--|-----|
| Host Device Selection for Recoveries | 296 |
| Example: Data Recovery Using a Shared Autochanger | 296 |
| Host Device Selection for Cloning | 298 |
| Labeling, Loading, and Unloading Volumes on Shared Autochanger Devices | 298 |
| Managing Volumes on Server Devices | 298 |
| Managing Volumes on Storage Node Devices | 298 |
| Backups and Recoveries on a Shared Autochanger Device | 299 |
| Automatic Unmounting of Volumes | 299 |
| Setting Compatible Block Size for UNIX and Windows NT | 299 |
| Problem Determination | 300 |
| Solutions | 300 |
| Autochangers Operations | 301 |
| ▼ To Designate Available Autochanger Slots | 301 |
| ▼ To Clean Autochanger Devices | 302 |
| ▼ To Use a Nondefault Slot for the Cleaning Cartridge | 303 |
| ▼ To Replace Cleaning Cartridges | 304 |
| Auto Media Management with Autochangers | 305 |
| How the Solstice Backup Server Selects a Volume for Writing | 306 |
| Automatic Mounting of Autochanger Volumes | 307 |
| ▼ To Mount Autochanger Volumes | 307 |
| ▼ To Mount Uninventoried Volumes | 307 |
| ▼ To Unmount Volumes | 308 |
| ▼ To Label Autochanger Volumes | 308 |
| Bar Code Labeling of Volumes | 310 |
| ▼ To Label a Volume with Bar Codes | 310 |
| Bar Code Labeling Tips | 311 |
| ▼ To Use Unmatching Volume and Bar Code Labels | 312 |

| | |
|---|-----|
| Inventorying Autochanger Volumes | 313 |
| ▼ To Inventory Autochanger Volumes | 314 |
| I/O Errors | 314 |
| Using Autochangers with a Volume Import/Export Capability | 315 |
| ▼ To Deposit Volumes | 315 |
| ▼ To Withdraw Volumes | 316 |
| Tips and Suggestions for Using Autochangers | 316 |
| Autochanger Notifications | 317 |
| Recycling Versus Adding More Volumes | 317 |
| Using Pools with Autochangers | 318 |
| Device Calibration | 318 |
| ▼ To Disable the Element Status | 318 |
| ▼ To Reset the Autochanger | 319 |
| Optical Autochangers | 319 |
| Silos | 320 |
| How Solstice Backup Software Interacts with a Silo | 320 |
| ▼ To Install Silos | 321 |
| Silo Device Naming Conventions | 322 |
| StorageTek | 322 |
| IBM 3494 | 323 |
| EMASS/Grau | 323 |
| ▼ To Set up DAS to work with Solstice Backup Software | 323 |
| ▼ To Define a Silo Device | 324 |
| ▼ To Release a Silo Device | 324 |
| Shared Devices | 325 |
| Silo Device Cleaning | 325 |
| Media Management in a Silo | 325 |

| | |
|---|------------|
| Silo Slot Numbers | 326 |
| Using Silos with a Volume Import/Export Capability | 326 |
| Mounting and Unmounting a Silo Volume | 327 |
| Labeling a Silo Volume | 328 |
| Allocating a Silo Volume | 328 |
| ▼ To Allocate a Silo Volume | 329 |
| Deallocating a Silo Volume | 329 |
| ▼ To Deallocate a Silo Volume | 330 |
| 11. Solstice Backup Cluster Support | 331 |
| What Is Clustering? | 331 |
| Example: General Cluster Configuration | 331 |
| Cluster Terminology | 332 |
| Types of Cluster Support | 333 |
| Cluster-aware and Cluster-unaware Solstice Backup Servers | 333 |
| Sun Cluster Environment | 334 |
| Data Backup in a Cluster Environment | 334 |
| ▼ To Back Up Data From a Virtual Cluster Client | 334 |
| Data Recovery in a Cluster Environment | 336 |
| Solstice Backup Cluster Licensing | 336 |
| Enabling Support for Solstice Backup Cluster Clients | 337 |
| Licensing | 337 |
| 12. NDMP Support for the Solstice Backup Server | 339 |
| Introduction | 339 |
| NDMP Connection Support | 339 |
| NDMP Support for the Solstice Backup Server | 340 |
| NDMP Connection Software and Installation Requirements | 341 |
| NDMP Client Support | 341 |

| | |
|--|-----|
| Silo Support | 342 |
| Limitations of NDMP | 342 |
| Legato Celestra Software and NDMP | 342 |
| NetApp Computers and NDMP | 342 |
| Important Notes on Swap Space | 343 |
| Configuring the Solstice Backup Server for NDMP Operations | 343 |
| Licensing the Solstice Backup NDMP Connection | 345 |
| Configuring Autochanger Support for NDMP Operations | 345 |
| Obtaining Information About NDMP Autochangers | 345 |
| ▼ To Find the Autochanger Handle and Pathname for Celestra Devices | 346 |
| ▼ To Find the Autochanger Handle and Pathname for NetApp Computers | 347 |
| ▼ To Create the ndmpjbconf_NDMPHostname File | 348 |
| ▼ To Use the inquire Command with NDMP Devices | 350 |
| ▼ To Configure an NDMP Autochanger | 350 |
| Example: How to Configure an NDMP Autochanger for NetApp Computers | 351 |
| NDMP Effects on Solstice Backup Functionality | 354 |
| Configuring the Solstice Backup Server for NDMP Operations | 355 |
| Client Resource Configuration | 356 |
| Application Information | 356 |
| Device Resource Configuration | 358 |
| The NDMP Workflow | 358 |
| NDMP Backup | 360 |
| Prerequisites | 361 |
| ▼ To Start a Scheduled Backup | 361 |
| Media Preparation | 361 |
| NDMP Client Configuration | 361 |

- Starting the Group Backup 362
- Viewing the Group Backup Details 362
- Viewing Client File Index Information 362
- NDMP Recovery 363
 - ▼ To Perform a Destructive Save Set Recovery 363
- NDMP and NetApp Backups to a Remote Tape Drive 364
- 13. SNMP Module 367**
 - SNMP Traps 367
 - Configuring Legato NetWorker SNMP Notifications 368
 - Command-Line Options for nsrtrap 368
 - ▼ To Modify the Preconfigured Legato NetWorker SNMP Notification 369
 - ▼ To Create a Legato NetWorker SNMP Notification 370
 - Configuring SNMP Management Software 371
 - Solstice Backup SMI Network Management Private Enterprise Code 371
 - Receiving Traps in Your SNMP Network Management Software 371
 - A. Notes and Troubleshooting Tips 373**
 - ▼ To Troubleshoot IP Errors 388
 - ▼ To Set Up a Hosts Table 389
 - ▼ To Set the IP Name Search Order 390
 - ▼ To Use ping to Verify Network Connections 391
 - ▼ To Use rpcinfo to Verify That Sessions Can Be Established 392
 - ▼ To Verify Firmware for Switches and Routers 394
 - ▼ To Test the Device Driver Installation 403
 - ▼ To Set Jukebox Attributes 409
 - ▼ To Configure the ATL Library 410
 - B. UNIX Platform-Specific Notes 415**
 - ▼ To Install an HP-PB or NIO Pass-Through Driver 417

- ▼ To Install a GSC, HSC, or PCI Pass-Through Driver 420
- ▼ To Shut Down the System 420
- ▼ To Start the System 421
- ▼ To Create a Device File 421
- ▼ To Test the Device Driver and Device File Installation 424
- ▼ To Test the Autochanger Connection 425
- ▼ To Activate Full Diagnostic Core Dumps on AIX 427

14. Glossary 431

Index 441

Figures

- FIGURE 1-1 Resource Configuration Order 7
- FIGURE 1-2 Daemon Processes and Programs During a Scheduled Backup 12
- FIGURE 1-3 Daemon Processes and Programs During a Recover Session 14
- FIGURE 1-4 Daemon Processes and Programs During a Save Session with a Storage Node 15
- FIGURE 1-5 Daemon Processes and Programs During a Recover Session with a Storage Node and Remote Storage Node 16
- FIGURE 2-1 Solstice Backup Administrator Program 32
- FIGURE 2-2 `nwbackup` Program 35
- FIGURE 2-3 `nwrecover` Program 36
- FIGURE 2-4 `nwarchive` Program 37
- FIGURE 2-5 `nwretrieve` Program 38
- FIGURE 4-1 One-Week Browse Policy 106
- FIGURE 4-2 Two-Week Browse Policy 107
- FIGURE 4-3 One-Week Backup Cycle; Three-Week Retention Policy 109
- FIGURE 4-4 Data Life Cycle in the Client Index and the Media Database 111
- FIGURE 5-1 Using Pool Configurations to Sort Data 136
- FIGURE 5-2 How the Solstice Backup Server Uses Label Templates 149
- FIGURE 6-1 How Solstice Backup Groups are Used to Back Up Multiple Clients 184
- FIGURE 6-2 Weekly Backup Cycle 195
- FIGURE 6-3 Using the Solstice Backup Default Schedule for Multiple Clients 196

| | | |
|-------------|--|-----|
| FIGURE 6-4 | Staggered Weekly Schedules for Multiple Clients | 197 |
| FIGURE 6-5 | Backups for October 1 Through October 7 | 206 |
| FIGURE 6-6 | Backups for October 1 Through October 14 | 206 |
| FIGURE 6-7 | Backups for October 1 Through October 15 | 207 |
| FIGURE 6-8 | Consolidated Backup | 211 |
| FIGURE 8-1 | Overview of Archive Operation | 247 |
| FIGURE 9-1 | Automatic Cloning Scheme | 274 |
| FIGURE 10-1 | Autochanger Device Pathnames | 290 |
| FIGURE 10-2 | How Autochanger Sharing Works | 294 |
| FIGURE 10-3 | Data Recovery Using a Shared Autochanger | 297 |
| FIGURE 11-1 | Basic Cluster Configuration | 332 |
| FIGURE 12-1 | Setup and Configuration Roadmap | 344 |
| FIGURE 12-2 | NDMP Local Configuration | 359 |
| FIGURE 12-3 | NDMP Three Party Remote Configuration | 360 |
| FIGURE 12-4 | Performing a NetApp Backup to a Solstice Backup Server | 365 |
| FIGURE 12-5 | Performing NetApp Backup to a Celestra Inline Image Server | 366 |

Tables

| | | |
|-----------|--|-----|
| TABLE 1-1 | Solstice Backup Resources | 6 |
| TABLE 1-2 | Server Daemons and Programs | 8 |
| TABLE 1-3 | Functions of Client Daemons and Programs | 9 |
| TABLE 3-1 | Maximum Parallelism Levels | 46 |
| TABLE 3-2 | Preconfigured Notifications | 50 |
| TABLE 3-3 | Events | 52 |
| TABLE 3-4 | Actions | 52 |
| TABLE 3-5 | Priorities | 53 |
| TABLE 3-6 | Storage Management Report Programs | 56 |
| TABLE 3-7 | Sample of rpcinfo Output for Solstice Backup | 65 |
| TABLE 4-1 | Save Set Status Values | 102 |
| TABLE 4-2 | Preconfigured Solstice Backup Policies | 112 |
| TABLE 4-3 | Solstice Backup ASMs | 128 |
| TABLE 4-4 | Preconfigured Directives | 130 |
| TABLE 5-1 | Preconfigured Pools | 137 |
| TABLE 5-2 | Solstice Backup Hierarchy for Resolving Pool Conflicts | 141 |
| TABLE 5-3 | Preconfigured Label Templates | 149 |
| TABLE 5-4 | Key Label Templates Attributes | 151 |
| TABLE 5-5 | Examples of Number Sequences for Volume Labels | 152 |

| | | |
|------------|--|-----|
| TABLE 5-6 | Using Label Template Components | 153 |
| TABLE 5-7 | Maximum Devices | 155 |
| TABLE 5-8 | Volume Modes | 170 |
| TABLE 5-9 | Volumes Window (Volume-Related Categories) | 172 |
| TABLE 5-10 | Volumes Window (Save Set-Related Categories) | 172 |
| TABLE 6-1 | Key Group Attributes | 185 |
| TABLE 6-2 | Preconfigured Solstice Backup Schedules | 194 |
| TABLE 6-3 | Key Schedule Attributes | 199 |
| TABLE 6-4 | Backup Levels | 203 |
| TABLE 6-5 | Advantages and Disadvantages of Backup Levels | 204 |
| TABLE 7-1 | Recovery Options | 219 |
| TABLE 8-1 | Archive Request Attributes | 254 |
| TABLE 9-1 | Accounting Incremental Schedule | 275 |
| TABLE 9-2 | Accounting Full-Clone Schedule | 275 |
| TABLE 10-1 | Autochanger Attributes | 288 |
| TABLE 10-2 | Jukebox Resource Sleep Attributes | 292 |
| TABLE 10-3 | Notification Messages | 303 |
| TABLE 12-1 | NDMP Terms | 340 |
| TABLE 12-2 | Supported Operating Systems | 341 |
| TABLE 12-3 | Required Application Information Values | 357 |
| TABLE 13-1 | nsrtrap Command Line Options | 368 |
| TABLE A-1 | Boot-Time File Locations | 396 |
| TABLE A-2 | Autochangers Maintenance Commands | 402 |
| TABLE A-3 | Startup Files | 412 |
| TABLE A-4 | Environment Variables | 413 |
| TABLE B-1 | Output of <code>ioscan -f</code> | 417 |
| TABLE B-2 | Output of <code>ioscan -f -C spt</code> | 420 |
| TABLE B-3 | Output of <code>lsdev -d</code> | 423 |
| TABLE B-4 | Output of <code>ioscan</code> When Robotics is Attached to an HP-PB Card | 423 |

TABLE B-5 Output of `ioscan` With `schgr` Driver Configured on the System 424

TABLE B-6 Output of `ioscan` Without `schgr` Driver Configured on the System 424

Preface

The *Solstice Backup 6.0 Administrator's Guide* contains information on how to configure and manage the Solstice Backup™ software.

This guide also provides information about how to use and administer the Solstice Backup software when it is installed in a typical networked environment. For more detailed technical information about Solstice Backup commands, refer to the online Solstice Backup 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:

- *Solstice Backup 6.0 Installation Guide*
- *Solstice Backup 6.0 Release Notes*

Using UNIX Commands

This document may 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 Command-line variable; replace with a real name or value | 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>Solstice Backup 6.0 Installation Guide</i> | 806-5608-10 |
| Release Notes | <i>Solstice Backup 6.0 Release Notes</i> | 806-5610-10 |

Accessing Sun Documentation Online

The `docs.sun.comsm` web site enables you to access a select group of Sun technical documentation on the Web. You can browse the `docs.sun.com` archive or search for a specific book title or subject at:

`http://docs.sun.com`

Ordering Sun Documentation

Fatbrain.com, an Internet professional bookstore, stocks select product documentation from Sun Microsystems, Inc.

For a list of documents and how to order them, visit the Sun Documentation Center on Fatbrain.com at:

`http://www.fatbrain.com/documentation/sun`

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 (806-5607-10) of your document in the subject line of your email.

Introduction

This chapter introduces you to the Solstice Backup software product line. It explains how Solstice Backup products back up and recover data, and how you license Solstice Backup software for use. It also describes related Solstice Backup products.

Solstice Backup Storage Management

It is important to back up the data on computer systems to protect them from the loss of valuable data. In a networked environment, where users depend on shared data and the amount of data grows rapidly as systems are added to the network, the need to manage data becomes crucial.

Solstice Backup products provide a solution for network storage management; it protects and helps manage data across an entire network of computers. The Solstice Backup software saves valuable administrator time by speeding up and simplifying daily operations of storage management. A graphical user interface (GUI) guides you through administering and configuring your network for storage management. As the Solstice Backup software manages your data, it creates a database of saved data, making it easy to locate data for recovery. Furthermore, as the network and amount of data grow, the Solstice Backup environment provides the capacity and performance to handle the load.

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

- Solstice Backup client software, which is a collection of processes and programs installed on the computers that contain data to be managed.
- Solstice Backup server software, which is a collection of processes and programs installed on the computer that performs the Solstice Backup services, including data backup, recovery, and archiving.

- Solstice Backup storage node software, which is a collection of processes and programs installed on computers that control their attached storage devices during Solstice Backup operations, including backup, recovery, and archive.

You can configure your Solstice Backup software to back up data to storage devices that are directly connected to the Solstice Backup server, or you can create a Solstice Backup *storage node*, which is a separate computer with one or more storage devices connected. The Solstice Backup server maintains the resources and files that keep track of Solstice Backup data. The storage node computer controls the storage devices and media.

You can direct backup data, archive data, and migrate data to specific collections of volumes, according to particular data characteristics. As data is written to storage media, the Solstice Backup server creates a database to track the location of all managed data. The Solstice Backup 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 point in time specified by the user. In response to a data recovery request, the Solstice Backup storage management system locates the volume that contains the requested data and either directs a device to mount the appropriate volume for retrieval, or sends a message to the operator to request the volume by name. In this way, the Solstice Backup system manages all storage volume operations.

Cross-Platform Management

The Solstice Backup software is a cross-platform product for storage management. You can administer a Solstice Backup server from a workstation running UNIX, Windows®, or Novell NetWare®, if the workstation has the Solstice Backup client software installed and is connected by a network. Additionally, a Solstice Backup client on one platform can back up its data to a Solstice Backup server of a different platform.

The Solstice Backup server can direct and conduct administration services for any client or server on the network that has the Solstice Backup software installed and is recognized by the Solstice Backup server. The Solstice Backup administration program interface offers you only the options that are relevant to a particular Solstice Backup client, so you do not need to know the configuration of individual computers before you start a monitoring session.

Performance

Solstice Backup software has many standard and optional performance features:

- Parallelism, which allows several *save streams* to flow to the server or storage node at the same time.
- Multiplexing, which lets more than one save stream write to the same device at the same time.
- Client parallelism, which lets the client send more than one save stream at a time.
- Session management, which enables you to control the number of save streams per device to maximize the performance of each device.
- Backup to file devices and optional subsequent staging to near-line or offline volumes.
- Optional Solstice Backup software additions, such as:
 - Networker Autochanger Module
 - Networker Silo Software Module
 - Networker Archive™
 - Legato NetWorker Modules
 - Legato NetWorker SNMP (Simple Network Management Protocol)
 - Solaris Clustering
 - Legato License Manager

Ease of Use

Solstice Backup software provides tools to make protection of critical data easy to manage. With the Solstice Backup software, you can:

- Use either the graphical Solstice Backup interface or the command line to manage Solstice Backup operations.
- Administer and configure Solstice Backup functions from any computer on the network.
- Grant permission to provide the capability for recovery of one client's data to another client computer of the same operating system.
- Obtain immediate answers to questions by accessing online help and man pages.
- Take advantage of the automatic media management feature to allow the Solstice Backup server or storage node to label and mount volumes as needed for backups.

- Use the Tech Dialog and technical bulletins on the Sun web site to find answers to common questions.

Scalability

You can add to your current Solstice Backup software purchase as your 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 nonfilesystem data.
- Add support for remote storage nodes to control backup devices, while the data management tasks remain centralized on a controlling Solstice Backup server.

How the Solstice Backup Software Works

The *Solstice Backup server* is the computer on your network that runs the Solstice Backup server software, stores the client indexes, media database, and resource files; and provides backup and recovery services to the clients on the network. You can connect storage devices to the Solstice Backup server or attach them to a *storage node*. The *Solstice Backup client* must have the client software installed and must be included in the server's list of clients. All server and storage node computers that you want to back up a client to must be listed in the client's *nsrhost* file, which is created when you install the client software.

The Solstice Backup server backs up client data in increments called *save sets*. A save set typically comprises all the backed-up data from a single file system or logical disk. Save sets are saved to *volumes* (for example, tapes, magnetic disk, or optical disk) mounted in backup devices attached to the server. The server uses a *pool* to sort specific data to preselected volumes to help manage your data and collection of volumes.

The Solstice Backup server maintains records of the client save sets and the volumes in a *client file index* for each client and in a *media database* on the server. The Solstice Backup software uses these records to track the saved files and the volumes on which the files are stored. The client file index and media database contain the tracking information for all data controlled by the Solstice Backup server, whether the data is written to a device attached to the server or a device attached to a storage node.

When the Solstice Backup server backs up files, you might receive a request to mount a writable volume in the server's backup device. When a Solstice Backup user recovers files, you receive a request to mount a specific volume by its name. To fulfill either request, you only need to mount the requested volume or volumes in the device attached to the server. If you back up to an autochanger, the requested volume is automatically loaded if it is stored in the autochanger.

This guide uses the term *autochanger* to refer to a variety of robotic libraries, including jukebox, autoloader, carousel, library, near-line storage, and datawheel devices.

The Solstice Backup software supports many of the latest backup devices that provide the highest performance and most reliable solutions for your network backup and recovery requirements. Refer to the *Compatibility Guide*, available from www.legato.com, for the current list of supported devices.

Solstice Backup Resources and Attributes

Solstice Backup performance and functionality are configured through Solstice Backup *resources*. Resources are the building blocks the Solstice Backup software uses to license, configure, and operate Solstice Backup clients and servers to back up, store, and recover data. Solstice Backup resources are represented in the Administrator program using windows and icons.

Most resources can be created, edited, and deleted by administrators with the proper permissions. These resources contain the attributes for creating and editing resources. Solstice Backup *attributes* are individual components of a resource. Attributes are represented by the various text boxes, drop-down lists, and so on, contained in a resource dialog box. For example, a client would be a Solstice Backup resource, and the client's name would be an attribute of that resource. Most resources also have default settings or multiple preconfigured settings to simplify configuration and operation.

The Solstice Backup product comes packaged with preconfigured and default resource settings. However, to maximize your Solstice Backup server performance, you can modify these resources to optimize the performance of your particular server(s). Many of the features and commands found in the Solstice Backup interface, and much of the information contained in this *Administrator's Guide*, have to do with resource creation and configuration.

TABLE 1-1 lists Solstice Backup resources and their descriptions. Some of the resources described in **TABLE 1-1** also exist as attributes of other resources. These interdependent relationships affect the order in which you implement the

configuration, because component resources must be defined before they can be listed as attributes of other resources. Solstice Backup software explicitly enforces this rule.

TABLE 1-1 Solstice Backup Resources

| Resource | Action |
|----------------|---|
| Registration | Defines the Solstice Backup features enabled for this Solstice Backup server. The Registration resource is described and explained in Chapter 2 . |
| Server | Identifies the local Solstice Backup server characteristics. Only one server resource can exist per Solstice Backup server. The Server resource is described and explained in Chapter 3 . |
| Notification | Names and defines a trigger-and-action pair of events: a specific Solstice Backup event is associated with a specific server response, defined either by a default setting or one you customize yourself. The Notification resource is described in Chapter 3 . |
| Client | Identifies a unique package of data items (filesystems, disk volumes, directories, or files) for backup and assigns specific backup instructions to the package. Client resources are described and explained in Chapter 4 . |
| Policy | Names and defines a time period that can be associated with administrator-specified data lifecycle policies for data. The Policy resource is described and explained in Chapter 4 . |
| Directive | Defines file processing activities that the Solstice Backup server performs during backup of a client. Directive resources are described and explained in Chapter 4 . |
| Device | Describes a device that is either locally connected or connected to a storage node that is managed by the Solstice Backup server. Device resources are described and explained in Chapter 5 . |
| Pool | Names a collection of volumes and specifies the acceptance criteria that backup data must meet to be written to this named collection. The Pool resource is described and explained in Chapter 5 . |
| Label Template | Provides instructions for the Solstice Backup software to follow whenever the storage management application needs to create a label for a volume. Each pool is associated with a specific label template. The Label Template resource is described and explained in Chapter 5 . |
| Staging | Names and defines the Solstice Backup policy of moving data from one storage medium to another. The Staging resource is described and explained in Chapter 5 . |
| 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. Group resources are described and explained in Chapter 6 . |

TABLE 1-1 Solstice Backup Resources (*Continued*)

| Resource | Action |
|-------------|---|
| 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. The Schedule resource is described and explained in Chapter 6 . |
| Autochanger | Describes a device that can mount and write to multiple volumes automatically. The Autochanger resource and the usage of silos is described and explained in Chapter 10 . |

FIGURE 1-1 graphically describes the interrelationships between the Solstice Backup resources. Resources with outgoing arrows are both independent resources and attributes of the resources to which the arrows point.

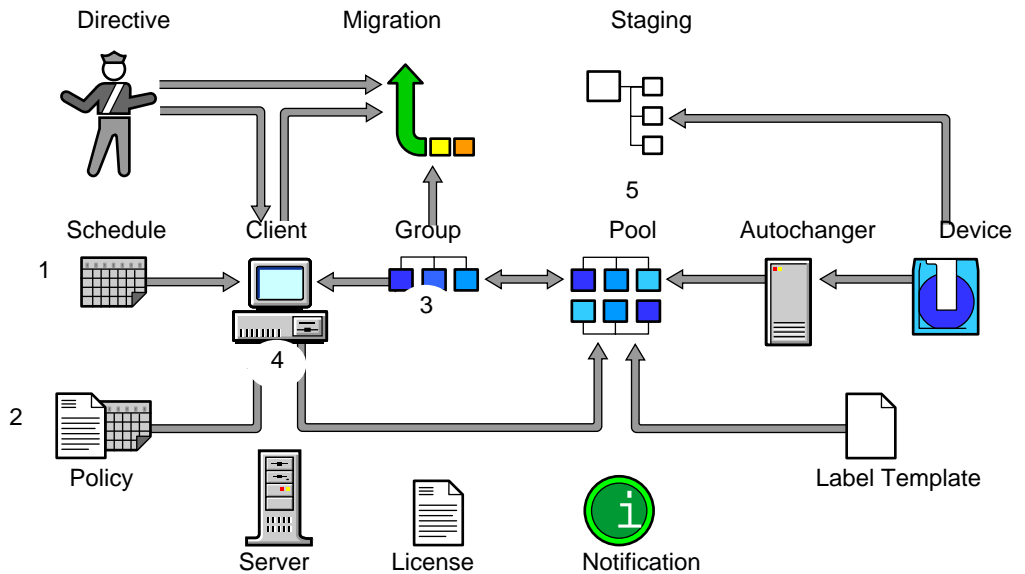


FIGURE 1-1 Resource Configuration Order

Note: In **FIGURE 1-1**, notice that Directive, Schedule, and Policy resources must be defined before they are available as attributes in the Client resource. The Solstice Backup administration feature contains preconfigured resources you can start using immediately. You can create customized resources later to optimize your Solstice Backup performance and reliability.

When you start customizing your configuration, keep in mind that resources that are attributes of other resources must be defined first. Since Schedule and Policy are two resources that do not depend on anything else, begin with these:

1. Create your customized Schedule, which becomes an attribute of the Group resource.
2. Create your customized Browse or Retention Policy, which becomes an attribute of the Client resource.
3. Customize your Group resource, which becomes an attribute of the Client resource.
4. Customize your Client resource, which becomes an attribute of the Pool resource.
5. Customize your Pool resource.

Solstice Backup Daemons and Programs

Solstice Backup client/server technology uses a network Remote Procedure Call (RPC) protocol to back up your data. The Solstice Backup server software consists of several daemons and programs that oversee the backup and recover processes, as well as storage management client configurations, a client file index, and a media database. The Solstice Backup client software includes the `nsrexecd` daemon and user interface programs.

The Solstice Backup 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 and the media containing the backed-up data.

This section provides a description of the Solstice Backup server and storage node daemons and programs, which contact the client for a backup and maintain the server's client file index and media databases. The Solstice Backup man pages contain further details about the Solstice Backup daemons and programs.

[TABLE 1-2](#) describes the server daemons and programs.

TABLE 1-2 Server Daemons and Programs

| Daemon/Program | Function |
|-----------------------|--|
| <code>ansrd</code> | This daemon monitors an active <code>save</code> or <code>recover</code> session; agent process spawned by <code>nsrd</code> in response to a <code>save</code> or <code>recover</code> session. |
| <code>asavegrp</code> | This daemon monitors the progress of individual save sets; agent process invoked by <code>savegrp</code> . |
| <code>nsrck</code> | This daemon checks the consistency of the client file index; invoked by <code>nsrd</code> whenever the consistency of the client file index needs to be confirmed. |
| <code>nsrd</code> | This daemon provides an RPC-based <code>save</code> and <code>recover</code> service to Solstice Backup clients; master Solstice Backup daemon. |

TABLE 1-2 Server Daemons and Programs (*Continued*)

| Daemon/Program | Function |
|----------------|--|
| nsrim | This daemon automatically manages the server's client file index; invoked at the end of a <code>savegrp</code> . |
| nsrindexd | This daemon provides a method for inserting entries in the client file index based on information passed by the <code>save</code> program. |
| nsrmmmd | This daemon provides device support, generates mount requests, and <i>multiplexes</i> save set data during a multiclient backup (<code>nsrd</code> can start several <code>nsrmmmd</code> daemons, up to the number specified in the device's Target Sessions attribute); media multiplexor daemon. |
| nsrmmdbd | This daemon provides media and save set database management services to the local <code>nsrd</code> and <code>nsrmmmd</code> daemons and records entries in the media database; media management database daemon. |
| savegrp | This program runs a group of Solstice Backup clients through the <code>save</code> process. |

The `nsrd` master Solstice Backup server daemon is responsible for several tasks:

- Starting other daemons
- Allocating media daemons on server and storage node machines
- Authorizing backup and recover services for the client
- Contacting clients for scheduled backups
- Maintaining Solstice Backup configuration information
- Monitoring backup and recover sessions
- Maintaining server statistics and message logs

The `nsrd` Solstice Backup server daemon calls on the `nsrexecd` Solstice Backup client daemon and several client-side programs when it receives a scheduled or on-demand backup request. The `ansrd` temporary server agent daemon starts on the Solstice Backup server to monitor the progress of the backup session.

[TABLE 1-3](#) describes the client-side daemons and programs.

TABLE 1-3 Functions of Client Daemons and Programs

| Daemon/Program | Function |
|----------------|--|
| nsrexecd | This daemon authenticates the Solstice Backup server's remote execution request and executes the <code>save</code> and <code>savefs</code> commands on the client. |
| recover | This program browses the Solstice Backup server's client file index and restores the specified file to primary disk storage. |

TABLE 1-3 Functions of Client Daemons and Programs (Continued)

| Daemon/Program | Function |
|----------------|--|
| save | This program sends specified files in a multiplexed data stream to the Solstice Backup server for backup to media by <code>nsrmmmd</code> and entry in the client file indexes and media database by <code>nsrindexd</code> . |
| savefs | This program will save a filesystem (using the <code>save</code> program) to a Solstice Backup 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 Solstice Backup Software Backs Up Data

When you configure a *backup group* on the Solstice Backup server, you schedule a start time for the backup group. The `nsrd` server daemon starts the server's `savegrp` program for the backup group at the scheduled time.

The `savegrp` program queries all client resources configured on the Solstice Backup server to determine:

- 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, determined by the parallelism value set on the Solstice Backup server
- When the most recent backup of the group occurred

If any of this information is not available on the Solstice Backup 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.

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 Clients resource. The client with the lowest value in the Client Priority attribute is contacted first.

If problems were encountered with the client file index during the previous backup session, `nsrd` invokes the `nsrck` daemon to check the consistency and state of the Solstice Backup server's client file indexes and media database. Then, `nsrd` starts the `nsrindexd` client file index insertion daemon.

There are three options: `-L`, `-R`, and `-T tempdir` for `nsrck`.

`-L` is used to specify the level of checking that will be used. Starting with the Solstice Backup 6.0 software, all recoveries of indexes are done with `nsrck -L7 client_name`. Entering no client name will recover all clients.

Refer to the *Legato Command Reference Guide* or the `nsrck` man page for more details of the seven different checking levels supported by `nsrck`.

The `nsrck` command recovers the lost index for a Solstice Backup server or client by locating the `index:client name save` set produced by the `savegrp` command at the end of a scheduled backup. The `nsrck` command queries the media database to determine which save sets to extract from which volumes to recover the index to the latest time. The `nsrck` command depends on the media and server resource files being undamaged; if this is not the case, they must be recovered using `mmrecov`.

When using `nsrck`, you would use the `-R` parameter to remove a client's index, for example, `nsrck -R -Y client`. `-Y` is used in conjunction with `-R` to indicate that you really do wish to remove the online file index). All previous command line options for `nsrck` will continue to work.

The `nsrck -T tempdir` option is used to specify a different directory to use for work space during conversion. It will enable the conversion to use the temporary directory specified as a work space for converting indexes. It is not recommended to use `/tmp`, since the `/tmp` directory contents are lost if the computer is rebooted. If this temporary space is on a separate disk, the conversion will go faster. This is also useful if the client indexes are on file systems that are nearly full.

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 Solstice Backup server or on a storage node controlled by the Solstice Backup server.

The `nsrmmmd` media daemon:

- Sends a message to the console of the Solstice Backup 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 Solstice Backup server's media database.

Any time there is a lull in save set activity from the client, the Solstice Backup 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 Solstice Backup server's configuration, to use the backup devices to their maximum potential.

The `savegrp` program repeats the process for each item on the server's work list until all clients in the group are backed up. Before the `savegrp` is completed, if the Solstice Backup server is part of the group being backed up, `nsrim` is invoked and the Solstice Backup server's `bootstrap` file is backed up. When the bootstrap backup is completed, a bootstrap printout is sent to the default printer configured for the Solstice Backup server. Keep the bootstrap printout in a safe place in case you need to restore the Solstice Backup server.

The final results of the `savegrp` execution are sent back to the server and are included in the `savegroup` completion report, which is sent through e-mail to `root`.

FIGURE 1-2 shows how all the Solstice Backup client and server daemons and programs interact during a scheduled backup.

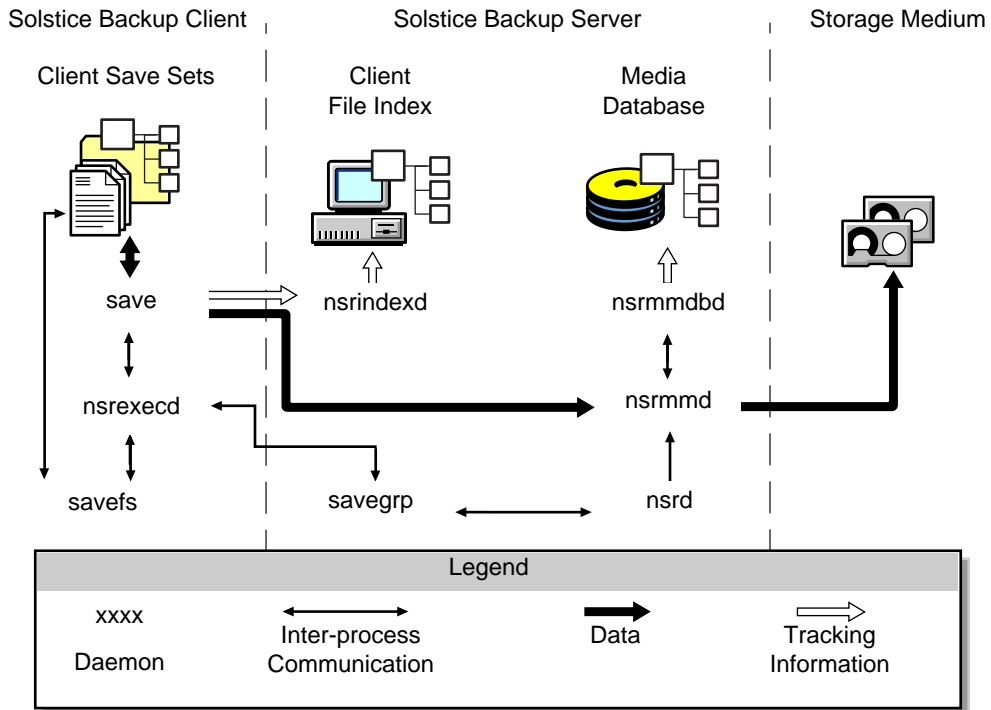


FIGURE 1-2 Daemon Processes and Programs During a Scheduled Backup

How the Solstice Backup Software Recovers Data

When the Solstice Backup server receives a `recover` request from a client, the server's `nsrd` daemon contacts the server's `nsrmmmd` media daemon. The `nsrmmmd` daemon contacts the server's `nsrmmdbd` media database daemon to determine which media contain the save set requested by `recover`. After the save set's media location is obtained, `nsrmmmd` 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 `nsrmmmd`. The media daemon forwards the save set to the client's `recover` program, which restores the data to the client's filesystem.

When the server's `nsrmmdbd` media database daemon cannot locate the required volumes, or when there are no `nsrmmmd` daemons (enabled devices) on the Solstice Backup storage node or server, the `recover` request fails. The Solstice Backup software displays the following message on the client machine that requested the `recover`:

```
NSR server client: no matching devices on server-or-storage-node
```

Solstice Backup software also displays the following message for the Solstice Backup 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
```

FIGURE 1-3 shows how the Solstice Backup server and client daemons and programs interact while recovering data to a Solstice Backup client.

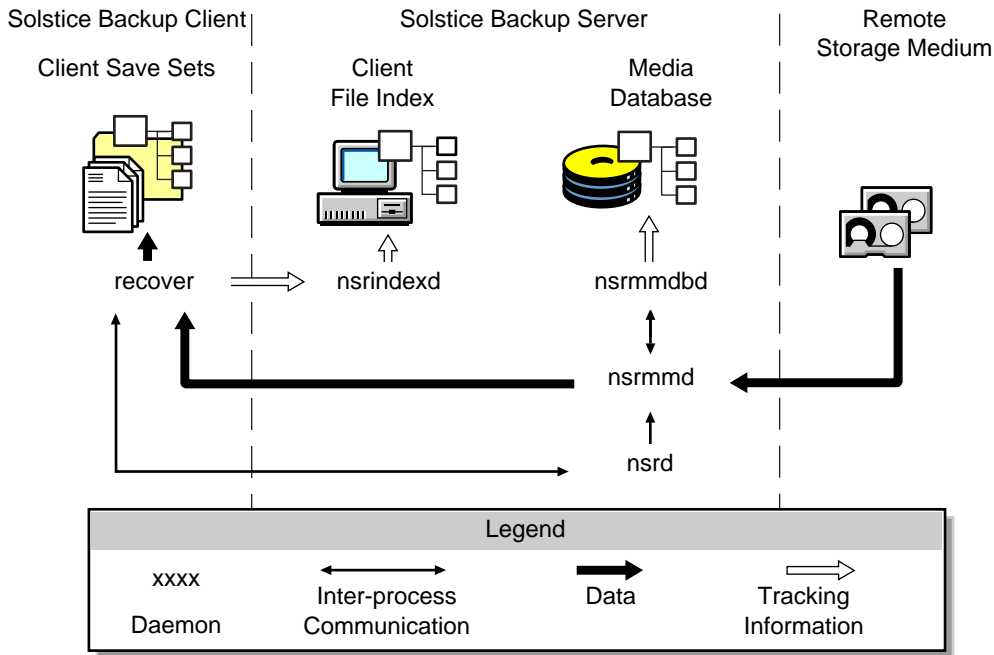


FIGURE 1-3 Daemon Processes and Programs During a Recover Session

Solstice Backup Operations with Storage Nodes

A storage node is a machine that contains Solstice Backup client and media management daemons and programs. A storage node is connected to one or more storage devices that are used in Solstice Backup operations, such as backup and recovery. Media management daemons on the storage node machine read and write data on storage volumes.

Backup and recover operations with storage nodes are very similar to backup and recover operations on the Solstice Backup server. The major difference is where the data resides.

A storage node is controlled by a Solstice Backup server (or servers). The Solstice Backup server's `nsrd` master daemon starts an instance of `nsrmmmd`, the media management daemon on the storage node and during backup. The Solstice Backup server routes appropriate data to each media management daemon according to the storage node affinity attribute in the Clients resource.

After data is written to the storage devices connected to storage node machines (remote devices), the Solstice Backup server tracks the location of the data. On the Solstice Backup server, the `nsrindexd` daemon writes entries in the client file indexes to track the location of each file in a save set, and the `nsrmmdbd` daemon writes entries in the media database to track the location of each save set on media.

FIGURE 1-4 shows the client's save process contacting the storage node's `nsrmmmd` process through the server's `nsrmmmd` process.

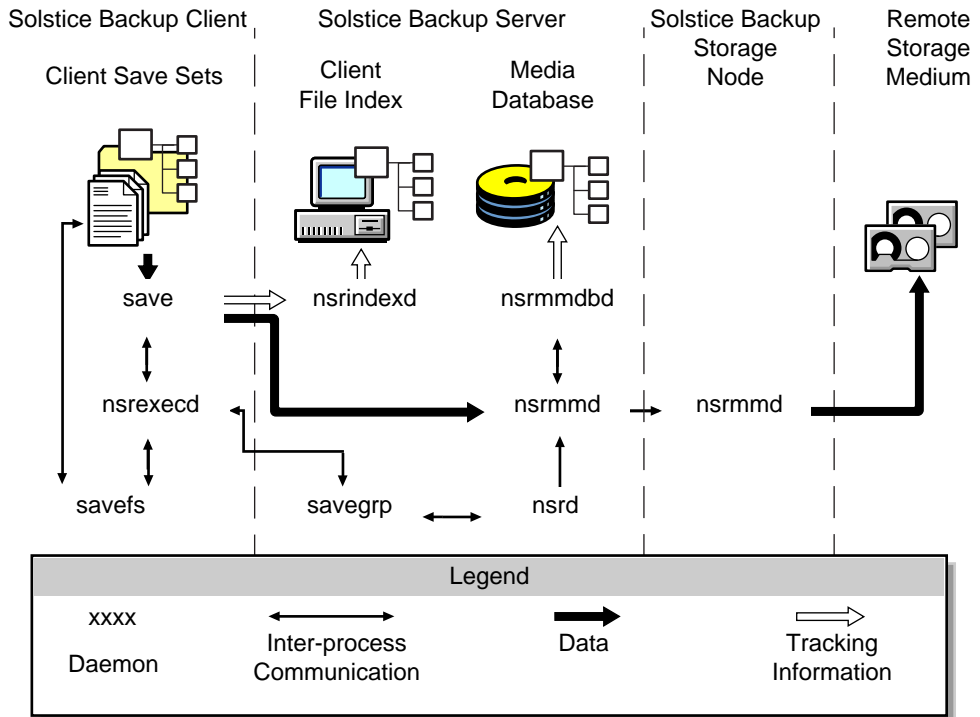


FIGURE 1-4 Daemon Processes and Programs During a Save Session with a Storage Node

When the Solstice Backup server receives a request from the client's `recover` program, it contacts `nsrmmmd`, the media management daemon, on the storage node instead of on the server machine. The `nsrmmmd` daemon reads the data from media and sends the data to the Solstice Backup client, where the `recover` program rebuilds the data.

FIGURE 1-5 shows an example of how a recover process works in a Solstice Backup setup that includes a storage node and a remote storage device.

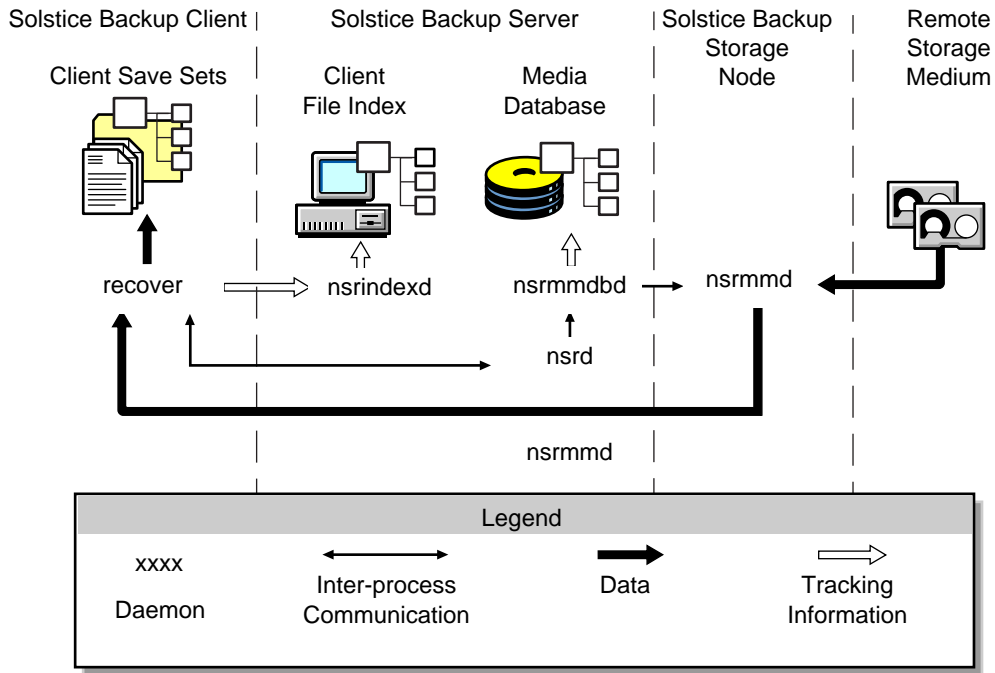


FIGURE 1-5 Daemon Processes and Programs During a Recover Session with a Storage Node and Remote Storage Node

Solstice Backup License Modes

The Solstice Backup product has different license modes to allow you the highest level of flexibility in evaluating and purchasing the software. You can use the software products in one of three software license modes:

- Evaluation mode
- Enabled mode
- Authorized mode

See [“Licensing” on page 39](#) for details on how to enable and register Solstice Backup products.

Evaluation Mode

Evaluation mode software is distributed on a free, trial basis for 30 days. You can use and evaluate the software without entering an enabler code for 30 days after you install the software on your computer. If you decide that you want to continue to use the Solstice Backup software, you must purchase the appropriate enabler code for all the features you want to keep. If you do not enter the proper enabler codes before the evaluation period ends, the Solstice Backup software does not allow further backups or configuration to take place. You can still, however, recover any data that was written to a Solstice Backup device during the 30-day evaluation.

When you first install your Solstice Backup software, all the basic Solstice Backup functionality and add-on modules are enabled for the duration of the evaluation period. However, if you already have a Solstice Backup product installed and enabled and you want to evaluate additional product features, you must enter a separate evaluation enabler for each additional feature.

Enabled Mode

Enabled mode requires a you to input an enabler code provided to you by Sun. The code that enables the Solstice Backup server is referred to as a *base enabler*. Enabler codes are generic to the product release and are entered on the Solstice Backup server. Enabled mode enables you to use the Solstice Backup software for 45 days beginning the day you enter the enabler code. After you enter the enabler code for the product, you must register the product to continue to use the software after the 45 days expire.

Your product registration information contains the host ID of the Solstice Backup server computer, which Sun incorporates into a unique authorization code to bind the license agreement to a specific computer. If you do not register the product, and then obtain and enter the authorization code during the 45-day enabled period, the Solstice Backup software does not allow further backups or configuration to take place. You can still recover previously backed-up data after the software expires.

Authorized Mode

Authorized mode begins when you enter the authorization code provided by Sun when you purchase a license for running the Solstice Backup software and each add-on module. After you enter the authorization code for your Sun products, the software is available for permanent use for the given computer.

Solstice Backup Power Edition Evaluation Enablers

Do not install a Power Edition evaluation enabler on your system if you already have a base enabler installed. If you do, the Solstice Backup software will cease to function properly. You can evaluate Power Edition only on a system that does not already have Network Edition or WorkGroup Edition enabled.

If you purchase a Network Edition or WorkGroup Edition enabler for a server on which you have evaluated Power Edition, downgrade the evaluation enabler and enter your purchased base enabler with the following command:

```
# nsrccap -d new_base_enabler
```

Authorization of Update Enablers

If you have an enabler for an earlier version of the Solstice Backup software and are installing this release as an update, you received the update enabler code in your update letter. Enter the update enabler code as you would any other enabler code. You do not need to remove the original enabler code from your Solstice Backup server in order to enter the update enabler code.

Rehosting a Solstice Backup Server

If you need to transfer the server software to another computer (or transfer the computer to another IP address), you must get a *Host Transfer Affidavit* from Sun Customer Service, complete it, and submit the form to Sun. You will be sent you a new authorization code(s), which you must install within 15 days, or the Solstice Backup software will become disabled.

Caution – Before starting the Solstice Backup software on the new server, you must delete 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/authorization code, you will receive a copy violation error message and the software will become disabled.

Solstice Backup Product Family

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

Server Software Products

All Solstice Backup clients (and servers, which can function as their own clients and back up their own data) can be backed up by Solstice Backup server software.

Solstice Backup server products are available in three distinct versions:

- Solstice Backup WorkGroup Edition, which enables you to back up the server and up to three clients. Solstice Backup WorkGroup Edition does not include support for additional clients or optional software features such as Legato NetWorker Archive. If you decide later to purchase a higher level of Solstice Backup software, all the data backed up by Solstice Backup WorkGroup Edition is recoverable by any level of Solstice Backup software you purchase.
- Solstice Backup Network Edition, which enables you to back up the server and as many client workstations as you purchased client connections for. You can upgrade Solstice Backup Network Edition to include support for additional clients, as well as optional software features.
- Solstice Backup Power Edition, a server software product that has been specially tuned to protect and provide backup for environments that support VLDB (very large database) or large filesystem applications (in the terabyte range). You can upgrade Solstice Backup Power Edition to include support for high-speed devices as well as cluster servers and clients.

Client Software Products

Some Solstice Backup client software products are included with the server products. Refer to the *Solstice Backup Installation Guide* for the client support packages available, and for instructions explaining how to install the software. Support for platform-specific clients is also sold separately as Legato NetWorker ClientPak[®] software. Presently, Sun produces client software support for the following platforms and operating systems:

- UNIX: Solaris, SunOS, COMPAQ Tru64 UNIX, Dynix/PTX, HP-UX, AIX, IRIX, and UnixWare
- PC Desktops: Microsoft Windows NT Workstation, Microsoft Windows 2000 Workstation, Microsoft Windows 98, and Microsoft Windows 95

- Microsoft Windows NT Server and Microsoft Windows 2000 Server
- Novell NetWare/IntraNetWare

Other Products

Sun offers other products to enhance your storage management needs:

- The Networker Autochanger module, which enables you to fully use the automatic backup capabilities of the Solstice Backup software when connected to one or more autochangers.
- The Networker Archive application, which provides you with a sophisticated method to archive valuable data for long-term storage.
- The Legato NetWorker SNMP (Simple Network Management Protocol) module, which communicates notifications to SNMP-compliant network management stations.
- The Legato NetWorker Modules, which reliably back up the most widely used databases, including Microsoft Exchange Server and SQL Server™, Oracle®, INFORMIX®, Lotus Notes® and Sybase®.

Sun continues to develop enhancements to the Solstice Backup product line, including additional server, client, database, and add-on products. Access the Sun web site (www.sun.com) for the latest product information.

Solstice Backup Documentation

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

Documentation Suite CD-ROM

All Solstice Backup product manuals and guides are provided in PDF format on a separate Documentation Suite CD-ROM as part of the product Media Kit.

Web-Accessible Documentation

Solstice Backup product documents often are updated between releases, particularly *Release Supplements*. Periodically check the Legato web site to see if your Solstice Backup product documentation has been updated from your existing version.

Printed Solstice Backup and Legato NetWorker Product Manuals

You can also obtain a Solstice Backup product Documentation Kit for certain products for a nominal charge. A Solstice Backup product's documentation kit typically includes the *Administrator's Guide*, *Disaster Recovery Guide*, and other select titles. Contact your Solstice Backup sales representative for information on ordering Documentation Kits.

Administrator's Guide

The *Administrator's Guide* is your main resource for understanding how the Solstice Backup product works. The Administrator's Guide provides information on a wide variety of Solstice Backup topics, including the following:

- How the Solstice Backup software saves and recovers data
- Configuring Solstice Backup servers and clients
- Understanding and using Solstice Backup resources
- Configuring backup devices and media, including autochangers and silos
- Using specialized storage management tools and protocols, including:
 - Cloning
 - Archiving
 - Staging
 - Network Data Management Protocol (NDMP)
 - Simple Network Management Protocol (SNMP)
- Troubleshooting

Installation Guide

The *Installation Guide* is your single source for installation information, such as:

- How to install the software components on your server, client, and designated storage node machines
- How to configure your backup devices
- How to update or upgrade your current Solstice Backup software
- How to remove the software components

Review the *Installation Guide* thoroughly before you install the Solstice Backup software. After the installation is 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.

The *Installation Guides* for the Legato NetWorker ClientPak for NetWare, UNIX, and Windows client software provide instructions for installing the client software on other operating systems that are not supported as Solstice Backup servers, but *are* supported as Solstice Backup clients.

Disaster Recovery Guide

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

Review the *Disaster Recovery Guide* when you first install and configure the Solstice Backup software to be prepared with strategies to protect your data and operations. Keep a printed copy of the guide with your other recovery information so that it is available in case of a disaster.

Release Supplement

The *Solstice Backup Release Supplement* provides late-breaking information on software bugs, workarounds, and other items not found in the other documentation sources. The *Release Supplement* might also contain information on new and improved features that could not be included in this *Administrator's Guide* or other documentation.

Before you install the Solstice Backup software, review the latest *Release Supplement* posted on the Sun web site (www.sun.com) for any last-minute notes on installation procedures. After you successfully install the software, review the *Release Supplement* more thoroughly for notes on how to configure the Solstice Backup software, backup devices, and backup media, as well as for information on new product features.

Updated versions of this *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.

Online Help

The Solstice Backup Administrator program and other graphical user interfaces include an online help component. The online help contains specific information about the GUI, including context-sensitive help in each window that describes the purpose of the window and what information to enter in each field in the window. There are also help topics that describe Solstice Backup features as well as topics that provide step-by-step procedures for doing Solstice Backup tasks with the GUI.

To access the online help, select the Help menu displayed in the graphical user interface. You have four choices:

- On Window provides help on the current window.
- On Topic provides a list of help topics from which to choose.
- On Help provides information on how to use the online help.
- On Version provides information on the version of Solstice Backup software that is installed on your system.

The Solstice Backup Administrator also provides Help buttons in most of the dialog boxes. Click the Help button in a Solstice Backup dialog box to display the help topic associated with that Solstice Backup feature.

Online Manual Pages

The manual (man) pages are a reference for the syntax and function of the Solstice Backup commands that you issue from the command line. To view a man page, make sure that the *MANPATH* environment variable includes the path where you

installed the Solstice Backup man pages, and then enter the man command plus the name of the command at the shell prompt. For example, to view the man page for the `nsrjb` command, type:

```
man nsrjb
```

Error Message Guide

The *Legato Error Message Guide* describes messages and troubleshooting procedures for problems that you might encounter while using the Solstice Backup software. The messages documented in this guide are those which are the most troublesome and appear most frequently when using *Solstice Backup* products.

Command Reference Guide

A complete set of the Sun online manual pages is available to reference the syntax and function of the Solstice Backup commands you issue from the command line.

The *Legato Command Reference Guide* describes the commands and resources that you might encounter while running *Sun* products. This document is your best source of information on commands you enter at the command prompt, including command options. For instructions on how to use the *Command Reference Guide*, including tips and troubleshooting, view the README file.

Getting Started

This chapter leads you through some initial Solstice Backup tasks as you get acquainted with the product and start your first backups. These tasks are a good place to start if you are evaluating the Solstice Backup software or are new to the product. Later in this chapter is an introduction to the different graphical and text-based user interfaces. Finally, this chapter provides detailed instructions on how to license your Solstice Backup product.

Operating Solstice Backup Software for the First Time

Use the following information to evaluate the Solstice Backup software and familiarize yourself with the product. Whenever possible, the tasks in this section use default and preconfigured settings. Later in this *Administrator's Guide*, you will learn how to begin customizing Solstice Backup resources and attributes, and how to configure the Solstice Backup software to work for your particular needs.

This section provides a series of procedures to run the basic functions of the Solstice Backup software, including the following tasks:

- Starting the Solstice Backup Administrator program
- Labeling and mounting a volume
- Backing up your data

▼ To Start the Solstice Backup Administrator Program

Enter the following command at the prompt to start the Solstice Backup Administrator program:

```
# nwadmin &
```

▼ To Label and Mount a Volume

To label and mount a volume:

1. Insert an unlabeled or recyclable volume in the Solstice Backup server's storage device, or verify that these types of volumes are already present for the Solstice Backup server to access.
2. Select the appropriate device in the Devices list of the Solstice Backup Administrator program.
3. Select Media>Label or click the Label toolbar button to open the Label window.
4. Check the Mount After Labeling attribute and click OK.

The Solstice Backup server automatically labels the volume and then mounts the volume into the device.

Backing Up Your Data

When the backup volume is successfully labeled, try performing a backup. Each Solstice Backup server is also a client of itself, meaning that the Solstice Backup server will back itself up at the scheduled time of the group to which this server-based client is a member. For this first attempt at backing up data, the clients being used are limited to just the Solstice Backup server's client.

▼ To Configure a Scheduled Backup

To configure a scheduled backup:

1. Select Customize>Groups to open the Groups window.
2. Select Enabled in the Autostart attribute, and then click Apply.

Your Solstice Backup server is now configured to back itself up, starting at 3:33 a.m.

Now you can test the scheduled backup by starting it immediately.

▼ To Start a Backup Immediately

To start a backup immediately:

1. Select Customize>Groups to open the Groups window.
2. Select Start Now in the Autostart attribute, and then click Apply.

Your Solstice Backup server will back itself up immediately.

Note: The scheduled backup will still occur at 3:33 a.m. By starting it now however, you can see how the Solstice Backup server performs an automatic backup.

Next, you can follow the detailed progress of the backup you just started.

Caution – If you manually stop a scheduled backup, the Manual restart option is set automatically, overriding the Autorestart setting. To ensure that the scheduled backup restarts automatically, you should disable the Manual restart option in the Solstice Backup Administration GUI.

The Manual restart option is enabled each time you stop a scheduled backup manually. Consequently, you must disable the Manual restart option each time you have stopped a scheduled back up manually.

▼ To Monitor and Verify Data During a Backup

You can monitor and verify your data while the backup you just started is in progress.

To monitor and verify data during a backup:

1. Select Server>Group Control or click the Group Control toolbar button to open the Group Control window.
2. Select the Default entry in the Name attribute.

The status of this entry's status is displayed in the Group Control window, indicating that it is in the process of being backed up.

3. Click Details to open the Group Control Details window.

The Solstice Backup Administrator displays more detailed messages about the progress of the backup. When a backup is complete, the Group Control Details window displays the successful, unsuccessful, and incomplete save sets.

Note: If you have a stand-alone storage device, make sure that you leave the volume in the device. Otherwise, the Solstice Backup server will not be able to begin the scheduled backup for the Default group at 3:33 a.m.

▼ To Add Clients

Now that you have successfully backed up the Solstice Backup server's client, you can add clients to the Solstice Backup server for the various computers in your enterprise. To add a client computer to the Solstice Backup server's backup group, each of the computers must have the appropriate Solstice Backup client software installed before the Solstice Backup server can back up their data.

You then must create a Client resource for that computer. Finally, you must add this Client resource to a Solstice Backup server's backup group. When the backup group starts, your client's data is backed up.

To add a client:

1. Ensure that each client computer has the appropriate Solstice Backup client software installed.
2. Select Clients>Client Setup to open the Clients window.
3. Click Create.
4. Enter the new client's hostname in the Name attribute.
5. Accept the default settings for the other attributes for now.
6. Click Apply.

Once successfully added, the client's name and icon appear in the Clients window.

7. Repeat these steps for each Solstice Backup client you want to be backed up by your Solstice Backup server.

When you have added some more clients to your Solstice Backup server's default backup group, you can go ahead and start the backup immediately, or wait for the scheduled backup to occur at 3:33 a.m.

Customizing and Configuring the Solstice Backup Software

Now that you have operated some of the basic features and functionality of the Solstice Backup software using the default and preconfigured settings, you can now start to customize and configure your Solstice Backup settings for your environment. The Solstice Backup software allows you to configure some or all of its settings, and you can mix customized and preconfigured settings in any combination.

Administrator Privileges

Only users who have administrative privileges can create or apply configurations. Because only `root@server-name` has Solstice Backup administrative privileges by default, you must become root on the Solstice Backup server before you start any Solstice Backup administration programs or tasks. You can add other users to the list of authorized Solstice Backup administrators later.

Sequence of Solstice Backup Configuration Tasks

This section provides a suggested roadmap for you to follow to set up and configure the Solstice Backup software for your environment. Cross-references to instructions found elsewhere in this *Administrator's Guide* are included.

Installation

Install and enable the Solstice Backup software according to the instructions in the *Installation Guide*.

Server Configuration

Configure the Solstice Backup server. See [“Server Configuration” on page 44](#) for more information.

Client Configuration

Configure the Client resource for the server and any additional clients that connect to the server. See [Chapter 4](#).

Devices Configuration

Configure one or more devices:

- For stand-alone devices, see [Chapter 5](#).
- For devices residing in autochangers and silos, see [Chapter 10](#).
- For filesystem devices, see [“Filesystem Devices” on page 159](#).

Licensing

Register your Solstice Backup software and enter the authorization code returned to you within 45 days of enabling it. You can register the Solstice Backup software either online, by fax, or by mail. See [“Licensing” on page 39](#) for more information.

Solstice Backup Resources

After you become familiar with basic Solstice Backup functionality using default and preconfigured settings, you can begin customizing the following resources to meet your specific storage management needs:

- Backup groups and schedules

You can assign your Solstice Backup clients to have their data backed up according to specific times and days by setting the appropriate values for the Groups and Schedules resources. See [Chapter 6](#) for more information.

- Pools and label templates

You can determine the criteria by which backed-up data is written to a particular set of volume, and also configure a template for naming those volumes using the Pools and Label Template resources. See [Chapter 5](#) for more information.

- Browse and retention policies

Set policies that determine the amount of time a particular backed-up set of data is still browsable and recoverable using the Policies resource. The Policies resource also helps you keep the size of your online indexes at manageable levels. See [“Browse and Retention Policies” on page 104](#) for more information.

- Directives

You can designate specific conditions and methods for backing up data using the Directives resource. See [“Directives” on page 127](#) for more information.

- Notifications

You can instruct the Solstice Backup software to inform you about the Solstice Backup server’s activities, such as software registration reminders, index size alerts, tape mount requests, and completion of scheduled backups by customizing the Notifications resource. See [“Notifications” on page 50](#) for more information.

Solstice Backup Administrator Interfaces

You can administer your Solstice Backup server software by the following methods:

- The graphical Solstice Backup Administrator program (`nwadmin`)
- The character-based Solstice Backup administration interface (`nsradmin`)
- The UNIX command line

Solstice Backup Administrator Program (`nwadmin`)

The Solstice Backup Administrator program ([FIGURE 2-1](#)) is a graphical user interface for X Windows environments. Here you can register, configure, and monitor the Solstice Backup servers, clients, storage nodes, and devices in your network.

To start the Solstice Backup Administrator program, enter the following command at the prompt:

```
# nwadmin
```

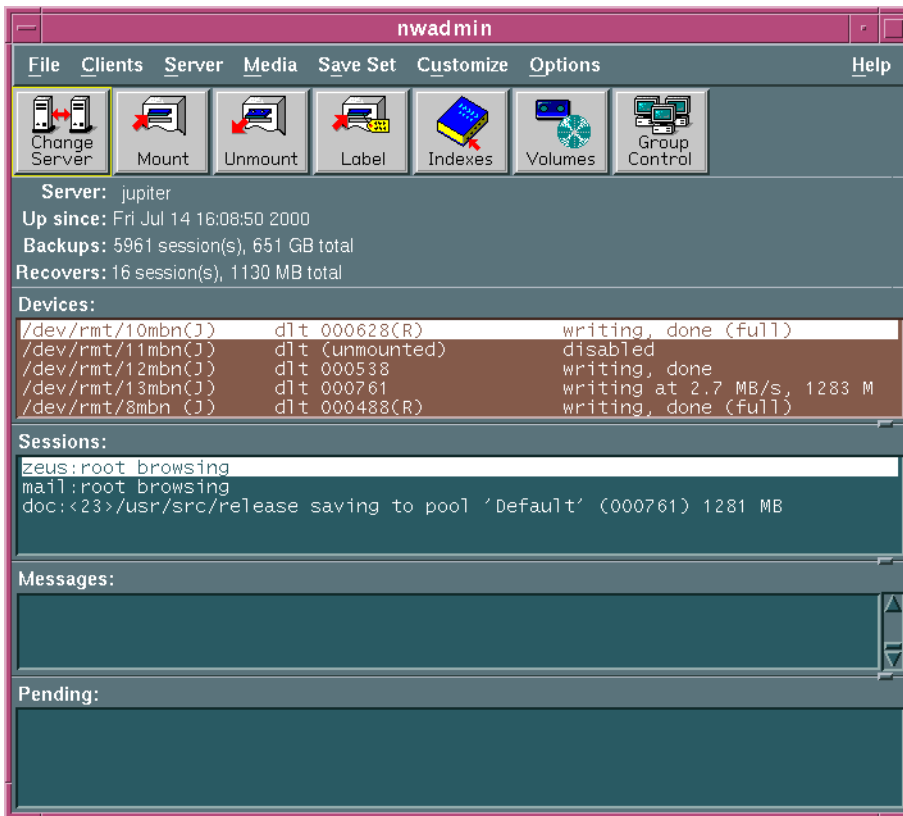


FIGURE 2-1 Solstice Backup Administrator Program

- For progress messages, watch the Sessions panel. This information is also written to the `daemon.log` file in the `/nsr/logs` directory.
- For completion messages, watch the Messages panel. This information is also written to the `messages` file in the `/nsr/logs` directory.
- For information on why a backup cannot progress, view the messages in the Pending panel. This panel shows requests to mount tapes and other conditions that require intervention.

The `nsrwatch` program also provides this information in a character-based format. For more information on the `nsrwatch` program, refer to the *Legato Command Reference Guide*.

Hidden Attributes

The Solstice Backup Administrator program and other graphical interfaces, such as the `nwbackup` and `nwrecover` programs, by default leave some attributes hidden from view in their respective resource windows. In most cases, these hidden attributes are read-only displays of more detailed Solstice Backup information, or they pertain to more complex aspects of the product.

You can configure most Solstice Backup displays to show these hidden attributes by selecting `View>Details` in the particular window. Changing values in these hidden attributes is not recommended unless you are instructed to in the documentation, or by a Solstice Backup technical support specialist.

Character-Based Solstice Backup Interface (nsradmin)

You can use the Solstice Backup character-based interface on any display, including those that do not support graphics. You can perform many of the same configuration and management tasks as in the Solstice Backup Administrator program.

You can start the character-based interface by entering the following command at the prompt:

```
# nsradmin -c
```

For more information about using the character-based interface, refer to the `nsradmin man` page.

UNIX Command Line

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

Solstice Backup Manual Backup, Recovery, and Archiving Interfaces

The UNIX version of the Solstice Backup 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`)

These task interfaces are all similar in appearance, but are invoked as separate executable programs.

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

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

`nwbackup`

You use the `nwbackup` program is used to perform a manual backup initiated from the client, rather than from the server's scheduled group backup (`savegrp`).

You start the `nwbackup` program by entering the following command at the prompt:

```
# nwbackup -s server
```

Use the `-s server` option to specify a particular Solstice Backup server on the network.

[FIGURE 2-2](#) displays the `nwbackup` program.



FIGURE 2-2 nwbackup Program

nwrecover

You use the `nwrecover` program to administer and monitor Solstice Backup servers.

You start the `nwrecover` program by entering the following command at the prompt:

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

Use the `-s server` option to specify a particular Solstice Backup server on the network. Use the `-c client` option to specify a particular Solstice Backup client on the network.

FIGURE 2-3 displays the `nwrecover` program.



FIGURE 2-3 `nwrecover` Program

nwarchive

You use the `nwarchive` program to archive files on a manual basis to a Solstice Backup server. Use of the `nwarchive` program requires the purchase of a license for the Legato Networker Archive Application. You must also enter that licensing information into the Solstice Backup server to fully enable the archiving functionality. See [“Licensing” on page 39](#) for general information on licensing the Solstice Backup software.

You start the `nwarchive` program by entering the following command at the prompt:

```
# nwarchive -s server
```

Use the `-s server` option to specify a particular Solstice Backup server on the network.

See [Chapter 8](#) for more information on the Solstice Backup archiving.

[FIGURE 2-4](#) displays the `nwarchive` program.

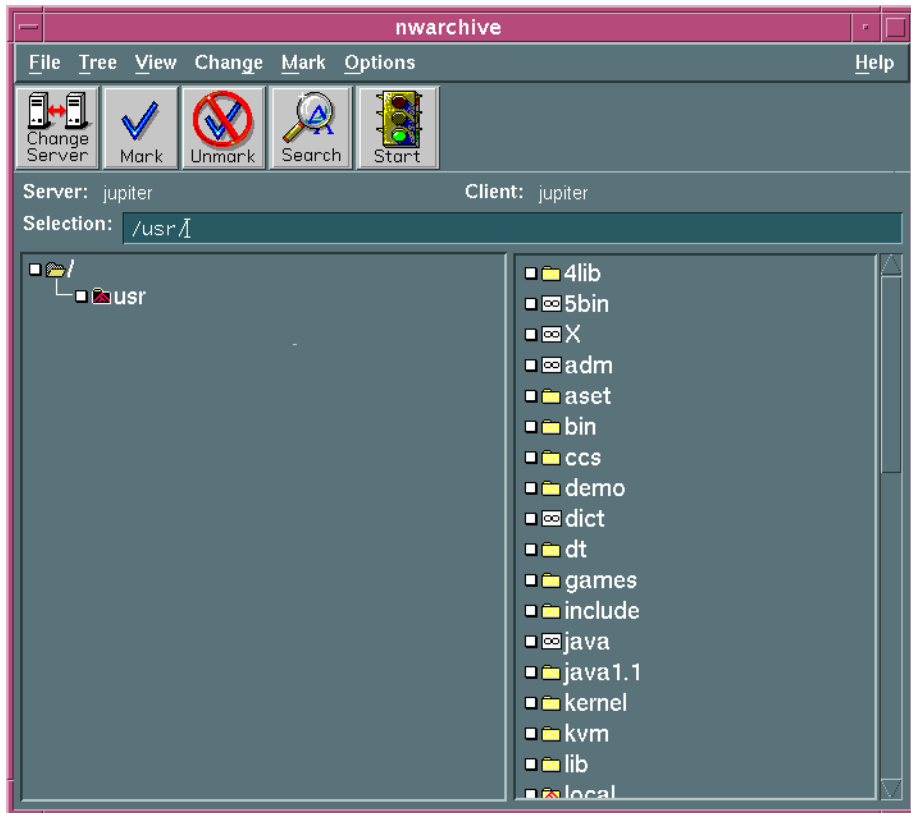


FIGURE 2-4 nwarchive Program

nwretrieve

You use the `nwretrieve` program to retrieve archived files on a manually from a Solstice Backup server.

You start the `nwretrieve` program by entering the following command at the prompt:

```
# nwretrieve -s server
```

Use the `-s server` option to specify a particular Solstice Backup server on the network.

FIGURE 2-5 displays the `nwretrieve` program.

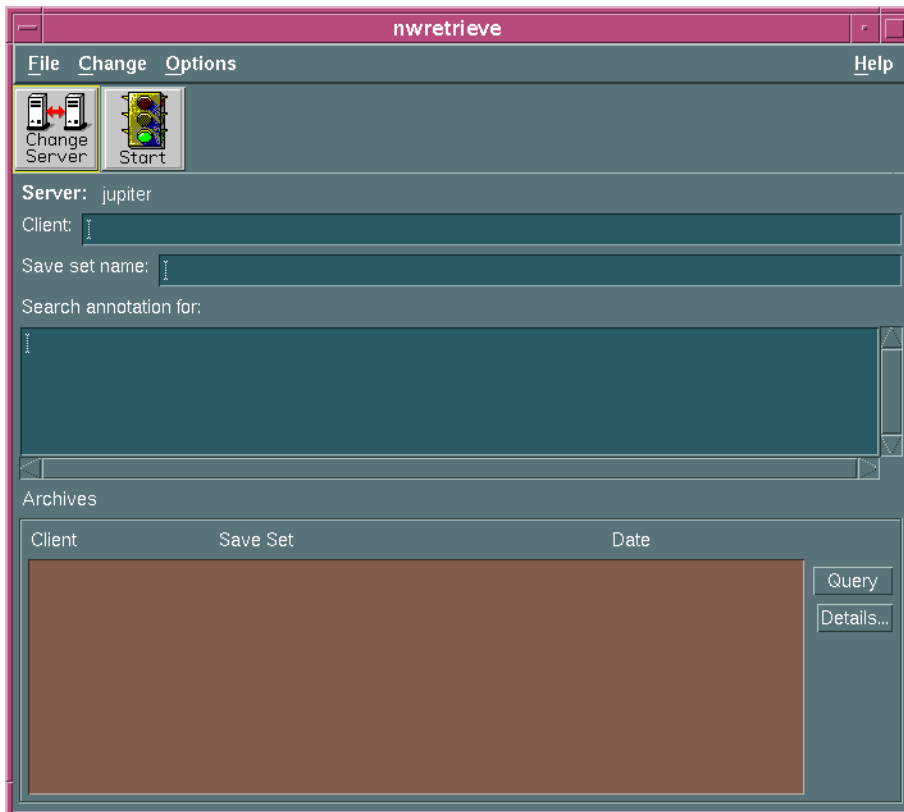


FIGURE 2-5 `nwretrieve` Program

Licensing

The following section describes the procedures for licensing your Solstice Backup software.

Overview

After you install the Solstice Backup server software, you can use it free of charge for up to 30 days. Following this free evaluation period, you can extend the functionality of your Solstice Backup server software for up to another 45 days by entering an enabler code. You can obtain Solstice Backup enabler codes either online, or by faxing or mailing information regarding your company and your Solstice Backup server to Sun.

You can only obtain an *authorization code* by purchasing the Solstice Backup software. If you fail to enter your authorization code by the end of your 45-day enabler code period, you will lose Solstice Backup functionality, including the ability to save and recover data. However, you can regain that functionality at any time by purchasing and entering an authorization code.

Once you have purchased your Solstice Backup product, the authorization code will be sent to you. You enter the authorization code in the same manner that you entered the enabler code.

▼ To Obtain Enabler Codes

To obtain enabler codes for your Solstice Backup products, contact the Sun License Center using one of the following methods:

- World Wide Web: <http://www.sun.com/licensing>
- E-mail: license@Sun.com
- Voice: 1-801-431-1595
- Fax: 1-801-431-3657

If you lose your license, you can request another via e-mail at:

lost-license@sun.com.

Also, because Solstice Backup licensing is based on the host id, an affidavit of migration must be obtained in order to move the Solstice Backup server to another machine. An affidavit of migration can be requested from the Sun License Center at the locations listed above. Once the affidavit is completed, the license center will issue a new Authorization code.

Once you have obtained your codes, you must enter them into your Solstice Backup software. See [“To Enter an Enabler Code” on page 40](#).

▼ To Enter an Enabler Code

Once you have obtained your enabler code, you must then enter it into the Solstice Backup software. You can complete this task either using the Solstice Backup Administrator program, or through the command line.

You can purchase additional enablers for added features, such as Legato Networker Archive as well as support for additional clients or autochangers that you add to your network after you install the Solstice Backup software. If you are entering multiple enabler codes, enter the base enabler code for the Solstice Backup server before you enter the enabler codes for additional features.

1. Start the Solstice Backup Administrator program (nwadmin).
2. Select Server>Registration to open the Registration window.
3. Click Create.
4. Enter the enabler code in the Enabler Code attribute and click Apply.

Your Solstice Backup server is now enabled for 45 days. You will still need to enter your authorization code to permanently enable Solstice Backup functionality.

▼ To Remove an Enabler Code

You can remove an enabler code at any time, whether or not it has been permanently authorized.

Caution – You cannot remove the code that enables the basic Solstice Backup server software, called the *base enabler code*. A base enabler can only be upgraded or downgraded.

1. Start the Solstice Backup Administrator program (nwadmin).
2. Select Server>Registration to open the Registration window.

3. Select the Solstice Backup product whose enabler code you want to delete from the Registration window. That product's enabler code is displayed in the Enabler Code attribute.
4. Click Delete to remove the enabler code.

Once you delete a Solstice Backup product's enabler code, it will no longer function until the enabler code is reentered.

▼ To Authorize the Solstice Backup Software for Permanent Use

Obtaining and entering the authorization code is the final step towards fully licensing your Solstice Backup software. Once you purchase your Solstice Backup software, you will be sent the authorization code. When entering the authorization code, you follow the same steps as when entering the enabler code. After you enter the authorization code, your Solstice Backup software is authorized for permanent use.

License Manager

Solstice Backup licensing can be managed locally at each Solstice Backup server, or it can be managed centrally using the License Manager, which provides a central control point (on a single computer) for managing the licenses of all your Solstice Backup clients and servers.

Installation of the License Manager is an option when installing the Solstice Backup software. The License Manager is required on only one Solstice Backup server. Refer to the *License Manager Administrator's Guide* for details on using the License Manager, which can be found on your Documentation Suite CD-ROM.

Server Management

This chapter describes operations that you manage through the Solstice Backup server and storage node configuration.

Solstice Backup Server Connectivity

The Solstice Backup Administrator program will, by default, open already connected to your own Solstice Backup server. However, you can change to another server at any time in order to configure that Solstice Backup server or monitor its activities. Remember that unless you have administrator privileges for that server, you cannot issue any commands; you can only monitor its activities.

▼ To Change Servers

To change the Solstice Backup server you are administrating:

1. Select Server>Change to open the Change Server window. You can also simply click the Change Server toolbar button.
2. Select a server from the Server list. If the server you want is not listed, try one of the following:
 - Click Update Server List to refresh the list.
 - Enter the name of the server in the Server attribute.
3. Click OK.

The Solstice Backup Administrator program now displays information for the new server to which you changed.

▼ To Terminate Your Server Connection

To quit the Solstice Backup Administrator program and terminate the server connection, select File>Exit.

Client/Server Communication Configuration

Communication between the Solstice Backup server and its clients is described by configuration values you enter in the Server resource and Client resource. The Solstice Backup server relies on full and accurate configuration of the network to implement features that protect data and ensure security.

Each Client resource should include both the DNS (domain name service) short name and long name in the Aliases attribute of the Client resource in the Solstice Backup Administration program.

For more details about how to set up Solstice Backup clients in the Client resource, see [“Solstice Backup Clients” on page 89](#).

To diagnose problems with network communications that affect Solstice Backup software performance, use the instructions in [“Client/Server Communications” on page 388](#).

Server Configuration

The following sections describe the various server configuration tasks.

▼ To Set Up the Server

To set up the Solstice Backup server:

1. Select Server>Setup Server to open the Server window.
2. Select View>Details to display hidden attributes.
3. Enter the number of save sessions backing up in parallel in the Parallelism attribute. See [“Parallelism” on page 46](#) for more information on this feature.

4. If you want to allow clients to make manual saves, select Enabled from the Manual Saves attribute.
5. If you want to allow a user retrieve archived data belonging to another author, select Enabled from the Public Archives attribute.
6. Your Solstice Backup server can also use the License Manager to manage its licenses if you enter the hostname of the License Manager server in the License Server attribute. For more information, refer to the *License Manager Administrator's Guide* on the documentation cd in your media kit.

Leave this attribute blank if not using the License Manager.

7. If you want the server to resynchronize its externally-managed licenses immediately, select Yes in the Update Licenses attribute.
8. Enter all of the addresses of those computers authorized to configure the Solstice Backup server in the Administrator attribute. See [“Server Access Management” on page 48](#) for more information.
9. Complete the Company Information attributes there. Providing this information is necessary for completing your registration form and obtaining your authorization code.

The product serial number is located on the Enabler Certificate you received from Sun Customer Service when you bought your Solstice Backup software. Enter the product serial number of the server only, not that of any add-on product.

10. Select an option from the Volume Priority attribute. Your choice determines the first priority of data to be written to volumes during a backup.
 - *NearLine Priority* gives autochanger volumes preference.
 - *OpenVault Priority* gives volumes managed by OpenVault preference.
11. Click Apply.

Performance Factor Controls

You can increase backup speed by setting the Solstice Backup server to *multiplex* data on a storage device. That is, data from more than one save set can be written to a single storage volume, and data from one save set can be spread across multiple volumes.

Note: Save sets that are multiplexed must, by definition, belong to the same pool of storage volumes.

Solstice Backup multiplexing optimizes and distributes the flow of data from multiple clients to all of the storage devices available to the Solstice Backup server. The parallelism and target sessions features allow you to control the Solstice Backup server's multiplexing functionality.

Parallelism

The Solstice Backup *parallelism* feature allows the Solstice Backup server to back up data from many clients in parallel, and to multiplex that data onto the same backup volume. The Parallelism attribute is located in the Server resource. See [“To Set Up the Server” on page 44](#) for more information on server configuration and the Set Up Server dialog box.

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 Solstice Backup software has a different default and maximum parallelism level. Each enabled storage node connected to the Solstice Backup server can increase the parallelism maximum. The maximum parallelism for any Solstice Backup server/storage node combination can vary, so refer to the documentation for that particular server for the maximum parallelism allowed. See [“Storage Nodes” on page 66](#) to learn more about storage nodes.

TABLE 3-1 displays the maximum parallelism levels for this release of the Solstice Backup server product, including the maximum for each enabled storage node.

TABLE 3-1 Maximum Parallelism Levels

| Solstice Backup Product | With Each Enabled Storage Node | Without Storage Nodes |
|-------------------------|---|-----------------------|
| WorkGroup Edition | Not applicable | Parallelism: 8 |
| | Not applicable | Devices: 2 |
| Network Edition | Parallelism: 32 Maximum = (32 x #nodes) + 32 | Parallelism: 32 |
| | Devices: 16 Maximum = (16 x #nodes) + 16 | Devices: 16 |
| Power Edition | Parallelism: 32 Maximum = (32 x #nodes) + 64 | Parallelism: 64 |
| | Devices: 16 Maximum = (16 x #nodes) + 32 | Devices: 32 |

To optimize 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:

- The server is taking too long to back up all the clients. Increase the parallelism.
- The server is so busy backing up clients that it is useless for any other tasks. Decrease the parallelism.

To “turn off” the parallelism so that only one client can back up to the server at a time, set the parallelism value to “1.”

Target Sessions

The Solstice Backup *target sessions* feature allows you to set the number of backup sessions accepted by an active device. The Target Sessions attribute is located in the Device resource. When you create or edit a device, you enter a value for the target sessions. Each device connected to a Solstice Backup server or storage node (including each individual device located in an autochanger or silo) can have a different target sessions value. See “Devices” on page 155 for more information about adding devices, and about the Device resource.

You can set the Target Sessions attribute 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 Solstice Backup server uses the next underutilized device for your backups.
- If all of the available devices are receiving the number of backup sessions configured by target sessions, the Solstice Backup server overrides the set value and uses the device with the least activity for the backup session.

Note: Recovery performance can suffer, because the data from one save set may be written to several volumes.

It is often more efficient for the Solstice Backup server to multiplex save sets to the same volume rather than write each save set to a separate device. For this reason, the Solstice Backup 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

You can add or delete Solstice Backup administrators for a given server at any time. Many Solstice Backup server tasks require the operator to have administrator privileges, so all those authorized to configure the server should have their addresses included in the server administrator list. Those who no longer have the right to administer the Solstice Backup server should be removed from the administrator list. See [“Administrator Permissions” on page 48](#) for more information.

You must have Solstice Backup Administrator privileges to add, modify, or delete Solstice Backup administrators. Every user or user group listed in the Administrator list box has permission to modify the Solstice Backup server configuration. If you are attempting to administer a Solstice Backup server and you receive the following error message:

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

you do not have permission to make configuration changes.

Administrator Permissions

Although any user can view the server's resources from a client computer, only users you specify in the Administrator attribute in the Server resource can add to or change the configuration of the Solstice Backup server, media devices, and autochangers. When you first install the Solstice Backup server, `root@server-name` is the only user authorized to change the Solstice Backup configuration.

To add other user IDs or computer names to the list of administrators, you must enter them in the Administrator attribute in the Server resource. See [“To Add Administrators” on page 49](#) for more information.

Valid entries in the Administrator attribute include:

- `user@hostname` or `hostname/user`
- `*@hostname` or `hostname/*`
- `user@*` or `*/user`

Note: If you use the `nsradmin` interface to input these entries, you must separate them by commas.

You can also add or restrict user privileges for individual clients. For more details, see [“Remote Access Rights for Other Clients” on page 94](#).

Caution – When you add a user to the Administrator attribute in the Server resource, that user has Solstice Backup administration privileges for only that Solstice Backup server. Administrators can change attributes for clients and other resources of the Solstice Backup server, but they have no special rights to client data for either backup or recovery.

▼ To Add Administrators

To add administrators to a Solstice Backup server:

1. Select Server>Server Setup to open the Server window.
2. Enter the user names you want to add in the Administrator attribute.
3. Click Add. The name you added appears in the Administrators attribute.
4. Click Apply.

▼ To Remove Administrators

To remove administrators from a Solstice Backup server:

Caution – Do not delete *all* administrators, especially *system@servername*. That administrator is required to start Solstice Backup services.

1. Select Server>Server Setup to open the Server window.
2. Select the administrator name you want to remove in the Administrator attribute.
3. Click Delete. The name you selected no longer appears in the Administrators attribute.
4. Click Apply.

Notifications

A *notification* describes a specific Solstice Backup event (which functions as a trigger) and a specific action that the Solstice Backup software implements in response to that event. The Solstice Backup server sends notifications to keep you informed about events taking place on the server.

The Solstice Backup software uses notifications to indicate which events to report, and how to report on them. Specific programs are set to be run when an event occurs. The Solstice Backup software sends the notifications to a variety of log files located in the */nsr/logs* directory. Depending upon the notification being sent, the Solstice Backup software might also display a notification in a Solstice Backup Administrator program status window.

You can either use the preconfigured notifications that come with your Solstice Backup server product, or you can create your own customized notifications. Third-party programs can also be used to send the information to other locations, such as an e-mail address or pager system.

Preconfigured Notifications

Solstice Backup software provides most of the notifications you need to stay well informed about Solstice Backup events.

[TABLE 3-2](#) lists the preconfigured notifications and their subsequent responses.

TABLE 3-2 Preconfigured Notifications

| Notification | Solstice Backup Response |
|-----------------------------|--|
| Bootstrap | Prints the bootstrap information on the Solstice Backup 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. |
| Index size | Sends e-mail to root: check the size of the client file index because it will soon exceed the space available. |

TABLE 3-2 Preconfigured Notifications (*Continued*)

| Notification | Solstice Backup Response |
|--|--|
| Log default | Directs the UNIX syslog facility (<code>/usr/bin/logger</code>) to log information from the <code>/nsr/logs/messages</code> 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 notifications through the configured trap. |
| Tape mount request 1 Tape mount request 2 Tape mount request 3 | Requests media be mounted in a device and: 1. Displays a pending message. 2. Sends e-mail. 3. Logs a message to <code>syslog</code> . |

Customizing Notifications

Notifications require three elements:

- Event
- Programs
- Priority

For example, every time a backup is completed, the Solstice Backup server sends a savegroup completion notification to the `/nsr/logs/savegrp.log` file that details which client save sets were backed up successfully. The server also displays this notification in the Messages window and the Group Control Details window. The associated priority helps you determine the importance of the event, so that you can take appropriate action.

Events

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

TABLE 3-3 Events

| Event Notification | Action Triggered |
|-----------------------------|---|
| 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 (for example, restarting the Solstice Backup software). |
| 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 is expired and needs replacing. |

Note: You cannot edit the events for an existing notification. To change either of these, you must delete the notification and recreate it with different settings.

Actions

TABLE 3-4 lists the programs that occur after an event notification takes place.

TABLE 3-4 Actions

| Programs | Description |
|----------|---|
| eventlog | Logs the notification message to the event Log. |
| nsrlog | Sends a message about an event to a file. |
| nsrlpr | Prints information to a printer. |
| nsrtrap | Sends notifications to a supported SNMP management console. |

Priorities

Each Solstice Backup notification has an associated priority. For example, the first time the Solstice Backup server sends a mount backup volume request, the priority assigned is “waiting”; the priority of the second request is “critical”; and the priority of the third request is “alert.”

[TABLE 3-5](#) lists the priorities upon which notifications are based.

TABLE 3-5 Priorities

| Priority | Description |
|-----------|---|
| Info | Supplies information about the server’s status. |
| Notice | Presents an important piece of information. |
| Warning | Presents information about a nonfatal error. |
| Waiting | Indicates that the server is waiting for an operator to perform a routine task. For example, 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 Solstice Backup software to fail, unless corrected immediately. |

Note: You cannot edit the priorities for an existing notification. To change either of these, you must delete the notification and recreate it with different settings.

Printing Notifications

The Solstice Backup server looks at two sources to determine which printer to use for printing: the printer that appears in the Action attribute in the Notifications window, and the printer that appears in the Printer attribute (displayed while using Details View mode) in the Groups window.

Note: The name of the printer used for printing by `nsrlpr` cannot contain spaces. Either use a shared printer with no spaces, or rename a shared printer so that its name has no spaces. You can test the `nsrlpr` program with a printer name as follows:

```
# nsrlpr -P printer_name text_file
```

where *printer_name* is the name of the selected printer and *text_file* is the name of the file you want to print.

If you receive these error messages:

```
Error: print server unreachable or specified printer does not
exist.

nsr1pr: The TCP/IP Print Server was stopped on exacto.

nsr1pr: Please start it.
```

then your TCP/IP print service may not be started. Restart these services on your server.

If you receive the following error message:

```
Error: print server did not accept request. Job aborted.
```

then your printer name is incorrect. Correct the printer name and try to print again.

Once you can print from the command line, change the Notification Action attribute to the print command:

```
# nsr1pr -P printer_name
```

Select Customize>Notifications and select the notification you want to print.

Note: You can also take advantage of other programs to notify you of an event, as long as the program supports reading from "standard input."

▼ To Create Notifications

To create a notification:

1. Select Customize>Notifications to open the Notifications window.
2. Click Create.
3. Enter the name of the new notification in the Name attribute.
4. Enter the UNIX command for the action to use for the notification in the Action attribute.
5. Complete any other attributes as necessary and click Apply.

▼ To Edit Notifications

You cannot edit the events for an existing notification. To change either of these, you must delete the notification and recreate it with different settings.

To edit a notification:

1. Select `Customize>Notifications` to open the Notifications window.
2. Select the notification to edit from the list displayed in the Notifications attribute.
3. Edit the attributes of the notification and click `Apply`.

▼ To Delete Notifications

To delete a notification:

1. Select `Customize>Notifications` to open the Notifications window.
2. Select the notification to delete from the list displayed in the Notifications attribute.
3. Click `Delete`.

You will be prompted to confirm that you really want to delete that notification. Click `OK` to confirm.

4. Click `Apply`.

Reports

The following sections describe various Solstice Backup reporting functionality you can use.

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:

- The success or failure of each save set participating in the operation
- The operation's save date and time
- The bootstrap *ssid*
- The bootstrap volume location (volume name, starting record number, and ending record number)

A second report, sent to the Solstice Backup server's designated default printer, repeats the bootstrap information as hard copy, which you should keep on-hand in a secure location. This printed report is the action triggered by the preconfigured notification *Bootstrap*.

Disaster recovery is much easier to perform if you have access to the bootstrap information in the most recent printed report. Refer to the *Disaster Recovery Guide* for further information on using the bootstrap during disaster recovery.

The `nsrinfo` program enables you to query the contents of the Solstice Backup client file index.

The `nsrwatch` program enables you to use a character-based interface to monitor Solstice Backup activity as it occurs.

Refer to the *Legato Command Reference Guide* and the man pages for these commands for more information.

Storage Management Application Reports

TABLE 3-6 lists the programs that the Solstice Backup software provides to query the contents of the storage management system. Refer to the *Legato Command Reference Guide* for descriptions of the most commonly used commands and options in more detail.

TABLE 3-6 Storage Management Report Programs

| Name | Report Generated |
|-----------------------|--|
| <code>mminfo</code> | The contents and mode of the storage volumes and/or the identification numbers and status of the stored save sets. |
| <code>mmlocate</code> | The user-defined location of storage volumes. |
| <code>nsrinfo</code> | The contents of the client file index. |
| <code>nsrmm</code> | The status of the storage devices known to the Solstice Backup software. |

Solstice Backup Server Statistics and Diagnostic Reports

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

Message Log Files

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

Firewall Support

Firewall support enables you to backup Solstice Backup clients and storage nodes that are separated from the Solstice Backup server by a *packet filtering* firewall. A packet filtering firewall looks at each packet entering or leaving the network and accepts or rejects it based on user-defined rules.

This section explains how Solstice Backup uses firewall ports, and provides guidelines on configuring these ports.

Service Ports

A service port is used by a Solstice Backup server or storage node to listen for client requests.

The default range used for service ports is 7937-9936.

Connection Ports

A connection port is used by the Solstice Backup processes to connect to the server during backup and recovery sessions.

The default range used for connection ports is 10001-30000.

▼ To Configure Firewall Support

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

To configure firewall support, you can either:

- Use Solstice Backup's default settings. To do this, refer to ["To Configure Firewall Support Using Default Settings"](#) on page 58.

Or

- Restrict the number of ports Solstice Backup uses. To do this, refer to ["To Restrict the Port Range Solstice Backup Uses"](#) on page 59.

▼ To Configure Firewall Support Using Default Settings

To configure the ports on the firewall, using Solstice Backup's default settings:

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

Caution – If the version of Solstice Backup Client Software, on the client computer or storage node, is earlier than Solstice Backup 5.5, you must 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 Solstice Backup host. To disable NAT refer to the documentation accompanying your router.
3. Shut down and restart the Solstice Backup server.
4. Ensure the **nsrexecd** service is running on the Solstice Backup server.

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 65.

▼ To Restrict the Port Range Solstice Backup Uses

To restrict the range of ports that Solstice Backup uses:

1. Determine the minimum port range required by each server, client and storage node. For more information, see [“Task 1: Determine Minimum Port Range Required” on page 59](#).
2. Change the default settings on the Solstice Backup server. For more information, see [“Task 2: Change Solstice Backup Default Settings” on page 63](#).
3. Complete the configuration of the firewall feature. For more information, see [“Task 3: Completing Configuration” on page 64](#).

Task 1: Determine Minimum Port Range Required

To determine the minimum service and connection port range required by each Solstice Backup server, client, and storage node, see:

- [“Server Port Range Requirements” on page 59](#)
- [“Client Port Range Requirements” on page 61](#)
- [“Storage Node Port Range Requirements” on page 62](#)

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 on determining the minimum number of service and connection ports required by a Solstice Backup server.

Service Ports:

The minimum number of service ports that you can assign to a Solstice Backup server is $10+n$, where:

- 10 is the minimum number of services registered on the Solstice Backup server.
If the number of services registered on the server is greater than 10, **nsrexecd** will use one port and the additional nine ports can be shared among the remaining services.
- n is the number of instances of the **nsrmmmd** service. An instance of the **nsrmmmd** service is allocated to each tape device configured on the server.

Example: Setting the Server's Service Port Range

If your configuration included four tape devices and 15 registered services, you would need to allocate 14 service ports to the Solstice Backup server. In this case, you could set the service port range for the Solstice Backup server, in the Solstice Backup Administrator program, to 8001-8014.

Connection Ports:

The minimum number of connection ports that you can assign to a Solstice Backup server is approximately the sum of the following:

- Number of:
 - Local and remote devices x 10.
 - Autochangers x 5.
 - Groups, or instances of **savegrp**, backing up to the server through the firewall x 5. For each **savegrp** instance that runs **nsrim**, an additional three ports will be required.
 - Solstice Backup Administration tools to be run remotely (such as **nwadmin**, **nwbackup**, and **nwrecover**) + 6.
 - Noncritical messages being posted.

The Solstice Backup server uses UDP ports for posting non-critical events. This number is dependent on the debug level and the number of errors. The larger the debug level set the more messages that will be posted. To turn debugging off, set the debug level to zero.

- Number of mechanisms used for copy protection.

The Solstice Backup Server uses UDP ports for copy protection. This number is dependent on the copy protection rules specified on the network.

- Parallelism value set in:
 - Server resource x 15.
 - Client or Server resource used during recovery sessions + 9.
 - 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.

Typically a server requires a minimum of 600 connection ports.

Example: Setting the Server's Connection Port Range

If the server required 600 connection ports, you could set the connection port range for the Solstice Backup server, in the Solstice Backup Administrator program, to 10001-10600.

Using the above server related examples, you would set the firewall rules to allow inbound and outbound TCP/UDP packets from the following ports:

- 8001-8014, for Solstice Backup server service ports
- 10001-10600, for Solstice Backup server connection ports
- 7937, for **nsrexecd**
- 7938, for Sun portmapper

Client Port Range Requirements

This section provides guidelines on determining the minimum number of service and connection ports required by a Solstice Backup client.

Service Ports:

The Solstice Backup client services do not bind to service ports as specified by the service ports range. Therefore you do not need to specify a service port range for a client.

The Solstice Backup client services, **nsrexecd** and Sun portmapper, will bind locally to ports 7937 and 7938.

Connection Ports:

The minimum number of connection ports that you can assign to a client is:

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

Typically a client requires a minimum of 100 connection ports.

Example: Setting the Client's Connection Port Range

If the client required 100 connection ports, you could set the connection port range for the Solstice Backup client, in the Solstice Backup Administrator program, to 10001-10100.

Using the above client example, you would set the firewall rules to allow inbound and outbound TCP/UDP packets from the following ports:

- 10001-10100, for Solstice Backup client connection ports
- 7937, for **nsrexecd**
- 7938, for Sun portmapper

Storage Node Port Range Requirements

This section provides guidelines on how to determine the minimum number of service and connection ports required by a Solstice Backup server.

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 than you would need to allocate 4 service ports to the Solstice Backup storage node. In this case, you could set the service port range for the Solstice Backup storage node, in the Solstice Backup Administrator program, to 8001-8004.

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

Typically a storage node requires a minimum of 150 connection ports.

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 Solstice Backup storage node, in the Solstice Backup Administrator program, to 10001-10150.

Using the above examples, you would set the firewall rules to allow inbound and outbound TCP/UDP packets from the following ports:

- 8001-8004, for Solstice Backup storage node service ports
- 10001-10150, for Solstice Backup storage node connection ports

- 7937, for **nsrexecd**
- 7938, for Sun portmapper

Task 2: Change Solstice Backup Default Settings

To configure the service and connection ports for the server, client and storage nodes, change the assigned default settings using **nwadmin** or **nsrports**.

▼ To Configure Service and Connection Ports using **nwadmin**

To configure the connection port and service port ranges using **nwadmin**:

1. Log on to the computer you want to configure ports for as *root*.
2. From the Solstice Backup Administrator program, select Options>Configure Ports.
3. In the Configure Ports dialog box, enter the client name (hostname or IP address) of the host and click OK.
4. In the Set Ports dialog box, enter a Service Ports range and a Connection Ports range.
5. Click OK in the Set Ports dialog box to save your changes.
6. Click Cancel to exit the Configure Ports dialog box.

▼ To Configure Service and Connection Ports using **nsrports**

To configure the connection port and service port ranges using **nsrports**:

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

- Use the **-s server** option to specify the system to contact.
- Use the **-a auth_server** to specify a Solstice Backup server. This option is required if **nsrports** is connecting to a remote system that is located on a different platform than the system on which the command is being executed.
- Use the **[-S]** option to set the system's service ports range to the specified range.
- Use the **[-C]** option to set the system's connection ports range to the specified range.

Task 3: Completing Configuration

To complete this configuration process:

1. Set the firewall rules to allow inbound and outbound TCP/UDP packets within the port range that you set on the Solstice Backup server.
2. Set the firewall rules to allow inbound and outbound packets from ports 7937 and 7938. These ports are used respectively by the Solstice Backup services **nsrexecd** and Sun portmapper.

Caution – If the version of Solstice Backup Client Software, on the client computer or storage node, is earlier than Solstice Backup 5.5, you must set the firewall rules to allow inbound and outbound packets from UDP 111 and 600-1023.

3. Ensure Network Address Translation (NAT) is disabled for each Solstice Backup host. To disable NAT refer to the documentation accompanying your router.
4. Shut down and restart the Solstice Backup server.
5. Ensure the **nsrexecd** service is running on the Solstice Backup server.

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 65](#).

Example: Sample rpcinfo Output

The sample `rpcinfo` output for Solstice Backup shown in [TABLE 3-7](#) shows port allocations for five devices in addition to the ten services required by the Solstice Backup server. A single port, 7937, is reserved for the **nsrexecd** program. Three ports, 7938–7940, are shared among nine other services. Five ports, 7941–7945, are allocated for the five **nsrmmmd** services. In total, nine ports are allocated for 15 services.

TABLE 3-7 Sample of rpcinfo Output for Solstice Backup

| Program Number | 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 | nsrmmdbd |
| 390107 | 5 | tcp | 7940 | nsrmmdbd |
| 390105 | 5 | tcp | 7940 | nsrindexd |
| 390104 | 105 | tcp | 7941 | nsrmmmd |
| 390104 | 205 | tcp | 7942 | nsrmmmd |
| 390104 | 305 | tcp | 7943 | nsrmmmd |
| 390104 | 405 | tcp | 7944 | nsrmmmd |
| 390104 | 505 | tcp | 7945 | nsrmmmd |

a. Programs that register themselves to the portmapper do not have to use their program name as the service name. For example, **nsrexecd** registers its service as the **nsrexec** service rather than the **nsrexecd** service.

Troubleshooting Port Configuration Problems

This section provides guidelines on troubleshooting common problems that 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 Solstice Backup 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 Solstice Backup 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 Solstice Backup server). A storage node has the physical connection and ownership of the attached devices, but the Solstice Backup server maintains the client file index and media database. Using the Solstice Backup software, you can route client data directly to a storage node's storage devices without the data first going to your Solstice Backup server.

From your Solstice Backup server, you can perform typical storage tasks, such as mounting and labeling volumes for the storage node devices, as well as configuring Solstice Backup resources associated with the storage nodes.

Note: Although any user can view the server's resources from any client computer, only users you specify in the Administrator attribute in the Server resource can add to or change the configuration of the Solstice Backup server, devices, and autochangers.

Prerequisites

To operate the Solstice Backup software with storage nodes, you must install the following software on your storage node:

- Solstice Backup client software (which includes the required storage node software). When you run the Solstice Backup installation script, you are given the opportunity to install storage node software (LGTONode).
- Device driver software for each remote device

Licensing

You must obtain and enter the following enabler codes on the Solstice Backup server to use the Solstice Backup software with storage nodes:

- Solstice Backup Server enabler code (included on the Server Enabler Certificate)
- 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 Solstice Backup 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 will have to 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 155](#). However, the storage node device name must follow this convention:

```
rd=storage_node_hostname:device_name
```

See [“Storage Nodes” on page 66](#) for more details on configuring storage node devices, including those in autochangers and silos.

The administrator of the storage node needs to be on the server's Administrator list (located in the Client attribute of the server's client resource) *only* 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 attribute.

For an autochanger or silo, you must manually add the storage node's hostname to the Administrator list before you define the devices with the `jbconfig` program. When the `jbconfig` program has finished running, you can remove the storage node's hostname from the Administrators list. If you need to configure a new autochanger later, you must add the hostname before you run the `jbconfig` program again. After you add the storage node's hostname to the Administrator list, one instance of `nsrmmmd` starts on the storage node for each device that it controls.

Caution – It is recommended that you do not attempt to convert an existing Solstice Backup server to a Solstice Backup 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 Solstice Backup server.

Installing Storage Node Software

The following sections provide information when installing new storage node software:

- On the Solstice Backup server and a new storage node computer
- On the Solstice Backup server only
- On the Solstice Backup server and an existing storage node

Installing on the Server and Creating a New Storage Node

When you install the storage node software to create a new storage node, first update your server to the same Solstice Backup release. Then, run the `jbconfig` program on the storage node computer to add a Jukebox resource. After you run the `jbconfig` program on the storage node computer, the Jukebox resource on the server displays values that the `jbconfig` program automatically selects based on the type of autochanger you configure.

Installing on the Server Only

When you update just the Solstice Backup server with the storage node software, existing storage nodes continue to operate as usual. However, these storage nodes are unable to use the features available with the updated `nsrjb` program.

Installing on the Server and Existing Storage Node

When 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.

After you install the software on the storage node, the Solstice Backup server's Jukebox resource configured for use by the storage node contains the new attribute values. See [“Devices and Autochangers” on page 400](#) for a description of the new attributes.

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 Solstice Backup 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 Node Affinity

The link between a Client resource and a list of storage nodes is called *storage node affinity*. You define storage node affinity in the Storage Nodes attribute in the Client resource. The default setting for the Storage Nodes attribute on most Solstice Backup Client resources is the Solstice Backup server. For the Client resource of a storage node computer, the default setting of the Storage Nodes attribute is the storage node and the Solstice Backup server. You can add the names of other storage nodes to the list. The Solstice Backup server uses the list in the Storage Nodes attribute to determine which device writes the data from each savestream.

During backup, only the devices attached to the storage node computer in the Storage Nodes attribute list are considered to receive that client's data. You cannot specify a different list of storage nodes to be used for different operations, but you can add and remove storage node names from the Storage Nodes attribute in the Clients resource at any time.

Storage Nodes and Cloning

Storage nodes can clone data in a manner similar to Solstice Backup servers. A hierarchical structure exists between the storage node and the Solstice Backup server to determine where storage node data is routed to be cloned.

The following list presents this hierarchy, from highest to lowest:

1. Cloning is defined by the storage node's Clone Storage Nodes attribute, also located in the Client resource of the Solstice Backup server:

The Client resource contain the Clone Storage Nodes attribute for cloning data sent to a storage node. Enter the name of each storage node you want to be available for cloning the storage node's data.

Like the Storage Nodes attribute previously described, the storage node will direct its data to the first storage node in the list that has an enabled device and is able to receive the data for cloning.

Unlike the Storage Nodes attribute, however, there is no default storage node name present here. If there are no storage nodes listed, the Solstice Backup software consults the server for cloning guidance.

2. Cloning is defined by the Clone Storage Node attribute:

If no storage node is listed in the Clone Storage Node attribute, the Solstice Backup software looks to the Clone Storage Nodes attribute for cloning instructions. The storage node(s) listed in the Clone Storage Nodes attribute determines where the data is cloned. If there is nothing listed in the Clone Storage Nodes attribute, the Solstice Backup software consults the Client resource's Storage Nodes attribute.

3. Cloning is defined by the Storage Node attribute:

If no storage node is listed in the Clone Storage Nodes attribute, the Solstice Backup software looks to the Storage Nodes attribute. Here the Solstice Backup software can either send the data to the default storage node *nrsserverhost* or to any other storage node you specify.

Online Index Structures and Management

The Solstice Backup software tracks the files it backs up in the client file indexes and the media database.

The client file indexes tracks the files that belong to a save set.

The media database tracks the following:

- Name of the volume,
- Backup dates of the save sets on the volume

- Filesystems in each save set

The Solstice Backup software 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 110](#).

The following sections provide information on the following online index topics:

- Understanding the online indexes and how they work
- Specifying and moving the location of the indexes
- Viewing index information

A Note on Upgrading from Solstice Backup 5.x Software to This Release

If you are updating from Solstice Backup release 5.x software to this release, your existing online indexes will be converted to the new model. For information specific to this conversion, see [“Index Conversion Issues” on page 78](#). Also refer to your *Installation Guide* and *Release Supplement* for any additional information.

Understanding 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. The database stores one entry for each volume, while the client file index stores one entry for each file saved on that volume. The Solstice Backup software determines which volume to mount for recovering a file by mapping the saved files to their volumes.

The online indexes grow with each backup, as entries are added for each newly backed-up file and saveset. As long as an index entry for a file remains in the client file index, you can recover the file. The structure of the client file indexes avoids operating system restrictions on file size and allows the client file index for a single client to continue to grow.

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

Automated Solstice Backup Index Actions

The Solstice Backup software performs the following online index actions:

- Inserts entries in the client file index for each file saved during a backup. For each new backup, the Solstice Backup software acquires more space from the filesystem for the new entries.
- Removes entries to free up space in the client file index. The Browse and Retention policies automatically determine when entries are removed from the index. You can also remove save sets manually by clicking Remove Oldest Cycle in the Index Save Sets dialog box.
- Immediately returns unused space to the filesystem, where it is available for other uses.

Client File Index Location Issues

During the initial client setup, the Solstice Backup software will normally designate a default location for the client's client file index on the Solstice Backup 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 a New Client's File Index Location

However, you can elect to designate another location for a client's index location using the Index Path attribute in the Client resource. This procedure is only possible when initially configuring a new Client resource, or before the first backup of a client occurs.

If you want to change the location of an existing client's file index, you must move that index. See ["To Move an Existing Client File Index's Location"](#) on page 73 for more information.

To designate a new client's file index location:

1. Create a new directory to store the client's file index in the Solstice Backup server (if the directory does not already exist).

2. Select Clients>Client Setup to open the Clients window.
3. Select View>Details to display hidden attributes.
4. Enter the full path of the location you specified in the Solstice Backup server in step 1 in the Index Path attribute, and click Apply.

▼ To Move an Existing Client File Index's Location

In some cases you might want to move a client's file index from its existing location to a new location you specify. For example, if the size of the client file index is too large to exist in its current location.

To move an existing client file index's location:

1. Create a new directory to store the client's client file index in the Solstice Backup server (if the directory does not already exist).
2. Select Clients>Client Setup to open the Clients window.
3. Select View>Details to display hidden attributes.
4. Enter the full path of the location you specified in the Solstice Backup server in step 1 in the Index Path attribute and click Apply.
5. Shutdown the Solstice Backup daemons and services.
6. Move the existing client file index to the new location you want by entering the following command at the prompt:

```
# mv /nsr/index/client_name /new_index_location
```

Make sure that the new location you select can sustain the current and anticipated size of the index.

7. Restart the Solstice Backup services and daemons.
8. Verify that the new directory contains the same number of indexes by entering the following command at the prompt:

```
# nsrinfo client_name
```

The Solstice Backup server will save all future client file index information to the new location you designated.

Note: If you move a client's file index to a new location, you will not be able to recover client file indexes from backups that occurred using Solstice Backup release 5.x or earlier software.

Accessing Online Index Information

You can access information on a variety of online index-related subjects within the Solstice Backup Administrator program. The following sections provide instructions on accessing this information.

▼ To View Index Information

Select Clients>Index or click the Indexes toolbar button to open the Indexes window.

The Clients attribute of the Indexes window displays the following information for each index:

- Name – the Solstice Backup client names.
- Size – the allocated size of the client file index. The allocated disk space automatically grows as the index size increases.
- Used – this category always displays a value of 100%, as the Solstice Backup software no longer uses disk space reclamation.

▼ To View Save Set Information

To view save set information:

1. Select Clients>Index or click the Indexes toolbar button to open the Indexes window.
2. The Indexes window appears displaying each client in the Clients list.
3. Select the client whose save sets you want to view in the Clients list.

The save set information for that client appears in the Save Sets list displayed in the lower part of the Indexes window.

The Save Sets list of the Indexes window displays the following information for each save set:

- Name – the save set name.
- Size – an estimated amount of the index space used by the save set in the client file index.
- Cycles – number of back-up cycles available for browsing. A cycle starts with a full backup and ends with the next full backup. A cycle includes the incremental and level 1 – 9 between the full backups (if any).

▼ To View Details of Each Instance of a Save Set Backup

The Instances window shows the instances, or versions, of a save set backup at different times. You can view details of the instances of any save set listed in the Indexes window. Use this information to determine how your resources are being

used. For example, you might need to see how large a save set is to plan the amount of disk space you need for the online indexes. You can also use the print feature of this window to generate a report that details your Solstice Backup client backups.

The following information is displayed for each save set instance:

- ID – the unique identification number of that particular save set.
- Files – the number of files backup up during that instance
- Size – the size of the backup.
- Date – the date of the backup.
- Level – the backup level.

Managing the Size of the Online Indexes

Over time, the size of your Solstice Backup server's online indexes can become prohibitively large. You can reduce the size of these indexes using the solutions suggested below.

▼ To Reduce Client File Index Size

You can reduce the size of your Solstice Backup 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. See [“To Remove the Oldest Save Set Cycles” on page 76](#) for more information.
- Remove volume-based entries from the client file index. See [“To Remove Volume-Based Entries from the Online Indexes” on page 77](#) for more information.
- Adjust the Browse Policy and Retention Policy attributes of clients backing to the Solstice Backup server to shorten the period of time 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, 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. See [“Browse and Retention Save Set Policy Modification” on page 115](#) for more information.

If the size of the client's online file index is still too large, consider moving the location of the index. By default the client's online file index is located in the `/nsr/index/client_name` directory. See [“To Move an Existing Client File Index's Location” on page 73](#) for more information.

▼ To Reduce Media Database Size

You can reduce the size of Solstice Backup server's media database by using one or more of the following procedures:

- Remove volumes that contain recyclable save sets from the Solstice Backup inventory. See ["To Remove Volume-Based Entries from the Online Indexes"](#) on page 77 for more information.

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. Leaving a volume's index entries in the media database prevents you from accidentally labeling another volume with the same name.

- Recycle volumes that contain recyclable save sets. See ["To Change the Mode of a Volume"](#) on page 176 for more information.

When a volume is recycled, it undergoes 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, using the `nsrmm -e` command. See ["Browse and Retention Save Set Policy Modification"](#) on page 115 for more information.

- Compress the media database. See ["To Compress the Media Database"](#) on page 78 for more information.

▼ To Remove the Oldest Save Set Cycles

To decrease the amount of space the index uses, use the Remove Oldest Cycle feature. 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 so that other entries can fit into the empty space without increasing the size of the index. You cannot remove a save set with fewer than two cycles using the Remove Oldest Cycle feature. The last full cycle might be needed for recovery.

After the Remove Oldest Cycle operation has finished, the statistics in the Index window are updated to reflect the current state of the file index.

Note: You cannot simultaneously reclaim space and remove the oldest cycle for a client, because the index is busy during either operation.

To remove the oldest save set cycles:

1. Select Clients>Indexes or click the Indexes toolbar button to open the Indexes window.
2. Select the save sets whose oldest cycle you want to remove and click Remove Oldest Cycle.

▼ To Remove 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 Solstice Backup configuration. See [“To Remove Bad Volumes from the Online Indexes” on page 177](#) for an explanation of how the Solstice Backup software handles bad volumes.

You can also use this feature to reduce the size of the online indexes by purging index entries associated with specific volumes.

The steps to remove index entries are the same regardless of whether you are removing a bad volume, trying to reduce index size, or both. The following sections provide information on how to remove volume-based entries from the online indexes.

To remove volume-based entries from the online indexes:

1. Select Media>Volumes or click the Volumes toolbar button to open the Volumes window.
2. Select the volume whose entry you want to remove from the media database.
3. Select Volume>Remove to open the Remove Volume window. Pick one of the following volume entry removal options:
 - File and Media Index Entries
 - File Index Entries Only

The next section describe each option.

4. Click OK.

The Solstice Backup 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 short period of time.

■ Removing Client File Index Entries

With this option, you remove just the entries contained in the client file index, which changes the status of the browsable save sets to recoverable. When save sets are marked recoverable, users may no longer browse and recover these files.

You would need to use the save set recover procedure to recover data once the client file index entries are removed. See [“Save Set Recovery” on page 229](#) for information on this procedure.

- Removing Client File Index and Media Database Entries

With this option, you remove both the client file index and media database entries, which removes all knowledge of the volume from the Solstice Backup software. 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 using the scanner program. See [“Scanner Program” on page 240](#) for more information on the scanner program.

Typically, you do not want to remove both the client file index and media database entries together 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 Solstice Backup software will simply access the cloned volume as needed, rather than accessing the original volume. The volumes’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 Compress the Media Database

You can free up more space on your server by compressing the media database.

To compress the media database:

1. Delete the file `/nsr/mm/.cmprssd`
2. Enter the following command at the prompt:

```
# nsrim
```

Index Conversion Issues

If you are updating from Solstice Backup release 5.x software to this release, you will be converting the media database and client file indexes to a new model. While this conversion is automatic as part of the installation, the information contained in this section can help you ensure the conversion is both successful and quick.

Changes in Disk Space Requirement from Previous Solstice Backup Releases

Once the conversion to Solstice Backup 6.x software is complete, the media database on the Solstice Backup server requires twice the disk space as it did in 5.x versions. The client file index will require 25 to 40% less space. This change in size is permanent.

Caution – It is strongly recommended that before updating your Solstice Backup software that you complete a successful, scheduled backup (not a manual backup) of the Solstice Backup server and clients.

Enhancing Client File Index Conversion

When you update Solstice Backup software from release 5.x to this release, the Solstice Backup server performs the following tasks:

- Converts the media database to the new model
- Converts each client file index to the new model

You cannot initiate any backups until the media database conversion is finished. Following its completion, you can back up data and perform recoveries of save sets. You cannot browse and select specific files for recovery until the corresponding client's file index is converted.

The Solstice Backup server will begin converting the client file indexes to the new model automatically using the `nsrck` program once the media database's conversion has finished. However, the default method of client file index conversion is to perform them serially, one client at a time. A Solstice Backup backup server with a large number of clients can take an extended period of time to convert all the client file indexes using this method.

Converting a Specific Client File Index

Even after the Solstice Backup server has started the process of converting the client file indexes one at a time, you can start the conversion of any specific client file index manually by starting an additional instance of the `nsrck` program to convert the index you indicate.

To start the index conversion of a specific client, enter the following command at the prompt:

```
# nsrck client_name
```

Once the conversion of that client file index is completed, you can browse the client file index and select specific items for recovery.

Converting Multiple Client File Indexes Simultaneously

You can also elect to start multiple client file index conversion in parallel by running multiple instances of the `nsrck` program simultaneously. Each instance of `nsrck` running will convert client file indexes serially until all the indexes are converted. Therefore, three instances of `nsrck` running in parallel could convert the client file indexes three times as fast as the default single instance, greatly speeding up the overall conversion process.

To start an instance of `nsrck`, enter the following command at the prompt:

```
# nsrck
```

Repeat this step for each instance of `nsrck` you want to run. By entering the `nsrck` command without specifying a particular client, each instance of the `nsrck` program running will look for the next unconverted client file index and convert it until no more are left.

CPU Usage for `nsrck` Instances

Each instance of `nsrck` running occupies a portion of your servers processing, memory, and space swapping capability. You must balance the speed with which you want to convert your indexes with other Solstice Backup server needs. We also suggest that you monitor your system performance using your favorite performance diagnostic tool to determine the optimum number of instances of `nsrck` to launch in parallel. On average, each instance of `nsrck` actively converting client file indexes consumes 100% of up to two CPUs in a Solstice Backup server. Additional instances of `nsrck` consume even more CPU power.

Server Maintenance Tasks

This section describes tasks you might need to perform after you install and configure your Solstice Backup server.

Message Log Management

The Solstice Backup server stores the messages generated by the Solstice Backup server daemons in a message log file in the `/nsr/logs` directory. When the log file becomes too large, you must delete some messages from the log. To automatically control the size of the log, you can use variables in the Solstice Backup startup script in the `/etc` directory or create a script that uses the operating system services.

▼ To Set the Startup Script to Trim Log Files

To modify the way that Solstice Backup services manage the Solstice Backup log files, change the following environmental variables in the Solstice Backup startup script, `/etc/init.d/networker`, before you start the `nsrd` daemon:

- To change the maximum size of log files, change the `NSR_MAXLOGSIZE` value. The default value for `NSR_MAXLOGSIZE` is 1024 KB.
- To change the maximum number of log files that are saved, change the `NSR_MAXLOGVERS` value. The default value is 4.

Every time the Solstice Backup server starts, it checks the size of the `daemon.log` file. By default, when the `daemon.log` file reaches the 1024 KB limit, it 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 see whether the log file has exceeded `NSR_MAXLOGSIZE`. If `nsrd` runs for a long time, the log file can still grow very large. To activate the trimming mechanism, enter `nsr_shutdown` to stop the Solstice Backup daemons, and then restart the `nsrexecd` and `nsrd` daemons.

▼ To Use To Solaris Services to Trim Log Files

You can use Solaris services to automatically manage the size of the Solstice Backup log files. The following example uses the operating system services available on the Solaris platform.

Solaris systems provide a two-part mechanism for managing the `syslog` message file (`/var/log/syslog`): a shell script (`/usr/lib/newsyslog`) and a **crontab** entry for root to periodically invoke the script.

You can modify the **newsyslog** script to manage and maintain a short history of the Solstice Backup server's log file. The modified script maintains a three-file history of the Solstice Backup server's `daemon.log` file.

To manage your Solstice Backup log file:

1. Use your favorite text editor to add the following lines to `/usr/lib/newsyslog`:

```
LOGDIR=/nsr/logs

LOG=daemon.log

if test -d $LOGDIR

then

    cd $LOGDIR

    test -f $LOG.1 && mv $LOG.1 $LOG.2

    test -f $LOG.0 && mv $LOG.0 $LOG.1

    test -f $LOG    && mv $LOG    $LOG.0

    cp /dev/null $LOG

    chmod 644 $LOG

fi
```

2. Shut down the daemons with the `nsr_shutdown` command, either manually or as an additional command in the `newsyslog` script.

The Solstice Backup server cannot use the new log file until you shut down and restart the Solstice Backup daemons. Make sure that the script does not run during a scheduled save.

3. Restart Solstice Backup using the following command:

```
# /etc/init.d/networker start
```

4. Add an entry to the `crontab` for root to control the frequency of running the `newsyslog` script. The entry shown in the following example invokes the `newsyslog` script every Saturday morning at 4:05 a.m., for example:

```
5 4 * * 6 /usr/lib/newsyslog
```

If you are running a system that does not have the `newsyslog` script and `crontab` entry to invoke it, create the `newsyslog` script manually and add the `crontab` entry for it. See the `crontab` man page for details on creating these entries.

▼ To Move Your Solstice Backup Server Software to a Different Computer

This section describes how to move the Solstice Backup server software from one computer to another:

1. Perform a full Solstice Backup backup of all the filesystems on the old Solstice Backup server.
2. Shut down the Solstice Backup daemons on the old server, using the `nsr_shutdown -a` command.
3. Make a `tar` tape of the entire `/nsr` directory from the old server, and reload it on the new server. If `/nsr` is a symbolic link on the old server, make sure that the new server has the `/nsr` symbolic link set up also.
4. Shut down your old server and disconnect all the devices.
5. Shut down the new computer, add the hardware devices to the new server, and restart both computers. Start the old computer first, and then the new one.
6. Install the Solstice Backup server software on the new server.

Note: If you have an autochanger, do not select the option to start the Solstice Backup daemons. Refer to the instructions in the *Installation Guide* to learn how to install and test the Solstice Backup device drivers.

▼ To Define the Index Entry

Because you created a new host, you must correctly define the index entry for the new host before you start the Solstice Backup daemons. There are two ways to define the index entry:

- Name the new server with the same hostname as the old server at the operating system level before you modify client resources.
- Create a new hostname for the new server with the same configuration choices as the old server.

To create a hostname for the new server:

1. Create a new hostname for the new server with the same configuration choices as the old server.
2. Delete the hostname entry for the old server.
3. Shut down the Solstice Backup daemons on the old server and the new server with:

```
# nsr_shutdown -a
```

4. Change to the directory containing the old server index entry with:

```
# cd /nsr/index
```

The entry for the new server hostname is empty.

5. Delete the entry for the new server hostname, as in:

```
# rmdir new_hostname
```

You must remove this entry, or the next step creates a subentry for the new server instead of the correct entry.

6. Rename the old index directory to the new server hostname, as in:

```
# mv old_hostname new_hostname
```

The Solstice Backup daemons start up on the new server.

You see the following messages on the new server:

```
new-server syslog: NetWorker Server: (notice) started

new-server syslog: NetWorker Registration: (notice) invalid auth
codes detected.

new-server syslog:

new-server syslog: The auth codes for the following licenses enablers
are now invalid.

new-server syslog: The cause might be that you moved the NetWorker
server to a new computer.

new-server syslog: You must re-register these enablers within 15 days
to obtain new codes.

new-server syslog:

new-server syslog: License enabler #xxxxxxx-xxxxxxx-xxxxxxx (NetWorker
Advanced/10)
```

▼ To Complete the Move

Once you have moved your software, you must reregister your Solstice Backup server. After you move the Solstice Backup server from one system to another, you have 15 days to register the new server with Sun. To reregister your new server, contact the Sun License Center and request a Host Transfer Affidavit. You must then complete and return the Host Transfer Affidavit to Sun. You will then receive a new authorization code, which you must enter into the Auth Code field of the Registration window. For information about contacting the Sun License Center, see [“Licensing” on page 39](#).

After you have successfully moved your server:

1. Verify that all the clients are included in the scheduled backups.
2. Use the Solstice Backup `recover` program to make sure that all the client indexes are visible and, therefore, recoverable.
3. Back up the indexes on the new server or perform a full backup of the new server as soon as possible.

▼ To Set Up an Old Solstice Backup Server as a Client

If you want to set up the old server as a client, first remove all the Solstice Backup software and the `/nsr` directory from the old server, and then reinstall the Solstice Backup client software.

▼ To Recover Bootstrap Data from Another Server's Tapes Using `mmrecov`

To recover a Solstice Backup 6.x server's bootstrap (media database and resource files) to another server:

1. Move the physical media to the recovering server.
2. Run `mmrecov` to recover the media database and resource files.
3. Shut down the Solstice Backup server and replace the default resource files with those that were recovered by copying or moving them from `/nsr/res.R` to `/nsr/res`.
4. Add the recovering server to the old server's remote access' list.
5. Run `nsrck -L7 old_server` to recover the server's indexes.

The command line program `mmrecov` will now recover the media index and server's resource files independent of the server's file index. Refer to the *Legato Command Reference Guide* and the `mmrecov` man page for more information about this command.

Device Parameter Environmental Variables

A number of new environmental variables are now available to adjust several Solstice Backup device parameters. For example, the amount of time allotted for a device to become ready after the loading of a volume is now a variable.

These new environmental variables are:

- `NSR_DEV_BLOCK_SIZE_`
- `NSR_DEV_TAPE_FILE_SIZE_`
- `NSR_DEV_DEFAULT_CAPACITY_`
- `NSR_DEV_LOAD_TIME_`

NSR_DEV_BLOCK_SIZE_

NSR_DEV_BLOCK_SIZE_ is organized in units of KB. The value set must be a multiple of 32KB, with a minimum value of 32KB. Maximums are determined by platform, SCSI driver, and device.

NSR_DEV_TAPE_FILE_SIZE_

NSR_DEV_TAPE_FILE_SIZE_ is organized in units of NSR_DEV_BLOCK_SIZE_ 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.

- the Solstice Backup 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.the Solstice Backup software

NSR_DEV_DEFAULT_CAPACITY_

NSR_DEV_DEFAULT_CAPACITY_ 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.

NSR_DEV_LOAD_TIME_

NSR_DEV_LOAD_TIME_ is the number of seconds that nsrmmnd polls and wait for the drive to become ready after the library inserts a tape into the device. If this value is too short, there may be unnecessary load failures. If it is too long, then labeling brand new tapes will take longer. The minimum allowable value is 10 seconds, the maximum value is 600 seconds.

Client Operations

This chapter describes how to configure and use Solstice Backup clients, and gives suggestions on how to best customize your client configurations to suit the needs of your environment.

Solstice Backup Clients

A Solstice Backup client is both a physical computer with Solstice Backup client software installed on it, and a Solstice Backup resource that specifies a set of files and directories that are included in a scheduled backup. As such, a single Solstice Backup client computer can have several individual Client resources specified that back up to the same, or even different Solstice Backup servers.

The concept of a client computer with multiple Solstice Backup Client resources is further explained in [“Multiple Clients from the Same Computer” on page 100](#).

Solstice Backup client software is available that can work with a variety of platforms. Solstice Backup products are heterogeneous, meaning that clients and servers of different platforms and operating systems can work together. The basic Solstice Backup client software contains backup and recovery capability for filesystem data. You can purchase additional modules to enable archiving, and backup of a variety of databases.

Client Configuration

This section provides an overview of how Solstice Backup clients are created and integrated into the storage management environment.

The Solstice Backup client software attempts to establish a connection with the server. Before accepting the connection, the Solstice Backup server verifies that the user who initiated the program has the necessary permissions to save or recover client files.

After you install the software on the Solstice Backup client computer, you create a Client resource on the Solstice Backup server that specifies your configuration choices for each Solstice Backup client. These choices determine the following:

- What data is backed up?
- Which schedule is being followed?
- What additional features, such as archiving, are enabled?

▼ To Create Clients

Before a client can be backed up by a Solstice Backup server, the client computer must have the appropriate Solstice Backup client software installed. Refer to the *Installation Guide* for instructions on installing the client software.

If you choose not to customize the configuration choices, the new Solstice Backup client is automatically assigned the default configuration. The default setting *All* for the Save Set attribute means that all the files on the client are backed up during a scheduled or manual backup.

Note: For UNIX clients, setting the Save Set attribute to All will not back up all locally mounted file systems. The only file systems backed up will be those listed in the following locations:

- Solaris and Sequent: `/etc/vfstab`
- HP-UX, Compaq Tru64, Linux, and SGI: `/etc/fstab`
- AIX: `/etc/filesystems`

To create a client:

1. Select Clients>Client Setup to open the Clients window.
2. Click Create.
3. Enter the client's new hostname in the Name attribute.
4. If you installed the Archive Application, the Archive Services attribute is automatically enabled. Otherwise, the Archive Services attribute is disabled.
5. Select a schedule in the Schedule attribute. See [“Schedules” on page 193](#) for more information.
6. Select a browse policy in the Browse Policy attribute. A typical default choice is *Quarter*. See [“Browse and Retention Policies” on page 104](#) for more information.

7. Select a retention policy in the Retention Policy attribute. A good choice is *Year*. See [“Browse and Retention Policies” on page 104](#) for more information.
8. Select a directive in the Directives attribute. See [“Directives” on page 127](#) for more information.
9. Select a group in the Groups attribute to which your client will be added. See [“Backup Groups” on page 182](#) for more information.
10. Enter the save sets that you want backed on this client in the Save Set attribute.
 - If you want all the save sets backed up, enter *All*.
 - 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.
11. Enter the user ID or hostnames of other clients that are allowed to back up or recover this client’s files in the Remote Access attribute. Leave this attribute blank if you only want this client to have those rights.
12. Enter the user ID of a client to allow remote access and operations on this client in the Remote User attribute.

The Password attribute is only used for computers running NetWare. The password entered in this attribute is used to perform remote commands on this client.
13. Enter the `savepnc` command, or the name of a program file you created that includes commands to execute before and after the Solstice Backup software backs up client data, in the Backup Command attribute.
14. Enter the nickname(s) for the client computers that queries can match in the Aliases attribute. Otherwise, queries match only the client name.
15. The Server Network Interface (SNI) field is for specifying the Network Interface on the server that connections for this client instance should attempt to use, in the instance of a private backup network for example.

The SNI should specify the unique hostname associated with the single network interface on the server to be used for saves and recovers. Be sure when using an SNI to include this as the server that will backup the client when configuring the client.

If left blank, NetWorker will default to the NIC associated with the servers hostname.
16. Enter each user that is authorized to use the archive services in the Archive Users attribute. If no users are listed, then only administrators are allowed to use the archive services on the client. The “\” and “@” characters are not allowed in the user names for this attribute.

17. Enter the name of each storage node to which the client can back up data in the Storage Nodes attribute. The client backs up to each active enabled storage node in the order they appear in the attribute. The default storage node *nsrserverhost* represents the server.

18. Enter the storage node names for cloned data in the same manner as step 17.

19. Complete any other attributes as necessary, and click Apply.

If you try to add more clients beyond the number of connections you purchased for your Solstice Backup server, you receive the error message:

Too many clients - maximum is *number*.

▼ To Edit Clients

To edit a client:

1. Select Clients>Client Setup to open the Clients window.
2. Select the client you want to edit listed from the list displayed in the Clients attribute.
3. Edit the attributes of the client, and click Apply.

You cannot change the name of a client using this process. To change a client's name, see ["To Change a Client's Name" on page 93](#).

▼ To Delete Clients

When you delete a client, the Solstice Backup 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 Solstice Backup server, defined as a client of itself.

To delete a client:

1. Select Clients>Client Setup to open the Clients window.
2. Select the client you want to delete from the list displayed in the Clients attribute.
3. Click Delete.

You will be prompted to confirm that you really want to delete that client. Click OK to confirm. The client is removed from the list of clients displayed.

4. Click Apply.

Note that even if you delete a client from the Solstice Backup server, the client's previously backed-up data is still accessible and recoverable directly from the volume containing the data using the scanner program. See [“Recovering Client File Indexes” on page 238](#) for more information.

▼ To Access 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. Under normal conditions you should leave this attribute empty when creating or editing a Client resource. Leaving it blank will cause the server to automatically assign a value.

The Solstice Backup 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, such as if all Client resources of the same name have the same ID.

If you create a Client resource, back it up a few time, and then create a new Client resource with the same client name, the Solstice Backup server recalls the client ID used previously for this name and reuses it for the new resource.

The only time you would typically need to access the client ID value of an existing client, or entering a value in the Client ID attribute, is for changing the name of a client or restoring a deleted client. The following sections address each of those tasks.

To access the client ID of a client:

1. Select Clients>Client Setup to open the Clients window.
2. Select View>Details to display hidden attributes.

▼ To Change a Client's Name

To change a client's name:

1. Note the Client ID of the original client.
1. Delete all instances of the original client.

2. Create a new Client resource. See [“To Create Clients” on page 90](#) for more information.
3. Enter the original client’s Client ID value in the Client ID attribute.
4. Complete the other attributes as necessary, and click Apply.

See [“To Access the Client ID” on page 93](#) for more information on understanding and accessing the Client ID attribute of the Client resource.

Example: Changing a Client’s Name

For instance, to change the name from *jupiter* to *zeus*, you would remove all instances of client *jupiter* and create the client *zeus* filling in the Client ID attribute of the Client resource with the value `74684a4a-00000004-87450361-37e8222d-37e8222d-0001000`.

▼ To Restore a Deleted Client

If you want to restore a client that has been deleted from the Solstice Backup server, recreate the Client resource using the same client name and the same attribute values as the deleted one. See [“To Create Clients” on page 90](#) for information creating clients. The Solstice Backup server remembers the client ID for the Client resource bearing this name, and inserts it into the Client ID attribute.

Remote Access Rights for Other Clients

Solstice Backup clients are preconfigured so that only the client itself can browse or recover its own files. However, you can authorize individuals to have remote access rights by entering their user ID in the Remote Access attribute of a client’s Client resource. When you enable remote access rights, authorized users can view, back up and recover files from other Solstice Backup clients.

To grant remote access rights:

1. Open the Client resource of your client.
2. Enter one of the following in the Remote Access attribute:
 - User ID and hostname (in the format `user@hostname` or `hostname/user`)
 - Netgroup name (if you are using NIS) in the Remote Access attribute in the Client resource.

Note: If your company is concerned about security, leave the Remote Access attribute blank, so that only the client itself can recover its backed-up files.

Remote User Attribute

This attribute has several uses. For those clients that are accessed through the `rsh(1)` protocol (new clients use `nsrexecd` instead), this attribute specifies the user login name the NetWorker server will use to authenticate itself with the client. The default value is `NULL`, meaning that `'root'` should be used.

When you run `savegrp -p` on the NetWorker server, the server runs commands on the client to determine which files to save. Note that when the `nsrexecd` protocol is used to access the client, the Remote User attribute is not used for authentication.

Certain clients, such as NetWare file servers, use this attribute, along with the password attribute, to access the files being backed up. Other clients that back up application data, such as Sybase databases, use this attribute along with the password to gain access to the application data. There might be a different value of this attribute for each resource that describes the same client.

See the remote access attribute under `NSR client` in the *Legato Command Reference Guide* for more information.

Solstice Backup Clients of Different Platforms

The Solstice Backup server can back up clients from a variety of platforms. This section provides configuration tips for configuring clients to enable them to back up to the Solstice Backup server.

To use clients of an operating system different from your Solstice Backup server, you must purchase and enable the appropriate ClientPak. See [“Solstice Backup Client Licensing” on page 133](#) for information about ClientPaks and how the Solstice Backup server checks each client before it allows a backup to begin.

Support for 64-bit filesystems exists for clients that run Solaris 7 or later, AIX 4.2, and HP-UX 10.20. You can archive, back up, browse, and recover files larger than two gigabytes for clients of Solaris 2.6, AIX 4.2, and HP-UX 10.20. If your clients are not 64-bit capable, you can browse files larger than 2 gigabytes, but you cannot recover them.

UNIX Clients

On all Solstice Backup clients for UNIX, you must manually update and verify certain files and paths, as follows:

- The */etc/hosts* file must contain the Internet address of the Solstice Backup client and the Solstice Backup server, unless you use DNS or Network Information System (NIS), for example:

```
127.0.0.1 localhost loopback

123.456.789.111 server server.domain.com

123.456.789.222 client client.domain.com
```

The Solstice Backup software does not automatically configure and update the */etc/hosts* file. You must manually edit the file and verify that the information in this file is accurate. Do not delete or “comment out” the entry for the localhost loopback.

- During installation of the SunOS, Solaris, AIX, and DYNIX/ptx client software, if you accepted the default directory when installing the Solstice Backup executables, the default directory should already be in your executable path. If you specified a different directory, add the directory to your executable path for root or Solstice Backup users.

When you install the HP-UX client software, you must manually add the directory to your executable path, even if you accept the default directory.

For most UNIX clients, the executable path is set in the PATH environment variable. Adding the directory containing the Solstice Backup executables to your executable path allows execution of Solstice Backup commands without entering the full pathname. For example, you would enter `nwbackup` instead of `/opt/nsr/bin/nwbackup`.

Windows NT Clients

On Solstice Backup clients for Windows NT, you must manually update and verify certain files, directories, and services, as follows:

- The `%SystemRoot%\Winnt\System32\drivers\etc\hosts` file must contain the Internet address of the Solstice Backup client and the Solstice Backup server, unless you are using DNS or Windows Internet Naming Service (WINS). The

HOSTS file is a simple ASCII text file with one line for each Internet Protocol (IP) address. The IP address is the first entry on the line followed by the hostname and all aliases for each computer, for example:

```
127.0.0.1 localhost loopback  
  
123.456.789.111 server server.domain.com  
  
123.456.789.222 client client.domain.com
```

Your %SystemRoot%\Winnt\System32\drivers\etc directory should contain a sample `hosts` file that gives details about adding entries to the `hosts` file. Do not delete or “comment out” the entry for `localhost loopback`.

If you are using DNS or WINS, verify that the DNS or WINS server has entries for both the Solstice Backup client and the Solstice Backup server.

- The `servers` file is typically in `C:\Program Files\nsr\res`. The Solstice Backup software uses the contents of this file to control who has the right to request a program to be executed on this client.
 - If you want this client to back up to other Solstice Backup servers, you must add the names of the additional Solstice Backup servers to this file. You can add only one server name per line.
 - If you want other clients to be able to perform directed recovers to this client, you will need to add their names to the `\nsr\res\servers` file. You can add only one client name per line.
 - If you want to allow any Solstice Backup server to back up this Solstice Backup client, delete the `servers` file.

Note: After you save your changes, you must restart the Solstice Backup Exec Service to make your changes take effect.

To allow any Solstice Backup server to back up this Solstice Backup client, delete the `servers` file.

- The Solstice Backup client for Windows NT must have the latest service pack from Microsoft applied.
- Make sure that the following services are running:
 - Solstice Backup Exec Service (`nsrexecd.exe`)
 - Solstice Backup Portmapper Service (also known as `portmap.exe`)

Solstice Backup Portmapper Service is an optional service for Solstice Backup clients. To enable this service, start it before Solstice Backup Exec Service.

- Save sets maintained on an NT DFS (disk file system) link are backed up during a server- or client-initiated backup; however, The Solstice Backup server does not traverse the links or back up the destination files. The server also does not permanently modify the last access time for save sets on a DFS link.

Windows 95/98 Clients

On Windows 95/98 clients, you must manually edit and verify certain files, directories, and services, as follows:

- The `hosts` file, typically found in `C:\Windows`, must contain the Internet address of the Solstice Backup client and the Solstice Backup server, unless you are using DNS or WINS. The `hosts` file is a simple ASCII text file with one line for each IP address. The IP address is the first entry on the line followed by the hostname and all aliases for each computer, as in the following example:

```
127.0.0.1 localhost loopback
```

```
123.456.789.111 server server.domain.com
```

```
123.456.789.222 client client.domain.com
```

- Your Windows 95 directory, typically `C:\Windows`, should contain a sample `hosts` file, named `host.sam`, that gives details about adding entries to an actual `hosts` file. Do not delete or “comment out” the entry for `localhost loopback`.

If you are using DNS or WINS, verify that this DNS or WINS server has entries for both the Solstice Backup client and the Solstice Backup server.

- The `servers` file is in `%SystemRoot\Program Files\Legato\nsr\res`. The Solstice Backup client software uses the contents of this file to control who has the right to request a program to be executed on this client.
 - If you want this client to back up to other Solstice Backup servers, you must add the names of the additional Solstice Backup servers to this file, one server name per line.
 - If you want other clients to be able to perform directed recovers to this client, you will need to add their names to the `\nsr\res\servers` file, one client name per line.
 - If you want to allow any Solstice Backup server to back up this Solstice Backup client, delete the `servers` file.

Note: After you save your changes, you must restart the Solstice Backup Exec Service to make your changes take effect.

To allow any Solstice Backup server to back up this Solstice Backup client, delete the `servers` file.

- The Windows 95 client must have the latest service pack from Microsoft applied.
- Make sure that the Solstice Backup Scheduled Backup (`wtcpschd.exe`) is running. Put a copy of Solstice Backup Scheduled Backup in the Startup folder to enable scheduled backup to run automatically.

NetWare Clients

On NetWare clients, you must manually update and verify certain files, directories, and services, as follows:

- The `SYS:ETC\HOSTS` file must contain the internet address of the Solstice Backup client and the Solstice Backup server. The `HOSTS` file is a simple ASCII text file with one line for each IP address. The IP address is the first entry on the line followed by the hostname and all aliases for each computer. The `HOSTS` file should also contain an entry for `localhost`, as in the following example:

```
127.0.0.1 localhost loopback

123.456.789.111 server server.domain.com

123.456.789.222 client client.domain.com
```

Caution – The TCP/IP hostname and the NetWare server name must be identical for the Solstice Backup for NetWare client. In this example, the NetWare server name replaces the value represented by “client.”

- TCP/IP must be loaded and bound correctly in the `AUTOEXEC.NCF`, for example:

```
load tcpip

load pcentw board=1 frame=ethernet_ii name=e_ii

bind ip to e_ii addr=137.69.8.2 mask=255.255.255.000
```

Load and bind TCP/IP before the Solstice Backup software is installed, or some configuration files are not properly updated.

- Other files that affect Solstice Backup operation on a TCP/IP network and that are automatically configured during the Solstice Backup installation are:

- SYS:ETC\RPCUSERS
- SYS:ETC\SERVICES
- SYS:ETC\RPC, SYS:ETC\GATEWAYS
- SYS:ETC\NET\NETWARE\SERVICES
- RPCNET.CFG (typically found in SYS:NSR)

Other RPC-based products can also use many of these files, so they might already exist on a client before you install the Solstice Backup software. If the files exist, the Solstice Backup software does not overwrite these files during installation. In most cases, files provided by other RPC-based software work with Solstice Backup software.

Multiple Clients from the Same Computer

The Solstice Backup 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 the save sets themselves are unique. For example, looking at a list of configured Solstice Backup clients, you might see several instances of a client named *oak*, but each instance would contain a unique collection of save sets.

Defining multiple clients from the same computer or filesystem can be useful for backing up specialized files such as databases. In this way, save sets from the same computer can be associated with different backup groups or can be backed up according to separate schedules and/or browse or retention policies.

Using Unique Client/Save Set Combinations

Save sets are collections of data to be backed up by the Solstice Backup server. The save sets associated with a specific client instance are visible as a scrollable list in the Save Set attribute of the client's Client resource. 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/save set backups. By redefining a large filesystem into multiple client/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/save set instances for the filesystem. You can then associate each client/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 necessary load balancing needed for your environment.

If different backup start times do not reduce the load adequately, you can associate the different client/save set instances with different backup schedules as well. 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/save set instance runs its full backup on a different day of the week. See [“How the Solstice Backup Server Uses Schedules” on page 193](#) for more information on schedules.

Caution – The `core` file is not backed up unless you specify it in the Save Set attribute of the Client resource.

Client/Save Set Combination

The same save set can appear in the Client resource's Save Set attribute for more than one client instance. 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 back up a specific filesystem, delete the default value *All* from the Save Set attribute and substitute the filesystem pathname. If you want 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 Solstice Backup 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. Select Media>Volumes or click the Volumes toolbar button to open the Volumes window. See [“To View Save Set Information” on page 172](#) for more information.

The save set status changes in the following situations:

- When the save set exceeds its browse policy. For more information about browse policy, see [“How the Browse Policy Works” on page 105](#).
- When the save set exceeds its retention policy, all save sets that are dependent on the save set also exceed their retention policies. For more information about retention policy, see [“How the Retention Policy Works” on page 107](#).
- When you manually change the save set status.

[TABLE 4-1](#) provides a list of all the possible values for save set status.

TABLE 4-1 Save Set Status Values

| 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 other save sets dependent on this save set for recovery have exceeded their retention policies. |

TABLE 4-1 Save Set Status Values (Continued)

| Status Value | Meaning | Description |
|--------------|------------|---|
| scann | scanned-in | The entry for this save set was restored using the <code>scanner</code> program. The policies were reset to the values that were applied when the save set was created (if it was created on a 6.0 server) or when the save sets were upgraded to the new 6.0 format (if the save set was created on a pre-6.0 server). For more information, refer to “Scanner Program” on page 240. |
| susp | suspect | An attempt to recover this save set failed. The <code>recover</code> program could not read all the blocks of the save set, for example, if there was a bad spot in the tape. |

▼ To Change the Status of a Save Set

The Solstice Backup 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 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 Solstice Backup server will always look for the original save set during recoveries. Only when the original is *suspect* or non-recoverable 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 is available. If a cloned volume containing the cloned save set is not available, then the *suspect* volume is used.

1. Select Media>Volumes to open the Volumes window.
2. Select the Volume whose save sets you want to change in the Volumes attribute.
3. Select Save Set>Change Status>Suspect *or* Normal depending on your needs.

Browse and Retention Policies

You use browse and retention policies to specify how long data is available for recovery. You can specify browse and retention policies for each client and save set.

How the Solstice Backup Server Uses Browse and Retention Policies

The Solstice Backup 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 Solstice Backup 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 to accommodate all file or directory renamings or deletions. 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 Solstice Backup server. During the period of the browse policy, users can browse backed-up data in the Solstice Backup recover program (`nwrecover`), and select individual files or entire filesystems for recovery. After the browse policy for a file is exceeded, the Solstice Backup 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 Solstice Backup 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 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 delete

The media database is the structure that tracks the location of save sets on storage volumes. The Solstice Backup server uses the retention policy to manage the longevity of Solstice Backup 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 Solstice Backup 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 Solstice Backup 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 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 193](#), and for more information about backup levels, see [“Backup Levels” on page 202](#).

Example: One Week Browse Policy

In [FIGURE 4-1](#), 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.

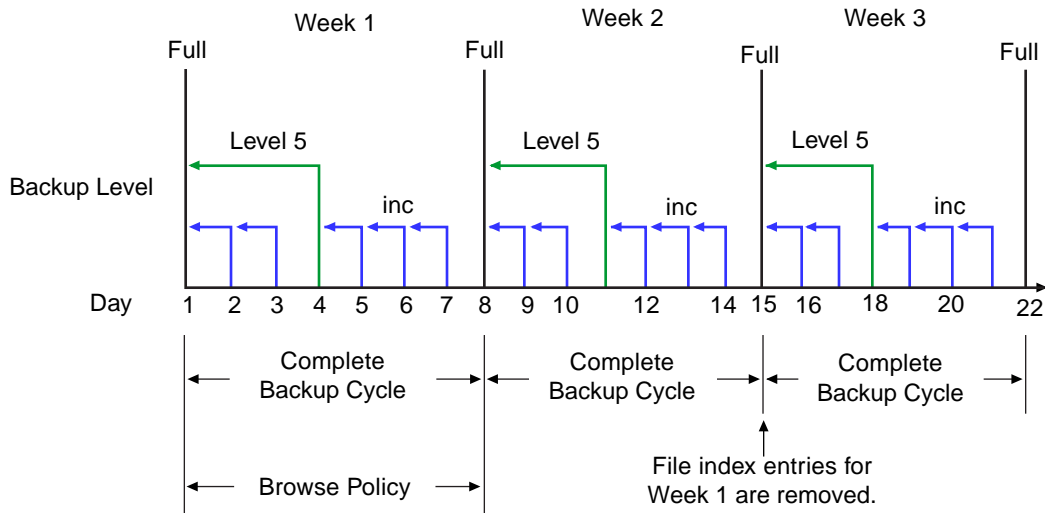


FIGURE 4-1 One-Week Browse Policy

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 the 5th 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](#), 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](#), the browse policy is two weeks, which is twice as long as the backup cycle (1 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.

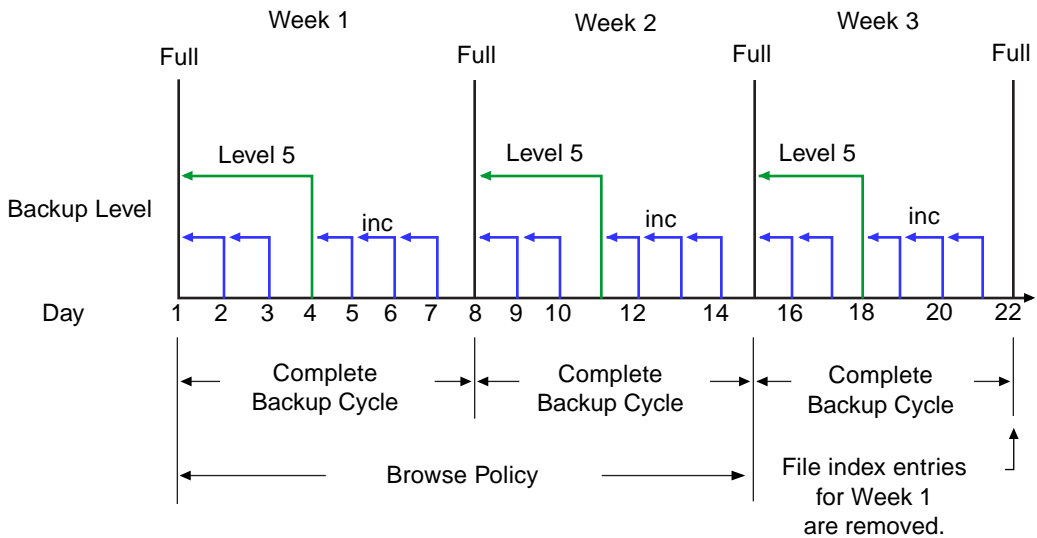


FIGURE 4-2 Two-Week Browse Policy

How the Retention Policy Works

The Solstice Backup media retention policy specifies a period of time 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 Solstice Backup 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 Solstice Backup 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 using either save set `recover` or the `scanner` command. 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:

- 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 Solstice Backup Server Selects a Volume for Relabeling” on page 169](#).

Save set's entries are also removed from the media database when you manually delete a volume from the Solstice Backup volume inventory. However, the data on that volume is still available for recovery 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 through the `nwrecover` program.

If you re-create the save set's entries in the media database, a user with Solstice Backup administration privileges can recover data with `save set recover`. See `recover` in the *Legato Command Reference Guide*, or see [“Scanner Program” on page 240](#) for more information on how to use the `scanner` program.

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.

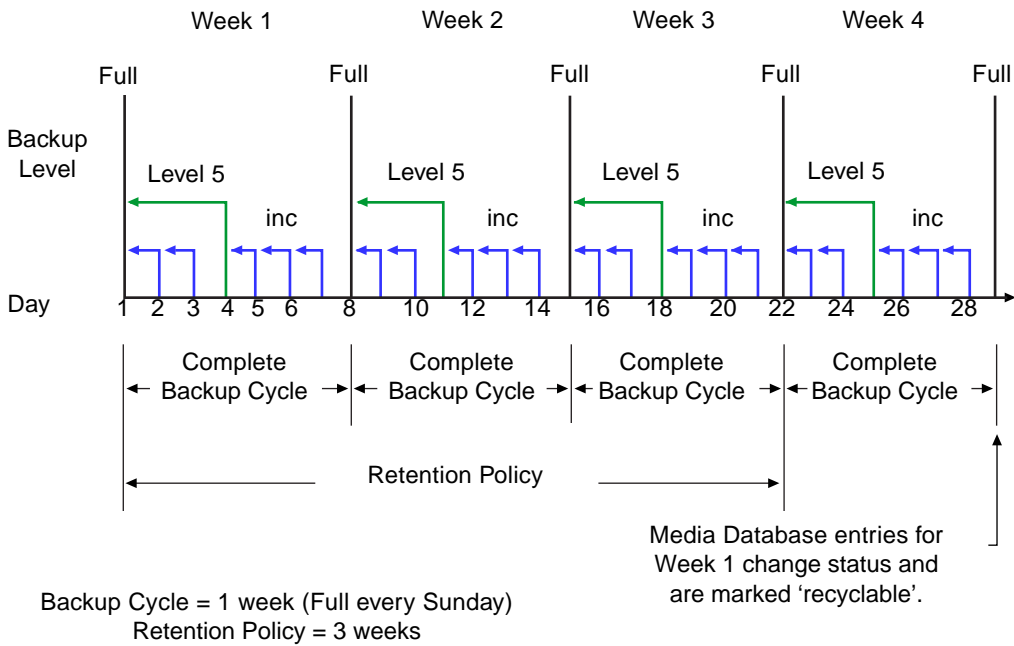


FIGURE 4-3 One-Week Backup Cycle; Three-Week Retention Policy

The save set entries for Week 1 have passed their browse policy and retention policy, but they remain available for recovery using the scanner program until you relabel the volume. When all the save set entries on a volume change status to recyclable, the volume mode changes from full or appendable to recyclable, and the volume is ready to be relabeled for reuse.

Caution – Once you relabel a volume, the data on the volume cannot be recovered.

See for the following section for more information on these topics:

- Storage volume modes, see [“Volume Status” on page 169](#)
- Schedules, see [“Schedules” on page 193](#)
- Backup levels, see [“Backup Levels” on page 202](#)

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 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 229](#) and ["Scanner Program" on page 240](#).

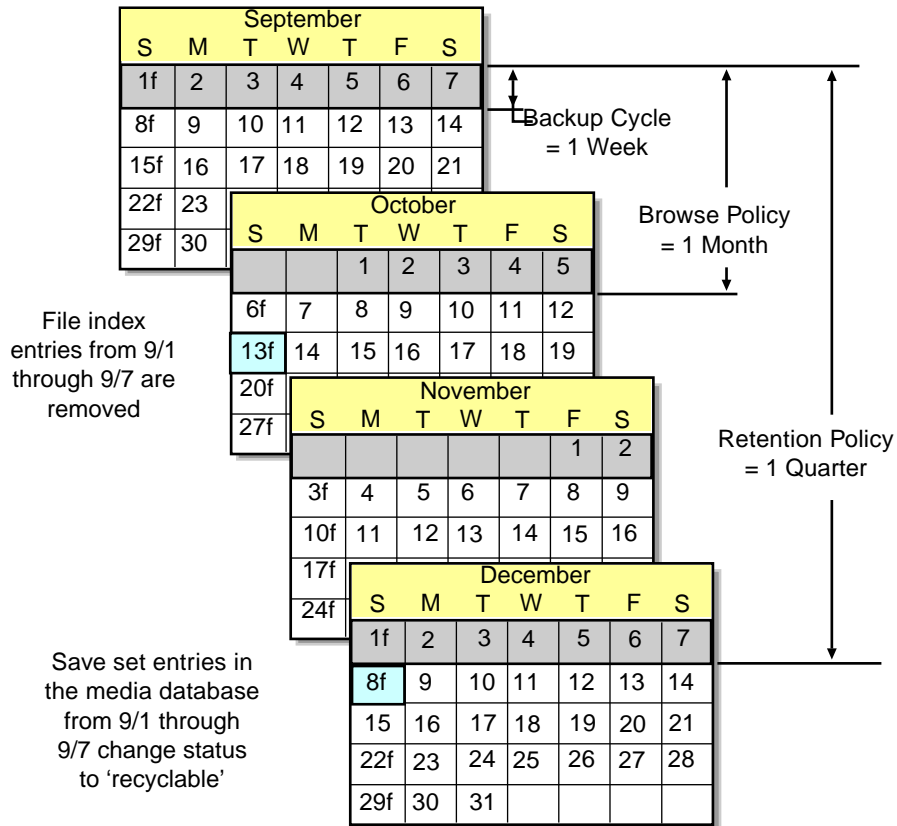


FIGURE 4-4 Data Life Cycle in the Client Index and the Media Database

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 Solstice Backup server to satisfy a volume mount request. After the volume is relabeled, all existing data on the volume is unavailable for recovery. See [“Auto Media Management” on page 162](#) for more information on that feature.

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 so, after relabeling, the information that the Solstice Backup server needs to locate and access all existing data on the volume is destroyed and neither the save set recover feature nor the `scanner` program are options. At this point, the volume is ready for new data. All existing data is inaccessible and is overwritten.

▼ To Assign Multiple Policies to a Single Client

Identical versions of a client/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 Solstice Backup server employs the Browse Policy and Retention Policy attribute values that correspond to the unique combination of the Client resource's Name, Save Set, and Group attribute values.

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 *Biweekly* and the retention policy as *Yearly*. If the group *Special* is backed up, then the *Biweekly* and *Yearly* policies are applied. If the group *General* is backed up, then the policies *Weekly* and *Monthly* are used.

Preconfigured Policies

[TABLE 4-2](#) lists preconfigured settings for the Policy resource attributes.

TABLE 4-2 Preconfigured Solstice Backup Policies

| Policy Name | Solstice Backup Availability |
|-------------|------------------------------|
| Decade | Available 10 years |
| Month | Available one month |
| Quarter | Available for three months |

TABLE 4-2 Preconfigured Solstice Backup Policies (*Continued*)

| Policy Name | Solstice Backup Availability |
|-------------|------------------------------|
| Week | Available one week |
| Year | Available one year |

▼ To Create Policies

To create a policy:

1. Select **Customize>Policies** to open the Policies window.
2. Click **Create**.
3. Enter the name of the policy in the **Name** attribute. 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, however, must be unique for the Solstice Backup server.
4. Select a period in the **Period** attribute.

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). Solstice Backup software defines a week as seven days beginning on Sunday, a month as the calendar month, and a year as the calendar year.

5. Enter a number in the **Number of Periods** attribute. This value indicates the number of periods applied to the policy. For example, if you choose *months* for the **Period** attribute and *three* for the **Number of Periods** attribute, then the policy lasts for three months (or a quarter).
6. Complete any other attributes as necessary, and click **Apply**.

▼ To Edit Policies

You cannot change the name of a policy.

The Solstice Backup 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:

1. Select Customize>Policies to open the Policies window.
2. Select the policy you want to edit from the list displayed in the Policies attribute.
3. Edit the policy, and click Apply.

▼ To Delete Policies

You cannot delete any policy currently assigned to a client.

To delete a policy:

1. Select Customize>Policies to open the Policies window.
2. Select the policy you want to delete from the list displayed in the Policies attribute.
3. Click Delete.

You will be prompted to confirm that you really want to delete that policy. Click OK to confirm.

4. Click Apply.

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:

- 1 week
- 1 month
- 1 year

A manual backup of *mars* with no additional modifications will automatically adopt a retention policy of 1 year and the corresponding browse policy assigned to the client instance which has the 1 year retention policy.

You can modify the browse and retention policies of manual backups, but only when performing the backup using the `save` command at the command prompt. See [“To Modify Manual Backup Policies” on page 116](#) for more information.

See [“Manual Data Backups” on page 213](#) for more information on manual backups in general.

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. You can also use the `save` program at the command line 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.

`nsrmm` updates browse and retention policies using these options:

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

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

Please note that browse time cannot be after retention time, the retention time cannot be before the browse time, and the browse time has to be after the *insertion time*. This 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 will fail if any of the following is true:

- Browse time > insertion time, if `-w` option was specified
- Retention time > insertion time, if `-e` option was specified
- Retention time \geq browse time

Both the browse time (specified with option `-w`) and the retention time (specified with the option `-e`) must be stated in time and date formats acceptable to `nsr_getdate`, which will convert them to formats acceptable to the Solstice Backup server. 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* or the man pages for those respective commands.

Note: With the exception of save time, save sets and volumes are no longer dependent on their corresponding Client resource for their browse and retention policies. This allows more control over whether the save set's file indexes or volumes need to be maintained, deleted, or recycled.

Example: Changing Browse and Retention Policies

These are examples of how a user might change browse and retention policies:

1. Change the retention time to midnight, January 1, 2005; change the browse time to midnight, January 1, 2001.

```
# nsrmm -s <3315861249> -e <24:00:00, 01/01/05> -w <24:00:00, 01/01/01>
```

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` command at the command prompt. You cannot modify manual backups using the graphical `nwbackup` program.

Use the `save` command with the following options in the same manner as modifying scheduled backup policies using `nsrmm`.

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

For example, if you wanted to perform a manual backup of `file_name` and give it a retention time of one month, you would enter the following command at the prompt:

```
# save -e "1 month" file_name
```

The `nsr_getdate` man page contains the terms and syntax to use when specifying browse and retention times. Refer to the *Legato Command Reference Guide* and the `nsr_getdate` and `save` man pages for more information.

Reports on Browse and Retention Policies for Save Sets

The `-p` option causes `mminfo` 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
- Save set ID
- Stored browse and retention dates (the term *undef* is displayed when connecting to a server using a pre-6.x version of Solstice Backup software)
- Client name
- Save set name

The `-v` and `-V` options have no effect on the columns included in this report.

Granularity

Prior to this release of Solstice Backup 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 Solstice Backup 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 `nsrim` 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.

▼ To Override Save Set-Based Browse and Retention Policies

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. Dependency rules still apply. However, this means that a save set is not marked *recyclable* until all save sets that depend on it are also marked as recyclable.

Backup Command Customization

You can customize your client backups by creating additional programs that affect the way your Solstice Backup server backs up client filesystem data.

For example, you can create a program that:

1. Shuts down either a mail server or a database before the Solstice Backup server performs a backup.
2. Restarts the mail server or database after the backup is completed.

In another example, you can create a program that:

- Configure the client's backup to use the `savenpc` program instead of the standard `save` program. To configure a client to run `savenpc`, enter `savenpc` in the Backup Command attribute in the Client resource for the client computer. The first time the client is backed up, `savenpc` will create 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. For further information, see [“Using `savenpc` With a Customized Backup Program” on page 120](#). Consider the following issues as you determine what level of customization works best for your environment:
 - Amount of disk space you have
 - Whether you have client data that does not need to be backed-up every time (for example, company e-mail)
 - Whether you want special messages sent (in addition to the `savegroup` completion reports) about the backups executed

- Configure the client's backup to use the `savepnpc` program instead of the standard `save` program. To configure a client to run `savepnpc`, enter `savepnpc` in the Backup Command attribute in the Client resource for the client computer. The first time the client is backed up, `savepnpc` will create a default backup program file, which you can then customize for future backups of the client. When a client is backed up using `savepnpc`, the instructions in the backup program are invoked just once during the backup session. For further information, see [“Using `savepnpc` With a Customized Backup Program” on page 120](#).

Using the `save` Command in a Customized Backup Script

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.

The instructions you include in your script are performed on a per-save set basis, rather than on a per-client basis. If you specify a save set value of *All*, your script is executed the same number of times as the number of filesystems on the client. The simplest implementation of a customized backup command is to create a special, separate client with a single save set listed in the Save Set attribute.

Caution – Unlike `savepnpc`, a new instance of the customized script, whose name you enter in the Backup Command attribute, is invoked for each save set listed in the Save Set attribute, similarly to `save` program. Remember this when you create a Client resource with a customized Backup Command for a database. A command to shut down the database is executed for each save set that you listed.

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 backup program name must begin with either the prefix *save* or *nsr* and must not exceed 64 characters.
- The backup program must reside in the same directory as the Solstice Backup `save` command.
- The Solstice Backup `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 Solstice Backup server cannot complete the remaining instructions.

- When you invoke the Solstice Backup `save` command, invoke the command with the following arguments: `save "$@"`. Doing so enables the `save` command in your batch file to accept the arguments usually passed to it by the Solstice Backup `savefs` program during a routine backup operation.

The program commands are to be placed in the following order:

- Run a pre-processing command before a client backup (optional).
 - Back up the data using the Solstice Backup `save` command (mandatory).
 - Run a post-processing command after a client backup (optional).
1. Use a text editor to create a program file in the directory where the Solstice Backup `save` command resides.
 2. Enter the name of the backup program in the Backup Command attribute of the Client resource.
 3. Back up the client to ensure that the backup command you created works.

Using `savenpc` With a Customized Backup Program

By using the `savenpc` program instead of the `save` program, you can run pre-processing and post-processing 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 is completed. The options for the `savenpc` command are identical to those for the `save` command. For further details about `savenpc` command options, see the `savenpc` command in [the Legato Command Reference Guide](#).

To execute `savenpc`, you must first enter “`savenpc`” in the Backup command text box in the Client dialog box, and run a backup of the client.

The first time a backup group with a client that uses `savenpc` runs, a standardized `<group-name>.res` file is created in the `/nsr/res` directory or `<nsrroot>\nsr\res` 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, pre-processing, post-processing, and timeout attributes:

```
type: savepnpc;
precmd: "echo hello";
pstcmd: "echo bye", "/bin/sleep 5";
timeout: "12:00pm";
```

Once the <group-name>.res file exists, you can use your favorite text editor to customize the file's attributes. Your customized instructions will be executed the next time the client is backed up.

Before performing a save operation on the client, the modified savepnpc program performs any pre-processing commands listed for the precmd attribute in the <group-name>.res file. Then it executes the save using the options you specified for the savepnpc command itself. After the last save operation is successfully completed, savepnpc performs any post-processing 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 your environment as part of your pre-processing (precmd) commands, especially the PATH variable. On UNIX clients, be sure to source the .profile, .cshrc, etc. login scripts.
- If you don't want to include your environment variables in the <group-name>.res file, then you must include full pathnames for all commands and files.
- You must escape any "\" characters in the <groupname>.res file. Therefore, a pathname like c:\mydir\myprogram.exe should be written c:\\mydir\\myprogram.exe. The following is an example of an actual <group-name>.res file:

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

It is not necessary to escape any "\" characters in scripts called in the <group-

name>.res. A suggested way to simplify the pathname issue is to put all of your commands in a script or batch file, and then put that script's full pathname on the precmd or postcmd line.

- Text written to standard output by the preprocessing commands will appear in the NetWorker completion notices. If this output is verbose, you might prefer to direct the output from the commands in your script to a log file for later examination. Output from the postprocessing commands is thrown away, 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.
- To add more than one command sequence to the precmd and postcmd attributes, use a comma (,) to separate the commands.
- A complete command line for an attribute must end with a semicolon (;).
- Do not use savepnp with groups that have a space in the group name.
- On a Windows-based client, do not use "@ECHO OFF" in the <group-name>.res file.
- Resident commands, for which there is no .exe file present, like time and dir, will not work as commands in your <group-name>.res file. The log will report that the executable file could not be found.
- If you believe your setup is correct, or the backup with savepnp was working but no longer does, or if a backup for a client is aborted, or the precmd fails, look for a /nsr/tmp/<group-name>.tmp or <nsrroot>\nsr\tmp\<group-name>.tmp file. Delete it if found. The .tmp file is a lock file that is created when savepnp runs, and is normally removed when the savesets for a system have been completed.

If a new savegroup session is started and there is an existing <group-name>.tmp file, savepnp will assume that the pre- and post-processing commands have already run and will skip them; however, it will run the filesystem save sessions.

- The timeout attribute (timeout) indicates a point in time at which your post-processing commands (postcmd) are to be run, regardless of whether all of your save sets have been backed up or not. The timeout condition must be specified in a format that nsr_getdate can understand (for details, see the man page for nsr_getdate, or read the description of nsr_getdate in [the Legato Command Reference Guide](#)).

Note: If an invalid time is entered for the timeout, the timeout is not executed, and there is no error message produced.

Note: The timeout attribute is optional. If you do not plan to use it, you can comment it out with a # sign, and it will be ignored, as shown here:

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


Note: The value of the timeout attribute may not be the exact time that post-processing actually commences. The `savepnp` program's `pstclntsave` subroutine uses a 1-minute polling interval to check for the completion of pre-processing tasks. Therefore, the `savepnp.log` file (see the sample log below) may show that post-processing was commenced up to 60 seconds *after* the designated timeout.

To customize pre/post command processing for various groups, you can copy any existing `/nsr/res/<group-name>.res` file to `/nsr/res/<your_new_group>.res`, and then edit this new file. If you do not have an existing `<group-name>.res` file you can activate your group for `savepnp` without the presence of this file and a default template will be created for you at `/nsr/res/<your_new_group>.res`, which you can edit.

On a UNIX client, any messages generated by `savepnp` are written to the `/nsr/logs/savepnp.log` file on the client.

The format of a UNIX client's `savepnp.log` file is similar to the following:

```
09/03/99 13:56:43 preclntsave: On mars:
```

```
Ran pre-processing cmd(s) successfully.
```

```
09/03/99 13:56:43 preclntsave: On mars:
```

```
Successfully spawned off pstclntsave.
```

```
09/03/99 13:57:43 pstclntsave: On Mars:
```

```
All savesets on the worklist are done.
```

```
09/03/99 13:57:51 pstclntsave: On mars:
```

```
Ran all post-processing cmd(s) successfully.
```

On a Windows client, any messages generated by `savepnp` are written to the `\<nsrroot>\logs\savepnp.log` file on the client.

The format of a Windows client's `savepnp.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.
```

The corresponding entries written to the `syslog` are similar to the following example:

```
Sep 3 13:56:43 jupiter syslog: Solstice Backup Savegroup:  
(info) starting default_savepnp (with 1 client(s))
```

Sep 3 13:57:51 jupiter syslog: Solstice Backup Savegroup: (notice) default_savepnpc completed, 1 client(s) (All Succeeded)

Sep 3 13:57:51 jupiter syslog: Start time: Tue Sep 3 03:33:01 1999

Sep 3 13:57:51 jupiter syslog: End time: Tue Sep 3 03:33:12 1999

Sep 3 13:57:51 jupiter syslog:

Sep 3 13:57:51 jupiter syslog: --- Successful Save Sets ---

Sep 3 13:57:51 jupiter syslog:

Sep 3 13:57:51 jupiter syslog: * mars:/testfile hello

Sep 3 13:57:51 jupiter syslog: mars: /testfile level=incr, 0 KB 00:00:01
0 files

Sep 3 13:57:51 jupiter syslog: mars: /space/nsr/index/mars level=9, 0 KB
00:00:01 0 files

Sep 3 13:57:51 jupiter syslog: jupiter: bootstrap level=9, 19 KB
00:00:02 6 files

Sep 3 13:57:51 jupiter syslog:

Sep 3 13:57:53 jupiter syslog: Solstice Backup index: (notice) nsrim has finished
checking the media db

▼ To Back Up a ClearCase® VOB

Technical Bulletin 306, which is found in the `bulletins.pdf` file included with your Solstice Backup software, provides a sample script that you can use as a reference to customize in order to back up a ClearCase VOB (version object base). Your customized script file must reside in the same directory as the Solstice Backup `save` program (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 configured for the ClearCase VOB. As a result, the Backup Command is invoked instead of the usual `save` program 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 is an example of a custom script that performs pre- and post-processing. This script locks a ClearCase VOB, performs the backup, and then unlocks the VOB.

```
#!/bin/sh

# export the SHELL that we are going to use

SHELL=/bin/sh

export SHELL

# export the correct PATH so that all the required binaries can be
found

case $0 in

/* ) PATH=/usr/atria/bin:/bin:/usr/bin:`/bin/dirname $0`

c=`/bin/basename $0`

;;

* )PATH=/usr/atria/bin:/bin:/usr/bin:/usr/sbin

c=$0

;;

esac

export PATH

# These are the valid statuses 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
```

Directives

Directives are special programs the Solstice Backup product applies to client/save set data to initiate additional data processing. A directive contains instructions to assist the backup process, maximize the efficiency of a backup, and handle special files.

Suggested uses of directives include:

- Reducing the amount of data you back up by skipping files and compressing data, possibly even eliminating the need to change volumes on the days you perform a full backup.
- Instructing the Solstice Backup server to always include a certain file for backup, regardless of the file's backup schedule.

Directives appear as selectable options associated with the Directives attribute in the Client resource. Backup directives provide a flexible method for backing up files.

Note: If you would like more information about directives, including algorithms and the correct syntax used to apply them, refer to `nsr` in the *Legato Command Reference Guide*.

▼ To Write Directives

You can write directives in the Directive resource and apply them to a specific client through the Client resource. Because every environment is different, it is impossible to prescribe directive-writing rules that work in every case. If you are having questions or difficulty writing a customized directive, contact Technical Support for further assistance.

Directives use Application Specific Modules (ASMs) to back up files. Directives tell the Solstice Backup server which ASM to apply to a particular file or groups of files. An ASM might perform something as simple as passing backup data to the server

unchanged. A directive might contain a more complex set of instructions that encrypt or filter data in a large directory (for example, a database) and save only the data that changed.

TABLE 4-3 describes the Solstice Backup ASMs.

TABLE 4-3 Solstice Backup ASMs

| ASM Name | Description |
|-------------|--|
| Always | Always backs up the specified files and ignores the scheduled level. This ASM is helpful if you have important client files that constantly change, and you want to assure that a full backup is always performed on the specified files. |
| Compressasm | Compresses the client files so that they use less network bandwidth and fewer volumes. Using the compressasm directive can reduce the space consumed on a volume by as much as 50 percent. If you use compressasm on all the files that comprise 8 GBs of data, the backup may fit on a single volume. If your storage device compresses data, you can also apply the compressasm directive; however, you gain only an insignificant amount of additional compression. |
| Nullasm | 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 Solstice Backup server uses nullasm to back up the online indexes during a scheduled backup. |
| Pw1 | Password-protects the named data items. <i>Note:</i> Available on Windows and NetWare versions of Solstice Backup software only. |
| Pw2 | Password-protects and encrypts the named data items. <i>Note:</i> Available on Windows and NetWare versions of Solstice Backup software only. |
| Skip | Omits certain files from the backup entirely. You can use pattern matching to apply the instructions to specific files in the directory tree. For example, you can instruct that the Solstice Backup server skip backing up <i>.obj</i> files. |

The correct syntax for directives is as follows:

1. Enter the disk volume or directory name to which you want the set of directives to apply. Surround the name with double-angle brackets. If the name contains a colon (:), surround it with quote marks. The two examples below are acceptable:

```
<< / >>

<< /tmp >>
```

while this example is not:

```
<< /tmp/ >>
```

2. Enter the name of the directive under the directory name, and include the pattern-matching instructions, if any. For example:

```
<< / >>  
  
skip: tmp_mnt
```

This instruction means “skip all files named tmp_mnt in the root directory.”

If a directive is preceded by a plus (+) sign, it applies to the named directory and all of its subdirectories. For example:

```
<< / >>  
  
+skip: core
```

To include all files, the following syntax examples are acceptable:

```
.  
  
*  
  
*.*
```

The following directive performs two operations.

- *nullasm* – skips backing up the /tmp directory but places entries for the /tmp directory files in the online index.
- *compressasm* – compresses the data in the /nsr directory.

```
<< / >>  
  
nullasm: TMP  
  
compressasm: NSR
```

Caution – In all instances where a wildcard character appears, there must be a space inserted between the wildcard and the next character, to avoid unexpected results.

For example, if you implemented the previous directive example that used the wildcard `*.*` to indicate that the directive should be implemented upon all files and you accidentally omitted the space between the asterisk (*) and dot (.) characters, the actual result would be that only files with a dot (.) character in the middle of their filename would be subject to the directive. In this instance, the intent was to have the directive act on all files: the asterisk (*) indicates all files, while the dot (.) indicates that the directive should also act upon all hidden files.

Preconfigured Directives

Besides writing customized directives, you can use the preconfigured directives that come with the Solstice Backup software. You can use the UNIX standard directive as a template that you can customize for your UNIX clients. You can modify or delete most of the preconfigured directives, with the exception of the UNIX standard directive, which can be modified but not deleted.

The Windows NT standard, DOS standard, and NetWare standard directives are shipped without specific directives coded, so you can use the UNIX standard directive as a template and customize this to your specific environment.

[TABLE 4-4](#) lists the preconfigured directives and their descriptions:

TABLE 4-4 Preconfigured Directives

| Directive | Description |
|-----------------------------------|---|
| DOS standard | Use to back up your Windows 98 and Windows 95 clients. |
| DOS standard with compression | Use to back up and compress your Windows 98 and Windows 95 clients. |
| NetWare standard | Use to back up your NetWare clients. |
| NetWare standard with compression | Use to back up and compress your NetWare clients. |
| NT standard | Use to back up your Windows NT and Windows 2000 clients. |
| NT standard with compression | Use to back up and compress your Windows NT and Windows 2000 clients. |

TABLE 4-4 Preconfigured Directives (*Continued*)

| Directive | Description |
|--------------------------------|--|
| UNIX standard | <p>Used for most of your UNIX clients, and when you do not need one of the other specialized directives. This selection:</p> <ul style="list-style-type: none"> • Applies the directive “+skip: core” to the <i>root</i> directory (/), thus skipping the backup of all <i>core</i> files. • Applies the <i>swpasm</i> directive to the <i>/export/swap</i> directory to back up the relevant information about all NFS-based and local swap files, but not the data in them. If your swap files are located in a different directory, modify this line to include the appropriate location. • Applies the <i>mailasm</i> directive to the <i>/usr/spool/mail</i> and <i>/usr/mail</i> directories to ensure that your mail files are backed up but not marked as read. If your mail files are located in different directories, modify these lines to include the appropriate locations. • Applies the <i>logasm</i> directive to the <i>/nsr/logs</i>, <i>/var</i>, <i>/usr/adm</i>, and <i>/usr/spool</i> directories. If you have log files located in different directories, modify these lines to include the appropriate locations. |
| UNIX standard with compression | Use to back up and compress your UNIX clients. |

▼ To Create Directives

To create a directive:

1. Select Customize>Directives to open the Directives window.
2. Click Create.
3. Enter a name for your new directive in the Name attribute.
4. Enter the directive syntax in the Directive attribute. Your directive syntax can include UNIX commands or Solstice Backup commands. See [“To Write Directives” on page 127](#) for instructions.
5. Complete any other attributes as necessary and click Apply

The directives you define in the Directive resource appear in the Client resource as selectable options for the Directive attribute. If you want to create your own customized directive, define it before you configure the client so that the directive name appears as an option.

Note: If you would like more information about directives, including algorithms and the correct syntax used to apply them, refer to *nsr* in the *Legato Command Reference Guide*.

▼ To Edit Directives

To edit a directive:

1. Select `Customize>Directives` to open the Directives window.
2. Select the directive you want to edit from the list displayed in the Directives attribute. The instructions of that notification is displayed in the Directive attribute.
3. Modify the instruction of the directive and click `Apply`.

You cannot change the name of a directive.

▼ To Delete Directives

To delete directives:

1. Select `Customize>Directives` to open the Directives window.
2. Select the directive you want to delete from the list displayed in the Directives attribute.
3. Click `Delete`.

You will be prompted to confirm that you really want to delete that directive. Click `OK` to confirm. The directive is removed from the list of directives displayed.

4. Click `Apply`.

You cannot delete preconfigured directives or any directives applied to a client.

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 can contain a value between 1 and 1000. 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 Solstice Backup server. If you do not specify a value 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.

Solstice Backup Client Licensing

When a backup begins, a Solstice Backup client passes its attributes to the Solstice Backup server. The Solstice Backup 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
- Solstice Backup version
- CPUs
- Enabler in use

When you enable a Solstice Backup server, the enabler also allows you to use a specific number of clients of the same platform type. For example, a Solstice Backup 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 Solstice Backup 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 Solstice Backup server enforces the number and type of licensed clients at backup time, 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.

At one time, there were two types of client connection licenses: *server* and *workstation*. Sun no longer makes this distinction. Now, only *universal* licenses are sold. These licenses are valid for any computer, regardless of the computer's operating system.

The `nsrlic` Program

Solstice Backup software uses the command-line program `nsrlic` to generate reports about all licenses currently active on a Solstice Backup server. The `nsrlic` program shows a breakdown of the types of workstations and servers that are in use. In addition, `nsrlic` has a `-v` option, which will provide a list of all connected clients, as well as additional license information. For more information about the `nsrlic` command, see `nsrlic` in 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 Solstice Backup server (for example, Solaris, AIX or Windows NT) are listed under universal.
- All others (for example, Windows 95/98) are listed as workstations. Workstation licenses will be “borrowed” from the universal licenses, unless the server has old workstation client connection licenses.

Device and Media Management

This chapter describes device and media operations you can perform through the Solstice Backup server.

Pools

A pool is a specific collection of media to which the Solstice Backup server writes data. The Solstice Backup 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 148](#).

How the Solstice Backup Server Uses Pools

The way you configure pools determines which volumes receive data. Each pool configuration contains a list of criteria that the data must meet for the data to be written to associated volumes.

When a scheduled backup occurs, the Solstice Backup 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 to see whether a correctly labeled volume is mounted on a storage device.

- If it is, the Solstice Backup server writes data to the volume.
- If it is not, the Solstice Backup server requests that such a volume be mounted and waits until an operator mounts the appropriate volume.

Example: Using Pools Configurations to Sort Data

FIGURE 5-1 illustrates how the Solstice Backup server uses pool configurations to sort data. The savestream contains data from a full backup that was performed on client/save sets in a group called *Accounting*. The Solstice Backup 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 mounted on one of the storage devices.

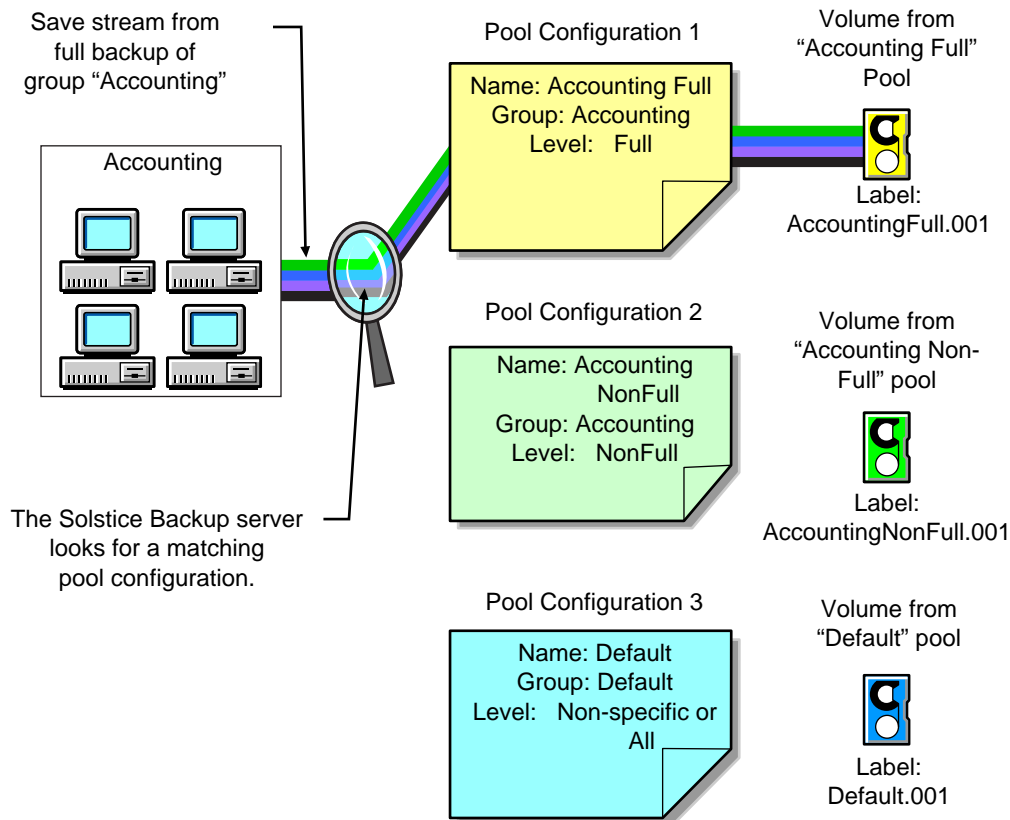


FIGURE 5-1 Using Pool Configurations to Sort Data

Solstice Backup Pool Types

The Solstice Backup software provides preconfigured pool types to keep different types of data separate. Solstice Backup 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 Archive pool. Cloned backup data is routed to the Default Clone pool, and cloned archive data is routed to the Archive Clone pool.

Preconfigured Pools

[TABLE 5-1](#) lists the preconfigured pools that come with the Solstice Backup product.

TABLE 5-1 Preconfigured Pools

| Pool Name | Solstice Backup Backup Operation |
|------------------|---|
| Default | Pool receiving all backed up data when another pool is not specified. |
| Default Clone | Pool for receiving data from clones of the Default pool's volume |
| Archive | Pool receiving all archived data when another pool is not specified. |
| Archive Clone | Pool for receiving data from clones of the Archive pool's volume. |
| Full | Pool for receiving data from Full backups. |
| NonFull | Pool for receiving data from incremental and other non-Full backups |
| Offsite | Pool for receiving data to be stored offsite |
| PC Archive | Pool for receiving data from PC Archive clients. |
| PC Archive Clone | Pool for receiving data from cloned PC Archive clients |

Default Pool

If the savestream does not match the criteria for any custom (user-created) pools, the Solstice Backup server directs the savestream to the *Default* pool. Even if you are using your own customized pools or other preconfigured pools, always ensure that appropriate storage devices and media are available for your Default pool for those cases when the Solstice Backup server directs savestreams there. If you do not have your 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
```

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 Windows systems) and periods must be escaped with the backslash character, so that a save set path of:

C:\SomeDir\MyFiles

would have to be written:

C:\\SomeDir\\MyFiles.

A filename of *MyFile.txt* would be written *MyFile\\.txt*.

Separate save set attributes with a comma and space. The following is an example of properly written save set entries in the save sets attribute:

```
/, /usr, c:\\Program Files\\legato, *\\.jpg
```

Refer to the *Legato Command Reference Guide* and the `nsr_regexp` and `nsr_pool` man pages for more information on these commands and expression matching.

How the Solstice Backup Server Uses Pool Criteria to Sort Data

When you configure the Solstice Backup server, you can create additional pools and sort data by pool type and any combination of the following criteria:

- Group (backup group)
- Solstice Backup 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 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 different pool than you send the backup data.

In the following example, the client file indexes are in `index:*`. To send the Solstice Backup server's bootstrap and all the client file indexes from this filesystem to the same pool, create a pool (in the Pools resource) with the following attributes:

```
name: Index;

pool type: Backup;

save sets: bootstrap, index:*;

levels: ;
```

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 Solstice Backup 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 Solstice Backup server has to determine to which pool 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 5-2 details the hierarchy that the Solstice Backup 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 5-2 Solstice Backup Hierarchy for Resolving Pool Conflicts

| Precedence | Group | Client | Save Set | Level |
|------------|-------|--------|----------|-------|
| Highest | x | x | x | x |
| | x | x | x | |
| | x | x | | x |
| | x | x | | |
| | x | | x | x |
| 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 your 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 Solstice Backup 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 Solstice Backup 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 Solstice Backup 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 Solstice Backup 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 Solstice Backup request for a Default pool volume. For the procedure to test your scheduled backup, see [“To Start Groups Immediately” on page 189](#).

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

Incremental Backups

If you want to create a separate pool for incremental backups, be aware that the Solstice Backup hierarchy of precedence affects the way the data is stored. 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 Solstice Backup 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 Solstice Backup 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 your data, you might have to use a large number of volumes to recover all your data. To store the client file indexes along with the incremental backup data and to speed the recovery operation, define the Level value in the Pools resource to accept both level 9 and incremental data.

You can use the Solstice Backup 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 Solstice Backup 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 the pool in which data from this client's save set is stored during a regularly scheduled backup operation.

If you do not create a customized pool to receive data from manual backups, the Solstice Backup 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 Solstice Backup 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. Refer to the *Disaster Recovery Guide* for further information on how the bootstrap is used during disaster recovery.

Clone Data

If you want to clone data, the Solstice Backup 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 Solstice Backup 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. See [Chapter 9](#) for more information on the Solstice Backup cloning feature.

Archive Data

If you want to use Networker Archive to archive data, the Solstice Backup software requires a specific pool to receive the archived data. You can then store these volumes off site, if you want. If you do not associate data to be archived with a customized archive pool, the Solstice Backup server automatically uses the preconfigured Archive pool. You must mount an appropriately labeled volume on a storage device for the archive process to proceed smoothly. See [Chapter 8](#) for more information on the Solstice Backup archive feature.

Migration Data

the Solstice Backup server the Solstice Backup server Archive and migration data are both written in a different format than regular Solstice Backup save set data. Therefore, archive and migration data must each be written to different volumes. Because of these differences, the client file indexes and bootstrap save set created during a PC archive, premigration, or migration operation are also not written to the same volume as the archived or migrated save sets. By default, they are written to a volume from the Default pool. If you need to direct the client file indexes and bootstrap to a volume pool other than Default, see [“Directing Client Indexes and Bootstrap to a Separate Pool” on page 139](#) for information.

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. See [“Client/Save Set Combination” on page 101](#) for more information about how a single save set can be written to more than one type of storage media.

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 off-site 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 Devices 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 Pools

To create a pool:

1. Determine a label template for your pool.

You can use one of the preconfigured label template, or create a new one. See [“To Create Archive Pools” on page 147](#) for a complete discussion of label templates.

If you do not select a label template, the Solstice Backup server notifies you that it will create a label template for your pool. Click OK to have the server create the label template.

2. Select Media>Pools to open the Pools window.
3. Click Create.
4. Enter a name for the pool in the Name attribute. This name must match either a label template you have created beforehand, or one of the default label templates.
5. Select Yes from the Enabled attribute.
6. Select the appropriate pool type from the Pool Type attribute.
7. Select the matching label template in the Label Template attribute.
8. Select the groups associated with this pool in the Groups attribute.
9. Enter the names of clients associated with this pool in the Clients attribute.
10. Enter the save sets associated with this pool in the Save Sets attribute.

See [“Expression Matching of Save Sets to Pools” on page 138](#) for information on the syntax of save set paths.
11. Select the backup levels associated with this pool in this Levels attribute.
12. Select the devices that are permitted to mount volumes associated with this pool.

13. Select whether or not client file index entries are generated for this pool in the Store Index Entries attribute.
14. Select whether or not automated media verification will be performed while data is written to a volume associated with this pool in the Auto Media Verify attribute. Refer to the section below on this attribute for more information.
15. Select whether or not volumes marked as *recyclable* can be used by other pools in the Recycle to Other Pools attribute.
16. Select whether or not this pool can recycle volumes from other pools in the Recycle from Other Pools attribute.
17. 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 writeable volume.
18. Complete any other attributes as necessary, and click Apply.

If any of these settings for a new pool match an existing pool, you receive the following warning message:

Pool *pool_name* has overlapping selection criteria.

Auto Media Verification

If the Auto Media Verify attribute in the Pools resource is enabled, the Solstice Backup 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. The data read is compared to the original data written. 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 Solstice Backup 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 Solstice Backup server is attempting to switch volumes, all save sets writing to the volume are terminated.

▼ To Edit Pools

To edit a pool:

1. Select Media>Pools to open the Pools window.
2. Select the pool you want to edit from the list displayed in the Pools attribute.
3. Edit the pool's attributes, and click Apply.

You cannot change the name of a pool.

▼ To Delete Pools

To delete a pool:

1. Select Media>Pools to open the Pools window.
2. Select the pool you want to delete from the list displayed in the Pools attribute.
3. Click Delete.

You will be prompted to confirm that you really want to delete that pool. Click OK to confirm. The notification is removed from the list of pools displayed.

4. Click Apply.

▼ To Create Archive Pools

If you want to archive data, the Solstice Backup server requires a specific pool to receive the archive data. You can then store these volumes offsite. If you do not associate data to be archived with a customized archive pool, the Solstice Backup server automatically uses the appropriate preconfigured archive pool. An appropriately labeled volume must be mounted on a storage device for the archive process to proceed.

To create an archive pool, follow these steps:

1. Create a new pool. See ["To Create Pools" on page 145](#) for instructions.
2. Set the Pool Type to Archive or PC Archive (depending on the type of client), and Store Index Entries to No. These two settings distinguish archive pools from backup pools.

Now 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 Solstice Backup server will write archive data only to an archive volume, and archive clone data only to an archive clone volume. See [“Pools” on page 135](#) for more information on pools.

Volume Labeling

The Solstice Backup software labels (initializes) each storage volume with a unique internal label that corresponds to a pool. During backup and other operations, Solstice Backup software can identify the pool to which a volume belongs by its label. Solstice Backup software applies a label template to create a unique internal label for each volume.

The Solstice Backup software uses label templates and pool configuration settings to sort, store, and track data on media volumes. If you need to recover data, Solstice Backup software prompts you for the specific volume that contains the required data, by volume name and sequence number.

How the Solstice Backup Server Uses Label Templates

The Solstice Backup server writes a given set of data to a specific pool. For the server to recognize that a particular volume belongs to the correct 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 you create in the Label Templates resource. You then associate a label template with a specific pool in the Pools resource. If you do not associate data with a specific pool, the Solstice Backup 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 your custom template to be available in the Pools resource.

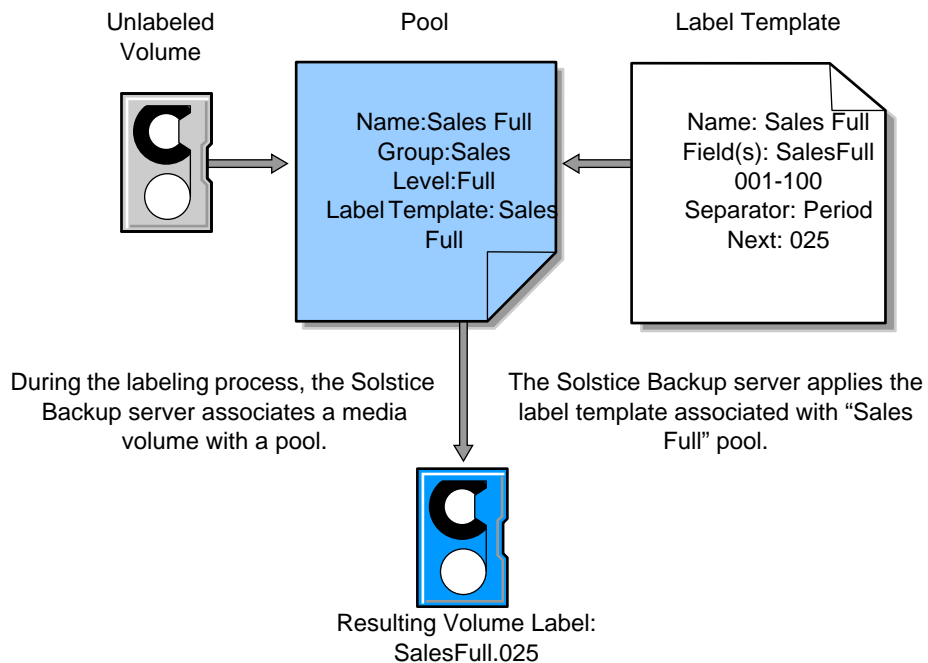


FIGURE 5-2 How the Solstice Backup Server Uses Label Templates

Preconfigured Label Templates

TABLE 5-3 lists the preconfigured label templates contained with the Solstice Backup server. These label templates correspond to the preconfigured pools.

TABLE 5-3 Preconfigured Label Templates

| Name | Description |
|---------------|--|
| Default | Corresponds to the pool receiving all backed up data when another pool is not specified. |
| Default Clone | Corresponds to the pool for receiving data from clones of the Default pool's volume. |
| Archive | Corresponds to the pool receiving all archived data when another pool is not specified. |

TABLE 5-3 Preconfigured Label Templates (*Continued*)

| Name | Description |
|------------------|---|
| Archive Clone | Corresponds to the pool for receiving data from clones of the Archive pool's volume. |
| Full | Corresponds to the pool for receiving data from Full backups. |
| NonFull | Corresponds to the pool for receiving data from incremental and other non-Full backups. |
| Offsite | Corresponds to the pool for receiving data to be stored offsite. |
| PC Archive | Corresponds to the pool for receiving data from PC Archive clients. |
| PC Archive Clone | Corresponds to the pool for receiving data from cloned PC Archive clients. |
| Two Sided | Corresponds to the pool for two-sided volumes. |

Key Label Templates Attributes

TABLE 5-4 lists key attributes for the Label Templates resource:

TABLE 5-4 Key Label Templates Attributes

| Attribute | Description |
|-----------|--|
| Name | <p>Keep the label name consistent with the pool name. This way you and your users can easily see how the data is organized. The same or similar names can be used. For example, you can create a label template called "AcctFull" for volumes that belong to a pool called "Accounting Full."</p> <p>You only can use alphanumeric characters when you create label templates. The Solstice Backup software does not allow the following characters in label template names:</p> <pre data-bbox="394 597 736 621">/ \ * [] () \$! ^ ' ; ' ~ < > & { }</pre> <p>In addition, you cannot use the following four characters, because they are used as separators in label templates:</p> <ul data-bbox="394 739 568 847" style="list-style-type: none"> • colon (:) • dash (-) • period (.) • underscore (_) |
| Fields | <p>A label template is made up of one or more fields. Each field, or component, provides a layer of specificity to your organizational structure. You can have as many components as you want, but it is best to keep the template simple, with few components. The total length of the label cannot exceed 64 characters.</p> <p>You can use four types of components:</p> <ul data-bbox="394 1065 929 1173" style="list-style-type: none"> • Range of numbers (for example, 001-999) • Range of lower case letters (for example, aa-zz) • Range of upper case letters (for example, AA-ZZ) • Character string (for example, Accounting) <p>Each range includes a start value, a dash (-), and an end value. The start value and the end value must have the same number of characters. For example, use "01-99," not "1-99," or "aaa-zzz," not "aa-zzz." (This rule does not apply to a list of character strings or words; character strings are separated by a blank space.)</p> <p>The order in which you enter each component of the Fields attribute is important. The Solstice Backup server applies each component in a left-to-right order, starting with the first one you enter. TABLE 5-5 illustrates how label templates use components to create a number sequence for volume labels.</p> |

TABLE 5-4 Key Label Templates Attributes (*Continued*)

| Attribute | Description |
|-----------|---|
| Separator | Choose which character 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 | <p>Choose the next sequence number to write on the label that the Solstice Backup 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 the start label value you want. The server continues to generate labels from that point on, according to the rules of the template. If you want the Solstice Backup server to generate the first label for you, leave this attribute blank.</p> <p>When the Solstice Backup 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.</p> |

[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 |
| <ul style="list-style-type: none"> • Character string • Range of numbers | SalesFull 001-100 | SalesFull.001,...SalesFull.100 | 100 |
| <ul style="list-style-type: none"> • 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 your 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 Solstice Backup server reaches the end of the template number sequence, it wraps around to the start value. In [TABLE 5-5](#), for example, after the server uses zz.99 for the 67,600th label, it uses aa.00 for the 67,601st label.

Label Template Naming Strategies

The Solstice Backup server is packaged with preconfigured label templates that correspond to the preconfigured pools. If you choose to create your own templates, you can include as many components in the Fields attribute as necessary to suit your organizational structure. However, it is a good idea to keep the template simple with few components. For example, if you create a label template for your Accounting Department, you can customize your label template in several ways, depending on the size of your storage system and media device capabilities.

TABLE 5-6 illustrates several ways you can use components to organize your labels.

TABLE 5-6 Using Label Template Components

| Type of Organizational Structure | Fields (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) |

▼ To Create Label Templates

To create a label template:

1. Select Customize>Label Templates to open the Label Templates window.
2. Click Create. An empty Name attribute is displayed.
3. Enter the name of the label template in the Name attribute.
4. Enter the components for your label template in the Fields attribute. Press [Return] after entering each component. You can use any or all of the following four components for your label template:
 - Range of number (for example, 001-999)
 - Range of lower case letters (for example, aa-zz)
 - Range of upper case letters (for example, AA-ZZ)

- Character string (for example, *Accounting*)

If you enter a range of numbers of characters, they change incrementally with each new label. For example:

First label:Engineering.001

Second label:Engineering.002

Third label:Engineering.003

5. Select a separator symbol in the Separator attribute. If no symbol is selected, the components will have no separators (for example, AA00aa).
6. Enter the next sequence number to write on the label the Solstice Backup server places on a volume (according to the template) in the Next attribute.
7. Complete any other attributes as necessary and click Apply.

▼ To Edit Label Templates

To edit a label template:

1. Select Customize>Label Templates to open the Label Templates window.
2. Select the label template you want to edit from the list displayed in the Label Templates attribute.
3. Edit the attribute values of the label template and click Apply.

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 Delete Label Template

To delete a label template:

1. Select Customize>Label Templates to open the Label Templates window.
2. Select the label template you want to delete from the list displayed in the Label Templates attribute.
3. Click Delete.

You will be prompted to confirm that you really want to delete that label template. Click OK to confirm. The label template is removed from the list of notifications displayed.

4. Click Apply.

You cannot delete the preconfigured label templates or a label template currently in use.

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 Solstice Backup 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 Solstice Backup server, the server will automatically recognize and display the change when it is restarted.

You must configure each device before the Solstice Backup server can recognize and use it. The Devices resource contains the attributes for each device that you must set. The instructions for configuring your devices differ depending on whether the device is stand-alone or is contained in an autochanger or silo.

Device Support by Solstice Backup Edition

Each edition of Solstice Backup software has a different maximum number of devices a server can administer. Each enabled storage node connected to the Solstice Backup server can increase the device maximum. The maximum number of devices for any Solstice Backup server/storage node combination is 256. Device levels for other Solstice Backup platforms 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 Solstice Backup release, including the maximum for each enabled storage node.

TABLE 5-7 Maximum Devices

| Solstice Backup Product | Server | Each Storage Node |
|-------------------------|--------|-------------------|
| Workgroup Edition | 2 | 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.

You are limited to using nonrewinding devices because the Solstice Backup server writes a file mark on the volume at the end of each backup and then appends data onto the volume based on the position of the file mark. If the device rewinds the media, the file mark position is lost and previously written data is overwritten by the next backup.

Use concurrent devices to spread backups evenly across your storage devices, and use the optional Legato Networker Autochanger Software Module for unattended backups to an autochanger.

Refer to the *Legato Compatibility Guide* for a list of the currently supported storage devices. You can obtain this guide by downloading it from the Legato web site, www.legato.com.

Supported Devices and Media Types

The Solstice Backup 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 your particular Solstice Backup server:

1. Select Media>Devices to open the Devices window.
2. Display the list of supported devices in the Media Type attribute.

This list displays abbreviations for all the device types supported in that Solstice Backup version. Select Help>On Window while the Devices window is open to view more detailed descriptions of each device type.

Filesystem devices are configured differently from traditional devices and are covered separately in this manual. See [“Filesystem Devices” on page 159](#) for information on file device types.

Nonrewinding Tape Device Usage

If you use tape drives as your storage devices, you must use nonrewinding devices. The Solstice Backup server writes a file mark on the volume at the end of each backup and then appends data onto the volume based on the position of the file mark. If the device rewinds the media, the file mark position will be lost and previously written data is overwritten by the next backup. The pathnames for these devices must follow the Berkeley Storage Device (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 Add Devices

You need to create a new device resource for each device stand-alone devices or autochanger device. When you elect to create a new device resource, some of the device attributes have default values. You can use these values if they apply to your storage device, or change them to a different value.

Caution – The Solstice Backup server automatically creates a 8mm tape device during its installation. If this device is incorrect, you will have to delete the existing Device resource and create a new one with the correct media type. The Media Type attribute of an existing Device resource cannot be changed. See [“To Delete Devices” on page 158](#) for information on deleting devices.

To add a device:

1. Select Media>Devices to open the Devices window.
2. Click Create.
3. Enter the device pathname into the Name attribute.
4. Select the appropriate media type from the Media Type attribute.

If you want to back up to a filesystem device configured on your disk, see [“Filesystem Devices” on page 159](#) for specific instructions and restrictions. For more information on the logical selection, see [“Logical Volume Backup” on page 160](#).

5. Select *Yes* in the Read-Only attribute if you want the device to read volumes only (as opposed to writing data to volumes).
6. Enter or select a value in the Target Sessions attribute. See [“Target Sessions” on page 47](#) for more information on how this attribute works.
7. Select *Yes* in the Cleaning Required attribute to enable the cleaning requirement feature.

The cleaning device options are available only if you are backing up to an autochanger. This attribute is only valid for autochanger tape devices.

8. Enter a date in the Cleaning Interval attribute 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:

- 4 days
- 2 weeks
- month (same as 1 month)

This attribute is only valid for autochanger tape devices.

9. Enter the date the device was last cleaned. The Solstice Backup server determines the next scheduled cleaning time using this date in conjunction with the value you entered in the Cleaning Interval attribute. This attribute displays the date of subsequent device cleanings.

This attribute is only valid for autochanger tape devices.

10. Select *Yes* in the Auto Media Management attribute if you want to enable this feature. See [“Auto Media Management” on page 162](#) for information on this feature.
11. Complete any other attributes as necessary and click *Apply*.

▼ To Edit Devices

To edit a device:

1. Select *Media>Devices* to open the *Devices* window.
2. Select the device you want to modify from the list displayed in the *Devices* attribute.
3. Edit the attributes of the notification and click *Apply*.

▼ To Delete Devices

To delete a device:

1. Select *Media>Devices* to open the *Devices* window.
2. Select the device you want to delete from the list displayed in the *Devices* attribute.
3. Click *Delete*.

You will be prompted to confirm that you really want to delete that device. Click *OK* to confirm.

4. Click *Apply*.

Filesystem Devices

One Solstice Backup data storage option is to save data to a computer's local disk, rather than use a tape or optical device to write the backed-up data to storage media. This method is known as *filesystem device* storage. Storing data to the filesystem greatly reduces the time it takes to both save and recover data.

You designate the Solstice Backup server to back up data to a filesystem device by creating a new Device resource for your Solstice Backup server the same was as for any other storage device. See [“To Add Devices” on page 157](#) for instructions on creating device resources.

You can have more than one file type device backing up a filesystem, as long as they have unique names. Each file type device counts as one storage device for the Solstice Backup license, so be sure you do not exceed your device limit. Additionally, if you have multiple file type devices backing up to the same disk, you should partition the disk and limit one device per partition.

▼ To Configure Filesystem Devices

To configure a filesystem device:

1. Create one directory for each disk (or partition) being used for file type device storage.

Creating a filesystem device requires you to create a directory in the disk filesystem that the Solstice Backup server will recognize as the device name (and the destination for the data), for example:

```
/export/backup.data
```

Note: Do not use the *tmpfs* directory for filesystem devices because the data might be overwritten.

1. Select Media>Devices to open the Devices window.
2. Click Create.
3. Enter the complete pathway for the directory in the Name attribute.
4. Select File from the Media Type attribute to indicate that the device is a filesystem device.
5. Enter or select a value for the Target Sessions attribute.

6. Leave the options for cleaning at their default disabled settings so that automatic cleaning is not invoked.
7. Do not enable the Auto Media Management feature for filesystem devices.
8. Complete any other attributes as necessary and click Apply.

You label and mount the file device in the same manner as for other offline media used for Solstice Backup backups. Use of a file as a device facilitates save set staging. See [“Save Set Staging” on page 177](#) for more information.

Restrictions

Note that NFS-mounted filesystems are not available for filesystem storage. Also, file type device storage volumes cannot be included in a pool that contains other types of volumes. Additionally, the Solstice Backup server does not allow volumes of other media types to be recycled or added to a pool of file type device volumes.

Use with Staging

Saving data to a file type device provides a great increase in the speed of backups and recoveries, but also uses a large amount of disk space. Therefore, most administrators backing up data to file type devices will want to have a staging policy in effect to move the data off the disk as soon as possible to make room for the next backup. See [“Save Set Staging” on page 177](#) for more information on creating staging policies and using staging in conjunction with file type devices.

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 Solstice Backup 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 Solstice Backup server probes (queries) the clients in a backup group (using the `savefs -p` command). Using this command determines:

- What data to back up.
- Where the data is physically located
- Whether there are logical volumes.

It 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:

1. First, the server allocates one save session per client in the backup group.
2. Then, if there are still save sessions available, it allocates one save session per physical disk on each client.
3. If, after that, 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.

Media Handling Error Resolution

The architecture of device drivers can produce media handling errors. For example, trying to mount and read a damaged tape can 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 shorten the number of times this type of loop can occur:

1. Select Media>Devices to open the Devices window.
2. Select View>Details to display the Devices window's hidden attributes.

3. Lower the value (set to 20 by default) in the Max Consecutive Errors attribute and click Apply.

When the number of errors meets the value set in this attribute, the device stops attempting the activity and becomes disabled. Once host running the Solstice Backup server is idle, remove any media from the disabled device and ensure that the device is in good working order.

Next, with the Devices window still open, re-enable the device:

4. Select the disabled device from the list in the Devices attribute.
5. Select *Yes* in the Enabled attribute.
6. Select *Yes* in the Write Enabled attribute, unless you want the volume to be read-only.
7. Click Apply.

Note that if your disabled device is part of an autochanger, you might need to reset the device by opening a terminal window, changing your path to the directory containing the Solstice Backup executables, and entering the following command:

```
# nsrjb -HV
```

Auto Media Management

The auto media management feature gives the Solstice Backup server automatic control over media loaded in the storage device. If you enable the auto media management feature in the Devices resource, the server automatically:

- Labels
- Mounts
- Overwrites a volume 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 305](#).

The Solstice Backup server considers a volume unlabeled in the following conditions:

- The volume has no internal label.

- The volume is labeled with information other than a recognizable Solstice Backup label.
- The volume is labeled with a Solstice Backup 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 Solstice Backup volumes are shared between devices with different densities.

If you do not enable the auto media management feature, the Solstice Backup 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 Solstice Backup 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 Solstice Backup 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 Solstice Backup 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 Solstice Backup server waits for 60 minutes before it automatically remounts the volume in the drive. This hour is considered a reasonable delay to give you or an operator time to unload the volume after unmounting.

Caution – The Solstice Backup server considers volumes that were labeled by a different application to be valid relabel candidates if auto media management is enabled. Once the Solstice Backup 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 Solstice Backup server maintains the client file index and media database. Using Solstice Backup storage node software, you can route client data directly to a storage node's storage devices without the data first going to your Solstice Backup server. See ["Storage Nodes" on page 66](#) for a complete description of storage nodes.

▼ 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 155](#). However, the storage node device name must follow this convention:

```
rd=storage_node_hostname:device_name
```

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 Solstice Backup 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 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 Solstice Backup server to a Solstice Backup 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 Solstice Backup server.

Installing the Software on a New Storage Node

When you install storage node software to create a new storage node, first update your server to the same Solstice Backup release. Then, run the `jbconfig` program on the storage node computer to add a Jukebox resource. After you run the `jbconfig` program on the storage node computer, 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 Solstice Backup Server

If you choose to update only the Solstice Backup 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 configuring a remote autochanger or silo device as a storage node remote device:

1. Verify that the storage node is listed in the Administrator attribute in the controlling Solstice Backup server's Server resource.

Use the *root@hostname* format for UNIX servers or *administrator@hostname* format for Windows servers, 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.

▼ 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. Select Media>Devices to open the Devices window.
2. Select View>Details to display the Devices window's 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 number of minutes specified by the Save Mount Timeout attribute, the storage node is locked out from receiving saved data for the number of minutes specified by the value of 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.

Caution – The Save Mount Timeout only applies to the initial volume of a save request.

Storage Management Operations

The following sections details the various tasks and components involved in the operation of storage volumes and its relationship with pools.

How the Solstice Backup Server Uses Volume Labeling

A volume label is a unique internal code applied the Solstice Backup 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, and refers to volume labels in the records stored in the media database. The Solstice Backup server uses the media database records to determine which volumes are needed for backing up or recovering data.

When it labels a volume, the Solstice Backup server first verifies that the volume is unlabeled. Then it labels the volume with the name specified in the Volume Name attribute, 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 Solstice Backup 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 take advantage of the preconfigured pools and preconfigured (and associated) label templates that come with the Solstice Backup product, or create your own pools, label templates, and pool-template associations. Customizing your own label templates gives you more control over your data storage organization.

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 Solstice Backup Server Selects and Mounts a Volume During a Backup

When a backup takes place, the Solstice Backup server searches for a volume from the appropriate pool to accept the data that needs to be written. The available volumes are:

- Volumes that are mounted on stand-alone devices.
- Volumes that are accessible to the Solstice Backup server through auto media management or available to the server via the autochanger.

If you try to back up files when an appropriate volume is not mounted, the Solstice Backup server requests a writable volume by displaying the following or a similar message:

```
media waiting (or critical): backup to pool 'Default' waiting for
1 writable backup tape or disk
```

If you are recovering data, the Solstice Backup server requests the volume name you need to mount in the device by displaying the following or a similar message:

```
media waiting: recover waiting for 8mm 5GB tape_volume_name
```

If you need more than one volume to recover your data, the Solstice Backup 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 Solstice Backup 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 Solstice Backup Server Selects a Volume for Relabeling

Backup data is destined for a specific pool. When the data is ready to be written, the Solstice Backup server monitors the active devices to locate a volume from the appropriate pool.

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 Solstice Backup server will consider in the following factors to guide its volume selection:

- Volume mode

The Solstice Backup 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 Solstice Backup 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 Solstice Backup server mounts the volume on the lowest-numbered drive. If there are no empty drives, the server will select a loaded idle drive and unload it to make room for the needed volume.

Volume Status

The Solstice Backup 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 Solstice Backup 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 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 field 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 176](#). See ["Client/Save Set Combination" on page 101](#) for more information about save set status.

[TABLE 5-8](#) lists the Solstice Backup volume modes and their definitions.

TABLE 5-8 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 are considered read-only. The mode can only be changed manually. |

Read-Only Mode

Changing the mode of a volume to Read Only prevents new data from being written to the volume. The volume mode, however, is still subject to the retention policies of its save sets. 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.

See [“To Change the Mode of a Volume” on page 176](#) to change the volume mode.

Solstice Backup Volume Operations

If you try to back up files when a volume is not mounted, you will 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 Solstice Backup 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

Select Media>Volumes or click the Volumes toolbar button to open the Volumes window.

The Volumes window displays the information for the server’s volumes, including the volume name, the amount of data written to each volume, the percentage currently used, and 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 5-9 lists the volume-related categories displayed in the Volumes window.

TABLE 5-9 Volumes Window (Volume-Related Categories)

| Category | Description |
|----------|--|
| Name | Volume name. |
| Bar code | Bar code 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. See “To Change the Mode of a Volume” on page 176 for more details on the type of volume modes. |
| Location | Location of the volume, entered by the user as a reminder. |

▼ To View Save Set Information

Select Media>Volumes or click the Volumes toolbar button to open the Volumes window.

The Volumes window lists the save sets stored on a volume. This window displays information such as the client name, the save set name, the size, and many other categories. You can also change the status of a save set from here.

TABLE 5-9 lists the volume-related categories displayed in the Volumes window.

TABLE 5-10 Volumes Window (Save Set-Related Categories)

| Category | Description |
|----------|--|
| Client | Name of the Solstice Backup 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. Manual backups are not displayed in this window. |

TABLE 5-10 Volumes Window (Save Set-Related Categories)

| Category | Description |
|----------|--|
| Status | Type of save set. For a listing of save set values and descriptions, see “Save Set Status Values” on page 102 . |
| 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 Volumes

To mount a volume:

1. Insert a volume in the Solstice Backup server storage device or ensure that volumes are already present for the server to access.
2. Select the appropriate device in the Devices list of the Solstice Backup Administrator program.
3. Select Media>Mount or click the Mount toolbar button.

After you label and mount a volume, the volume name is displayed in the Devices list of the Solstice Backup Administration 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 Solstice Backup 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 Berkeley Storage Device (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 Volumes in One Operation

If you have more than one storage device connected to your Solstice Backup server, you must first select the device you want to use for labeling from the Devices list in the Solstice Backup Administrator program. Relabeling a volume makes it impossible for the Solstice Backup 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 Solstice Backup server storage device, or ensure that these types of volumes are already present for the Solstice Backup server to access.
2. Select the appropriate device in the Devices list of the Solstice Backup Administrator program.
3. Select Media>Label or click the Label toolbar button to open the Label window.
4. Select a pool in the Pool attribute. This action will also generate a name for the Volume Name attribute. The following rules apply:
 - The Solstice Backup server assigns the next sequential label from the label template associated with the chosen pool if the volume is considered unlabeled. 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 Solstice Backup 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. Check 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 Solstice Backup server disregards the assigned browse and retention policies. Therefore, only an administrator can mark the volume recyclable.
6. Check the Mount After Labeling attribute and click OK. the Solstice Backup 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 Solstice Backup label is internal and machine-readable, it is a good idea to put an adhesive label on each volume that matches the internal volume label.

See [“Bar Code Labeling of Volumes” on page 310](#) for information on using barcode labels,

▼ To Label Without Mounting

You may prelabel volumes without mounting them. Follow the same procedures as for labeling and mounting in one operation, but deselect the Mount After Labeling attribute in the Label window.

Labeling Tips

Naming schemes vary from site to site. We suggest naming the volumes with the name of the Solstice Backup server followed by a three-digit number, for example:

```
jupiter.001
```

No matter what naming scheme you choose to use, keep in mind that a simple convention will be easy to understand by operators and administrators.

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

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

▼ To Unmount Volumes

To unmount a volume:

1. Select the appropriate device in the Devices list of the Solstice Backup Administrator program.
2. Select Media>Unmount or click the Unmount toolbar button.

Once the volume is unmounted from the device, the device's listing in the Solstice Backup 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 two methods:

- Mount the volume in the device. The volume name appears next to the device name in Solstice Backup Administrator program. See [“To Mount Volumes” on page 173](#) for more information.

or

1. Load the volume in a device.
2. Select Media>Label or click the Label toolbar button.

When the Label window appears, the name of the label appears in the Volume Name attribute.

3. Click Cancel once you determine the volume name to avoid accidentally labeling the volume.

▼ To Change the Mode of a Volume

The mode of a volume determines whether the Solstice Backup 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 believe you do not need the data on the volume. The Solstice Backup 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. Select the appropriate device in the Devices list of the Solstice Backup Administrator program.
2. Select the Media>Volumes or click the Volumes toolbar button to open the Volumes window.
3. Select Volume>Change Mode> and one of the following:
 - Appendable
 - Read Only
 - Recyclable

See [TABLE 5-8](#) for a listing of the volume modes.

4. Select a Mode button and click OK.

The changed volume mode is now displayed in the Devices list of the Solstice Backup Administrator program.

▼ To Remove 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 Solstice Backup 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 Solstice Backup 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 your 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. See [“To Remove Volume-Based Entries from the Online Indexes” on page 77](#) for instructions on how to perform these tasks.

Save Set Staging

Save set staging is a process of transferring data from one storage medium to another. Use save set staging to move save sets that have been backed up, archived, or migrated. You can create, edit, and delete staging policies similar to other Solstice Backup resources. When you enable a staging policy, the Solstice Backup 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 Solstice Backup 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.

Staging Data from a Filesystem Device

Staging is particularly well suited for data that has been backed up on filesystem devices. Staging allows the occupied disk space on filesystem 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 filesystem device. Staging also allows data to be moved off the file type device outside the backup period, ensuring that sufficient disk space is available for the next backup session. See [“Filesystem Devices” on page 159](#) for more information on backing data to filesystem devices.

▼ To Create Staging Policies

To create a staging policy:

1. Select Customize>Staging to open the Stage window.
2. Click Create.
3. Enter the name of the staging policy Name attribute.
4. Select Yes in the Enabled attribute if you want the staging policy to start functioning right away. You have the option of enabling or disabling staging at any time.
5. Enter or select a number in the Max Storage Period attribute. This value indicates the maximum number of days for a save set to be in a volume before it is staged to a different storage medium.
6. Enter or select a number in the High Water Mark (%) attribute. This value is the percentage of available space on the Solstice Backup server that is filled with data. Once this value is met, staging will begin. Staging will continue until the low water mark level is met (see next step).
7. Enter or select a number in the Low Water Mark (%) attribute. This value is the percentage of available space on the Solstice Backup server that is filled with data. Once this value is met, staging will cease.
8. Select one of the choices from the Save Set Selection attribute. This attribute allows you to determine which save sets are selected for staging.
9. Select the destination pool for the staged data from the Destination Pool attribute.

10. Check the filesystem device(s) that will perform the staging task. You must configure the filesystem device before you configure the stage policy in order for the device to appear as a choice. You can assign multiple devices to the staging policy, but a given device should not be controlled by more than one staging policy.

11. Select one of the choices in the Start Now attribute:

- Recover space
- Check file system
- Stage all save sets

Note that the choice you make will start immediately. If you do not want any action to start now, leave this attribute alone.

12. Click Apply.

▼ To Edit Staging Policies

To edit a staging policy:

1. Select Customize>Staging to open the Stage window.
2. Select the staging policy you want to edit from the list displayed in the Stage attribute.
3. Edit the attribute values of the staging policy and click Apply.

▼ To Delete Staging Policies

To delete a staging policy:

1. Select Customize>Staging to open the Stage window.
2. Select the staging policy you want to delete from the list displayed in the Stage attribute.
3. Click Delete.

You will be prompted to confirm that you really want to delete that staging policy. Click OK to confirm.

4. Click Apply.

Backup Scheduling Tools

This chapter provides information about the Solstice Backup product's scheduling features, including automated group backups and customizable backup schedules.

Overview of Solstice Backup Scheduling

Scheduled backups of client data is the result of two Solstice Backup resources working in tandem:

- Group
- Schedule

The Group resource, also known as backup groups, are a resource that specifies a specific time of the 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 network traffic load associated with backups does not interfere with regular work activities. Any client assigned to a specific group will have its specified save sets backed up during that time.

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 is also specified in this resource. Because scheduled backups are configured in terms of weekly or monthly, the type of backup scheduled for a specific 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.

Together, the Group and Schedule resources provide you with the maximum flexibility in scheduling large numbers of differing client backups each day, and doing so in a method that causes the least amount of disruption and inconvenience to your enterprise. The rest of this chapter describes these resources in greater detail.

Backup Groups

Solstice Backup backup groups provide a way for you to designate what time a client's scheduled backup starts. Using groups, you can schedule the backups to take place in the middle of the night, or some other time when network traffic is low. 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. You can also stagger groups at different times of the day to avoid an excessive strain on the network or the Solstice Backup server. See ["Schedules" on page 193](#) for more information about creating schedules that stagger the days that different clients perform full backups.

Each client backing up to a Solstice Backup server must be assigned to a back up group, either a preconfigured group or a group you create yourself. In many cases, the Solstice Backup 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 configuration window, which is where you associate a group with a pool. Remember, you can assign a client's save sets to more than one group.

How the Solstice Backup Server Uses Groups

The client save sets 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.

Any time a Solstice Backup client is added to a Solstice Backup server, either by creating a new Client resource or by adding an existing client, that client is assigned to a backup group. The group determines the time during the day when the assigned client, and any other clients also assigned to the same group, have their data backed up to the server.

Example: Using Groups to Back Up Multiple Clients

FIGURE 6-1 illustrates how the Solstice Backup server uses backup groups to back up multiple client save sets. In the example shown, three client computers—*Oak*, *Elm*, and *Fir*—are part of the group named *Weekly Full*, which starts its automatic scheduled backup at midnight. Client *Oak* runs a full backup of all its save sets every Monday and incremental backups of its save sets on the other days; client *Elm* runs a full backup of all its save sets on Tuesday and incremental backups on the other days; and client *Fir* 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 *Birch* and *Pine* and starts its backups at 7:00 p.m., when the computers in the Accounting Department are available for backup. Although the two client computers run full backups on the same day, computer *Pine* is scheduled to back up only the `/usr/home` save set; all the save sets on computer *Birch* 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 Solstice Backup 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: | Oak | Client: | Birch |
| Save Set: | All | Save Set: | All |
| Schedule: | Full, Monday incremental, all other days | Schedule: | Full, 1st of month Level 5, 15th of month incremental, all other days |
| Client: | Elm | Client: | Pine |
| Save Set: | All | Save Set: | /usr/home |
| Schedule: | Full, Tuesday incremental, all other days | Schedule: | Full, 1st and 15th of month incremental, all other days |
| Client: | Fir | | |
| Save Set: | All | | |
| Schedule: | Full, Wednesday incremental, all other days | | |

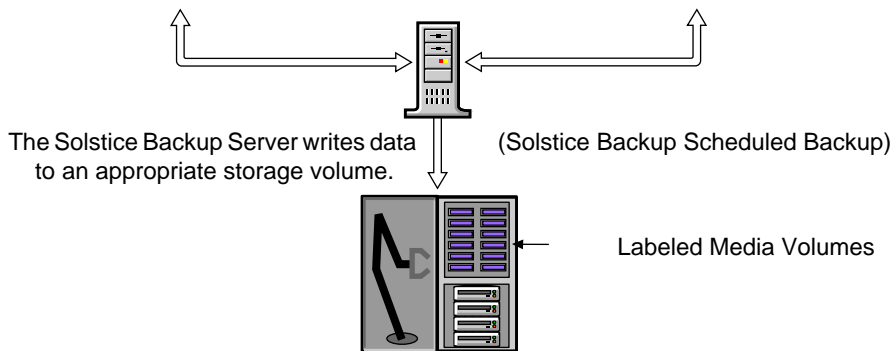


FIGURE 6-1 How Solstice Backup Groups are Used to Back Up Multiple Clients

Preconfigured Groups

The Solstice Backup product ships with a single preconfigured group named *Default*. To ensure that all data is backed up, the Solstice Backup server automatically adds all clients to the Default group. However, you must enable the Default group for the Solstice Backup server to back it up. Depending on your needs, you can keep a client in the Default group, or you can put the client in one or more customized groups.

You can 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

Depending on your needs, you can keep a client in the Default group, or you can put the client in a customized group.

TABLE 6-1 displays some of the key attributes for the Group resource.

TABLE 6-1 Key Group Attributes

| Attribute | Action |
|----------------------|---|
| Start Time | The Default group is preconfigured to start its daily backup at 3:33 a.m. This time can be changed, if desired. |
| Backup Time Interval | This group attribute dictates how often a group will start a scheduled backup. The default value is once a day, but you can change the value so that a client's group backups run more often. See "To Set Group Backup Time Intervals" on page 186 for instructions for modifying this attribute. |
| Autostart | You must enable the Autostart attribute for the Default group, and any other group you create, before the Solstice Backup server can run a scheduled backup. |
| Client Retries | If the Solstice Backup server is unable to make a connection with a client, the client will retry connecting up to the value specified in the Client Retries attribute before the backup is considered a failure. The first retry will not occur until after an attempt has been made to at least contact each client in the group. |
| Inactivity Timeout | The Inactivity Timeout attribute in the Group resource specifies the number of minutes that the Solstice Backup server waits for evidence of backup activity on the client. If the server has not received status information for longer than the time specified, the server abandons the backup operation for the save set. |

Abandoned Group Backups

The backup of an abandoned save set might be completed, but the automated report from `savegrp` will not show that the backup is completed. For example, if the client is being backed up over a network filesystem (NFS) connection and the NFS server crashes and reboots, the Solstice Backup backup hangs until it times out. The Solstice Backup server marks the save set "abandoned," and when the NFS server comes back up, the backup continues and is completed.

▼ To Create Groups

To create a group:

1. Select Customize>Groups to open the Groups window.

2. Click Create.
3. Enter the name of the group in the Name attribute. Do not include spaces in a group name.
4. Select Enabled from the Autostart drop-down list.
5. Enter a new time in the Start Time attribute, unless you want to keep the default time of 3:33 a.m.

Note: When you select a start time for your different groups, make sure you schedule them far enough apart so that one group has completed backing up before the next group starts.

6. Enter or select a number in the Client Retries attribute. This value specifies the number of times the Solstice Backup server will attempt to back up a failed client.
7. Select Yes in the Clones attribute to automatically clone the group's backed-up data. If you select No, do not complete the subsequent clone-related tasks in the subsequent steps and skip to step 10.
8. Select the appropriate clone pool for the cloned data in the Clone Pool attribute (requires cloning to be enabled).
9. Select the appropriate migration clone pool in the Migration Clone Pool list (requires cloning to be enabled).
10. Enter the name of the printer on which `savegrp` will print the bootstrap save set information in the Printer attribute.
11. Complete any other attributes as necessary and click Apply.

When you finish creating a backup group, the group appears as a choice in the Client resource where you assign each client individually to a group.

▼ To Set Group Backup Time Intervals

The Solstice Backup 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 Interval attribute is a hidden attribute, meaning that you must set the Solstice Backup Administrator program to display hidden attributes before you can access and change the Interval value.

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 backup twice a day. For example, a group with the default start time of 3:33 a.m. and an interval of 12:00 would backup twice a day, first at 3:33 a.m., and then again twelve hours later at 3:33 p.m.

For best results, use time interval values that make determining the group backup time easy to determine, such as 24, 12, or 6 hours. Also note that increasing the group backup time interval from once every 24 hours can result in a greatly added strain on your network, the Solstice Backup server, and associated devices.

To set group backup time intervals:

1. Select **Customize>Groups** to open the Groups window.
2. Select **View>Details** to display the Group resource's hidden attributes.
3. Enter the interval attribute value in hours, using hh:mm format, and click **Apply**.

▼ To Force Incremental Backups

You can instruct the Solstice Backup server to perform only incremental backups on groups that back up more than once every 24 hours. Forced incremental backups under these conditions might be advantageous if the amount of data changed during the 24 hour period does not warrant a more comprehensive backup level.

The Force Incremental attribute in the Group resource must be enabled to enact this feature. The Force Incremental attribute is a hidden attribute, meaning that you must set the Solstice Backup Administrator program to display hidden attributes before you can access and change the Interval value. Setting the Force Incremental attribute to *Yes* will cause only incremental backups to occur from that point on until the attribute value is changed to *No*. Forced incremental backups override all scheduled backup levels. This can effect volume recycling practices.

To force incremental backups on groups:

1. Select **Customize>Groups** to open the Groups window.
2. Select **View>Details** to display the Group resource's hidden attributes.
3. Select *Yes* in the Force Incremental attribute.
4. Click **Apply**.

Repeat these steps and select *No* in the Force Incremental attribute to disable the function and allow the scheduled backup levels to occur.

See [“To Set Group Backup Time Intervals” on page 186](#) for more information about configuration group backups that occur more often than once every 24 hours.

▼ To Edit Groups

To edit a group:

1. Select Customize>Groups to open the Groups window.
2. Select the group you want to edit from the list displayed in the Groups attribute.
3. Edit the attributes of the group and click Apply.

▼ To Delete Groups

To delete a group:

1. Select Customize>Groups to open the Groups window.
2. Select the group you want to delete from the list displayed in the Groups attribute.
3. Click Delete.

You will be prompted to confirm that you really want to delete that group. Click OK to confirm.

4. Click Apply.

Group Backup Management

Beside providing you with the ability to create and configure backup groups, the Solstice Backup product also provides a variety of tools and features to manage and monitor the operations and progress of group backups.

This section describes how to accomplish the following tasks involving Solstice Backup backup groups:

- Immediately start a scheduled backup
- Stop a scheduled backup
- Restart a scheduled backup
- View the details of a scheduled backup

▼ To Start Groups Immediately

When you start a scheduled backup group manually (on demand), the Solstice Backup server runs the backup at the level of the next scheduled backup, which can be full, level 1-9, consolidated, or incremental.

To start a group immediately:

1. Select Server>Group Control or click the Group Control toolbar button to open the Group Control window.
2. Select the name of the group you want to start immediately in the Groups attribute.
3. Click Start to initiate the group backup.

When you start the group in this manner, the Solstice Backup server overrides the groups' scheduled start time and immediately backs up the clients in the group.

▼ To Stop Groups

To stop a group immediately:

1. Select Server>Group Control or click the Group Control toolbar button to open the Group Control window.
2. Select the name of the group you want to stop in the Groups attribute.
3. Click Stop to stop the group backup.

After you initiate the stopping of the group backup, the Solstice Backup server completes its backup of the current save set, halts the rest of the scheduled backup, and displays *Not Finished* in the Status attribute in the Group Control window.

Caution – If you manually stop a scheduled backup, the Manual Restart option is set automatically, overriding the Autorestart setting. To ensure that the scheduled backup restarts automatically, you should disable the Manual Restart option in the Solstice Backup Administration program.

The Manual Restart option is enabled each time you stop a scheduled backup manually. Consequently, you must disable the Manual Restart option each time you have stopped a scheduled back up manually.

▼ To Restart Groups

To restart a group:

1. Select Server>Group Control or click the Group Control toolbar button to open the Group Control window.
2. Select the name of the group you want to restart in the Groups attribute.
3. Click Restart to reinitiate the group backup.

After you initiate the restart of a group backup, the Solstice Backup 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. Select Server>Group Control or click the Group Control toolbar button to open the Group Control window.
2. Select the group whose details you want to view in the Groups attribute.
3. Click Details to open the Group Control Details window.

The Group Control Details window displays the status of client save sets in the backup process in one of three message fields:

- 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 Solstice Backup server has successfully backed up.
- Failed Save Sets – displays the client save sets that the Solstice Backup 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 Solstice Backup server failed.

- The Solstice Backup client failed.
- The network connection failed.

▼ To Preview a Group Backup

You can simulate a backup for a specific group by utilizing the Group Control Preview feature. This feature helps you identify potential problems before the Solstice Backup server runs an upcoming group backup. The simulated results are then displayed in the Group Control Details window.

1. Select Server>Group Control or click the Group Control toolbar button to open the Group Control window.
2. Select the group whose backup you want to preview in the Groups attribute.
3. Click Preview to open the Group Control Details window.

Save Group Completion Message

When the backup is completed, the Solstice Backup server generates a report about the success of the scheduled backup. The server sends the root user an automatic notification and displays the same information in the Solstice Backup Administrator program.

Bootstrap Generation and Printout

When the backup group includes the Solstice Backup server, the server generates a special save set called the *bootstrap*, which includes the media database and configuration files. The bootstrap information is essential for recovery from a disaster. Refer to the *Legato Networker Disaster Recovery Guide* for information on how the bootstrap is used during a disaster recovery operation.

By default, the bootstrap is printed to the Solstice Backup server's default printer. To change the default printer, change the Printer attribute in the Group resource.

A bootstrap printout is created with any scheduled backup of a group that includes the server, or after other scheduled backups if the server is not in an active group. A bootstrap printout is generated whether the scheduled backup is initiated automatically or manually.

▼ To Save the Bootstrap to a File

To save the bootstrap to a file:

1. Select Customize>Notifications to open the Notifications window.
2. Select Bootstrap in the Notifications attribute. The Action attribute displays the following value:

```
/usr/bin/lp -s -c -t bootstrap -d <printer_name>
```

3. Change the Action attribute value to the following in order to create a file containing the bootstrap information:

```
/bin/cat >> /directory/filename
```

▼ To Send the Bootstrap File Via E-mail

To e-mail the bootstrap file to a user ID, enter the following value in the Action attribute instead:

```
/usr/ucb/Mail -s nwserver_bootstrap user@hostname
```

System Console Log

The UNIX system log displays messages passed from the Solstice Backup server. When the Solstice Backup 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
```

How the Solstice Backup Server Backs Up Open Files

You can back up most open files using the Solstice Backup product. If an open file changes during a scheduled backup, the Solstice Backup server “notices” 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.

Schedules

The Solstice Backup server backs up each client computer across your network according to a backup schedule. A client’s backup schedule tells the Solstice Backup 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 your 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. See [“Client Configuration” on page 89](#) for more information about client configuration.

How the Solstice Backup Server Uses Schedules

The Solstice Backup 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 levels* you specify. Backup levels can be set to backup up a client's entire filesystem, or only data that has changed since the last full backup. See "Backup Levels" on page 202 for more information on how backup levels work.

You customize each client's backup schedule to best suit your needs in the Schedules dialog box. For example, some clients may have data you want to back 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 and/or level [1 – 9] backups on other days.

Preconfigured Solstice Backup Schedules

Solstice Backup software ships with preconfigured schedules. If these schedules fit your backup requirements, you can use them as is. Otherwise, you can create new ones to accommodate your site-specific needs.

The Default schedule can be modified, but not deleted. Schedules marked with an asterisk cannot be modified because they are based on predefined lists of dates that are initialized each time the server is started. All other schedules can be deleted or modified (except for a schedule's name, which can never be modified). If you want to change the name of a schedule, you must delete it and recreate it, giving it a new name.

TABLE 6-2 describes the preconfigured schedules.

TABLE 6-2 Preconfigured Solstice Backup Schedules

| Schedule Name | Solstice Backup Backup Operation |
|-----------------------------|---|
| Default | Completes a full backup every Sunday, incremental backups on all other days. You cannot delete this schedule. |
| Full Every Friday | Completes a full backup every Friday, incremental backups on all other days. |
| Full on 1st Friday of Month | Completes a full backup on the first Friday of the month, incremental backups on all other days. The Solstice Backup product ships with the overrides already preset for this schedule. The overrides for this schedule carry over from year to year. You cannot modify the values of this schedule. |
| Full on 1st of Month | Completes a full backup on the first calendar day of the month, incremental backups on all other days. |

TABLE 6-2 Preconfigured Solstice Backup Schedules (*Continued*)

| Schedule Name | Solstice Backup Backup Operation |
|------------------------------------|--|
| Quarterly | <p>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.</p> <p>When customizing a quarterly schedule, use the Month period to set the level backups, then use an override to set each quarterly full backup on the calendar. You cannot modify the values of this schedule.</p> |
| Consolidate Every Friday | Completes a consolidated backup every Friday; completes incremental backups on all other days. You cannot modify the values of this schedule. |
| Consolidate on 1st Friday of Month | Completes a consolidated backup on the 1st Friday of the month; completes incremental backups on all other days |
| Consolidate on 1st 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. You cannot modify the values of this schedule. |

Example: Weekly Backup Cycle

The period of time from one full backup to the next full backup is called a *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 your Solstice Backup server. Depending on the size of your 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 your clients. The Default schedule tells the Solstice Backup 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.

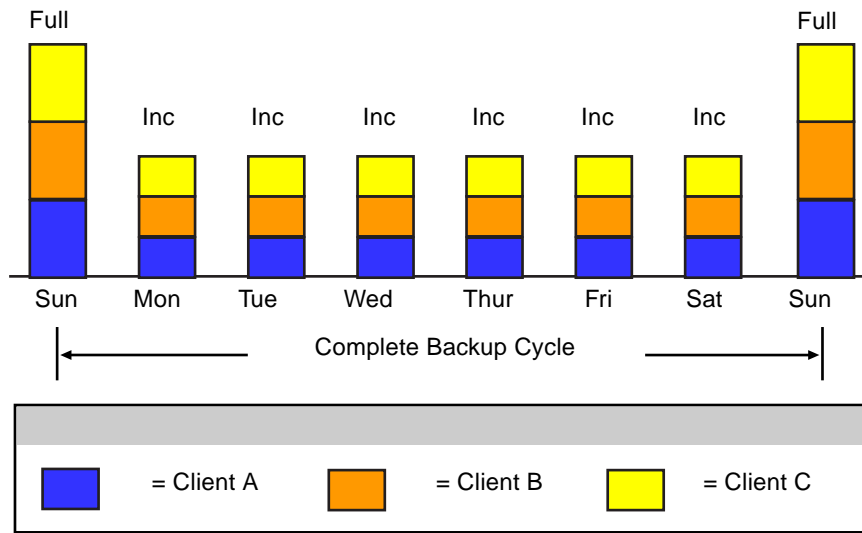


FIGURE 6-3 Using the Solstice Backup Default Schedule for Multiple Clients

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-level backup. See [“Save Set Consolidation” on page 207](#) for more details.

Since full backups transfer large amounts of data and typically take longer than other backup levels, you may want to stagger them throughout the week. For example, you could apply a schedule that performs a full backup for Client A on Sunday, 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 Thursday.

Example: Staggered Weekly Schedules for Multiple Clients

FIGURE 6-4 illustrates how you can use staggered backup schedules 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 your Solstice Backup server.

Schedule Planning Considerations

The Solstice Backup software makes it easy to set up your backup schedules. Deciding which schedules best fit your 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 want 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 equal to three months), you will 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 Solstice Backup server is determined by the browse and retention policies associated with each client. See [“Browse and Retention Policies” on page 104](#) for more information on data life cycle management.

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 5.5 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 addressing the different filesystems in the Client resource. For example, configure the first Client resource to back up one filesystem, `/usr`, with one backup schedule in one group. Then, configure the second Client resource to back up another filesystem, `/var`, 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 an explicit list are omitted from backup. This includes any new disk volumes that are added to the system. This risk of omission does not exist when you enter the special value “All” in the Save Set attribute. The Solstice Backup software automatically adds the new disk volumes to the backups.

Key Schedule Attributes

To create a customized backup schedule, you must define the following schedule configuration values in the Schedule resource in the Solstice Backup Administrator program.

[TABLE 6-3](#) lists the key attributes contained in the Schedules window.

TABLE 6-3 Key Schedule Attributes

| Attribute | Action |
|-----------|--|
| Name | <p>The name you give your customized schedule will appear in the Client resource as an attribute you can apply to a client/save set. You should choose a simple, descriptive name; for example, <i>Monday Full</i>.</p> <p>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.</p> |
| Period | <p>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.</p> <ul style="list-style-type: none">• When you select <i>Week</i> and set up a schedule, the backup level <i>Full</i> is applied to that day of the week for all weeks in the calendar year; for example, full backups every Sunday. <i>Week</i> is the default setting.• When you select <i>Month</i> and set up a schedule, the backup level is applied to that day of the month for all months in the calendar year; for example, full backups on the 15th of each month. |

TABLE 6-3 Key Schedule Attributes (*Continued*)

| Attribute | Action |
|-----------|---|
| Set Level | <p>In addition to full and incremental backups, you can use 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:</p> <ul style="list-style-type: none"> • Full • Incremental • Level [1 – 9] • Consolidated • Skip <p>See “Backup Levels” on page 202 for more information on how backup levels work. See “Save Set Consolidation” on page 207 for more information on how save set consolidation works.</p> |
| Override | <p>Specify a list of actions and dates overriding 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.</p> |

Note: You might also want to change the Force Incremental attribute value located in the Group resource. The default value for this attribute is *Yes*. This means an incremental backup will occur if the group is run more than once a day. Set this attribute to *No* to do more than one full backup per day.

Resource Configuration Order for Schedules

If you want to use your own customized schedule, you must configure the schedule before you can associate it with a client or save set in the Client resource. The start time for your automatic daily scheduled backup is determined by the backup group with which the client save sets are associated. The length of time that the data is available for browsing or recovery is determined by the browse and retention policies you configure for the client’s save sets, rather than by the schedule.

▼ To Create Schedules

To create a schedule:

1. Select Customize>Schedules to open the Schedules window.
2. Click Create.

3. Enter the name of the schedule in the Name attribute.
4. Select a weekly or monthly period for the schedule in the Period attribute. The value you select will cause each corresponding day of the week, or the month, to duplicate that backup level. Duplicate values are assigned based on the day of the week, not the date of the month.
5. Select a backup level for each day of the month displayed in the Schedules window. Click Previous Month or Next Month to see the backup calendars for preceding or future months.
 - If you selected the Week value for the Period attribute, whatever backup level you assigned for a given day is duplicated for the corresponding day of the other weeks in the month.
 - If you selected the Month value for the Period attribute, the backup level you selected for a given day of a given week is passed on accordingly to the other months. For example, if the first Sunday of a month was designated as a full backup, then the first Sunday of all the months would have that corresponding backup level.

See [“Backup Levels” on page 202](#) for more information on how backup levels work.

6. Complete any other attributes as necessary and click Apply.

You can set up as many Solstice Backup 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 Schedules

To edit a schedule:

1. Select Customize>Schedules to open the Schedules window.
2. Select the schedule you want to edit from the list displayed in the Schedules attribute.
3. Edit the attributes of the schedule and click Apply.

▼ To Delete Schedules

To delete a schedule:

1. Select Customize>Schedules to open the Schedules window.

2. Select the schedule you want to delete from the list displayed in the Schedules attribute.
3. Click Delete.

You will be prompted to confirm that you really want to delete that schedule. Click OK to confirm.

4. Click Apply.

▼ To Override a Client's Regular Backup Schedule

You can use a group's Level and Schedule attributes 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.

Backup Levels

Because it may not be practical or efficient for you to run full backups every day, the Solstice Backup 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 your 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 lists the backup levels supported 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 lower numbered backup. For example, a level 1 backup backs up all the files that have changed since the last full backup (considered a level zero). A level 3 backup backs up all the files that have changed since the last level 2 backup, level 1 backup, or full backup. A level 9 backs up all the files that have changed since the last level 8, 7, 6, 5, 4, 3, 2, 1, or full (0) backup. |
| 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. See “Save Set Consolidation” on page 207 for more information. |
| 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, but the Solstice Backup 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 Solstice Backup server to hang waiting 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 the Solstice Backup server Uses Backup Levels

A backup schedule defines what level backup the Solstice Backup 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 are considering using backup levels in a customized schedule, consider the following issues to help you make decisions that best suit 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 your schedules because they collect all the files that have changed over several days, or even weeks, into a single backup session. Using level backups can simplify but slow file recovery.
- 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 original 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 | <ul style="list-style-type: none"> • Faster recovery | <ul style="list-style-type: none"> • Slow backup • High server load • High load on client and network • Uses more volume space |
| Level | <ul style="list-style-type: none"> • Faster backup • Low load on server • Uses least volume space | <ul style="list-style-type: none"> • Slow recovery • Data can spread across multiple volumes |
| Consolidated | <ul style="list-style-type: none"> • Faster backup (from the client's perspective only) • Faster recovery • Low load on client and network | <ul style="list-style-type: none"> • Longest high load on server • Requires at least two volume drives • Uses most volume space |

The online client file indexes, server index, and media database are backed up whenever the Solstice Backup server is backed up. In general, they are backed up at the same level as the server.

For example:

- If the Solstice Backup server's backup is a full-level backup, the backup levels of the online client file indexes, server index, and media database are also full.
- If the Solstice Backup server's backup is a level 5, the backups of the online client file indexes, server index, and media database is also a level 5.

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

How the Solstice Backup Server 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 your level [1 – 9] backups help you 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 your data by using directives, which compress and eliminate unnecessary data from your backups. For example, you can use a directive that tells the Solstice Backup server to skip certain files or filesystems when performing a backup. For more information on directives, see [“Directives” on page 127](#).

The following three examples illustrate the concept of how Solstice Backup backup levels work and the data requirements for recovery in the event of data loss.

Example: Backup Levels Part 1

As shown in [FIGURE 6-5](#), 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 only two volumes: the full volume from October 1 and the new level 7 volume. You no longer need the data on the volumes 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

To continue the example illustrated in [FIGURE 6-5](#), [FIGURE 6-6](#), a level 5 backup on October 8, backs up everything that changed since the full backup on October 1. To fully recover from a disk crash on October 8, you need only two volumes: the full volume from October 1 and the new level 5 volume. You no longer need the data on the volume 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 volume includes that information.




FIGURE 6-6 Backups for October 1 Through October 14

Also, on October 11 a level 7 backup 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

Continuing the example, in [FIGURE 6-7](#) a level 5 backup on October 15 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 backup 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 two volumes: the level full volume from October 1 and the new level 5 volume.

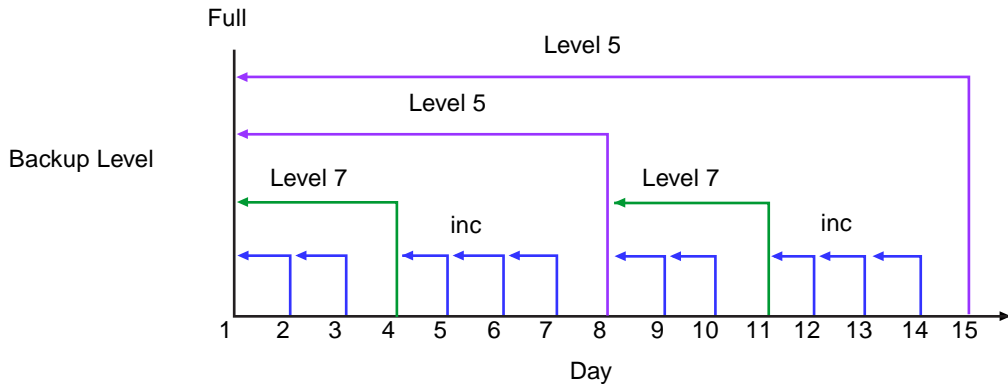


FIGURE 6-7 Backups for October 1 Through October 15

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 really a full level backup and the resulting save set of the consolidation process is the same as a full level save set. Although a consolidated backup takes place, there is really 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 Solstice Backup backup levels traditionally appear. For example, you cannot find *consolidation* as a potential backup selection in the Pool resource. The one exception is the Schedules attribute, where you can use the default Consolidation schedules provided, or create your own 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 restore.

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 recover.

Caution – If there are no existing Full backups the first time a consolidation backup begins, the consolidation backup defaults to a level of Full. Once a Full backup exists, subsequent consolidation backups are performed as a level 1 followed by consolidation with the existing Full.

Requirements

To initiate save set consolidation by the server, the client must be running a Legato ClientPak version of 5.5 or later. In version 5.5, the administrative interfaces within the Legato ClientPak software were updated to provide for initiation of save set consolidation. Save set consolidation will work for save sets of clients whose Legato ClientPak version is earlier than 5.5, but the consolidation must be initiated by the Solstice Backup server itself.

Save set consolidation requires at least two attached tape drives. For better performance, you should have three or more tape drives available for consolidated backups.

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 your 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.

Limitations of Save Set Consolidation

There are some limitations to the save set consolidation process:

- 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.
- You can only administer the server performing the save set consolidation using the graphical Solstice Backup Administrator program from a Solstice Backup release that includes the save set consolidation feature.

When one or more of these events occur, save set consolidation fails with an error message.

In addition, we recommend you 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 will be aborted. The transactional log assists save set

consolidation in putting back changes made to the online index. In most instances, the Solstice Backup server only makes online index changes when data has been written to tape.

▼ To Start Save Set Consolidation

To start save set consolidation:

1. Select the Customize>Schedules to open the Schedules window.
2. Click the button corresponding to the day during which you want the save set consolidation to occur from the calendar, and select Consolidate from the list of levels. The letter *c* now appears in the button you clicked earlier, signifying the newly created consolidated backup.

Just as with other backup levels, the Solstice Backup server will consequently invoke the save set consolidation backup on the scheduled date.

For a consolidated backup, the Solstice Backup server will first perform a level 1 backup; subsequently, it performs a consolidation of the level 1 backup and the most recent full save set.

▼ 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 Groups” on page 185](#).
2. Configure a Pool resource for consolidation backups. See [“To Create Pools” on page 145](#).
3. Select the name of the Group resource you created in step 1 as the value for the 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 illustrates the concept of how the Solstice Backup 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 consolidated level “c” automatically performs a level 1 backup, which backs up all the data that has changed since the last full back up. As part of its consolidation process, the Solstice Backup server merges this level 1 backup with the last full back up (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.

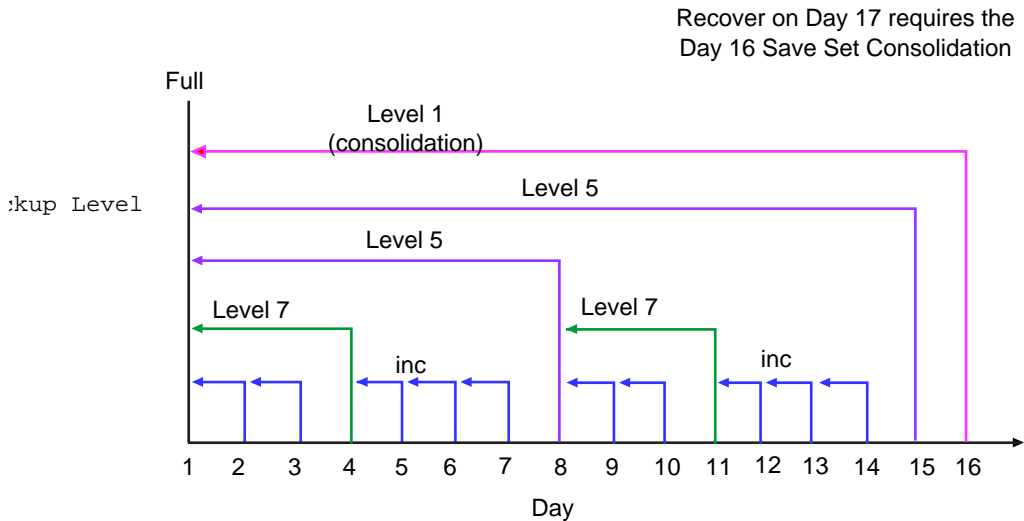


FIGURE 6-8 Consolidated Backup

Data Backup and Recovery

This chapter describes the method for backing up your 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 the location of your choice.

Manual Data Backups

The Solstice Backup server backs up your client data regularly using scheduled backups. Scheduled backups are the preferred backup method because they occur automatically, and also because scheduled backups can be recovered more easily.

Manual backups are intended for quick backups of a few files initiated by a user. They do not generate bootstrap files or back up indexes, so they cannot replace scheduled backups. Remember that you can manually start a scheduled backup group at any time by selecting the Start Group command. See ["" on page 181](#) for a complete description of scheduled backups and groups.

There may be times, however, when you want to back up one or more data items to the Solstice Backup server right away, instead of waiting for the scheduled nightly backup. For example, if you have just finished working on a spreadsheet, you might like to back the file up right away.

▼ To Manually Back Up Data

The `nwbackup` program provides a graphical interface through which you can perform manual backups. To view a display of the `nwbackup` program, see ["nwbackup" on page 34](#).

The `nwbackup` program is used to perform a manual backup initiated from the client rather than from the server's scheduled group backup (`savegrp`) that includes the client.

To start a manual backup:

1. Start the `nwbackup` program by entering the following command at the prompt:

```
# nwbackup -s server
```

Use the `-s server` option to specify a particular Solstice Backup server on the network.

The `nwbackup` program displays a representation of the client's filesystem.

2. Mark each directory and/or file for manual backup by selecting it and performing one of the following tasks:
 - Selecting the item and clicking the Mark toolbar button,
 - Clicking the “mark” box next to each item listed,
 - Selecting Mark>Mark.

When you mark an item for backup, a checkmark appears next to that item.

You can unmark a marked item in a similar way by selecting the item and performing one of the following tasks:

- Selecting it and clicking the Unmark toolbar button.
- Clicking the checked “mark” box next to the item to remove the checkmark.
- Selecting Mark>Unmark.

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. If you want to exclude certain files from manual backup, you must deselect those items.

3. Select File>Start Backup or click the Start toolbar button to begin the manual backup.

If you do not see the directories and files you requested for backup appearing in the Manual Backup Status dialog box, or if you receive a message such as the following:

```
Waiting for Server
```

and there is no apparent backup activity, the Solstice Backup server might not recognize your client, or there might not be a volume mounted in the storage device. You can end the backup by clicking Cancel.

When the Solstice Backup server has successfully finished the backup, the `nwbackup` program displays the following (or a similar) message:

```
Backup completion time: 2-15-97 3:27p
```

Caution – Sometimes there are bad files or errors on computer disk volumes that the Solstice Backup server cannot detect during a backup. It simply backs up the bad data. It is a good idea to run diagnostic programs regularly on your computers to correct disk volume errors.

Data Recovery

Recovery of backed-up data is performed using the `nwrecover` program. This program is similar to the `nwbackup` program discussed in the previous section. To view a display of the `nwbackup` program, see [“nwbackup” on page 34](#).

Every time the Solstice Backup server backs up a file to media (for example, magnetic tape), it creates an entry for the file in an index. The Solstice Backup server uses this index to map the backed-up files to volumes. If your computer has not been backed up yet, you will see an error message when you try to recover data.

Using `nwrecover`, you can browse backed-up files and filesystems from Solstice Backup clients and recover files as they were at a specific point in time when they were backed up. The versions of files and filesystems available in the `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 104](#).

If remote access is enabled, users may also be able to recover files from other client computers. For more information on remote access, see [“Remote Access Rights for Other Clients” on page 94](#).

When recovering data onto a filesystem where other versions of the same data already reside, there is the likelihood of a file conflict. You can determine what course of action to take each and every time a conflict occurs during a recovery, or

you can instruct the Solstice Backup server in advance what action to take for all conflicts, You can also simply redirect the recovered data to a different directory where there is no danger of conflicts.

▼ To Recover Data

To recover data:

1. Start the `nwrecover` program by entering the following command at the prompt:

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

Use the `-s server` option to specify a particular Solstice Backup server on the network. Use the `-c client` option to specify a particular Solstice Backup client on the network.

The `nwrecover` program displays a representation of the client's filesystem.

2. Determine how you want the Solstice Backup server to handle one of these conflicts:
 - Cases where there is a conflict between a recovered file and one that already exists on the filesystem of the client receiving the recovered data.
 - If you want the recovered data to be relocated to a different directory than the one it originally existed in at the time of the backup.

See [“To Relocate Recovered Data” on page 218](#) for more information on resolving recovered data conflicts.
3. Mark each directory and/or file for manual backup by selecting it and performing one of the following tasks:

- Selecting the item and clicking the Mark toolbar button,
- Clicking the “mark” box next to each item listed,
- Selecting Mark>Mark.

When you mark an item for backup, a checkmark appears next to that item.

You can unmark a marked item in a similar way by selecting the item and performing one of the following tasks:

- Selecting it and clicking the Unmark toolbar button.
- Clicking the checked “mark” box next to the item to remove the checkmark.

- Selecting Mark>Unmark.

You can designate the entire filesystem, specific directories, or individual files for recovery. Note that if you select a directory for backup, all of its nested subdirectories and files are also selected for recovery. If you want to exclude certain files from recovery, you must deselect those items.

Note: if you want to recover data from a scheduled backup other than the most recent one, then you must change the browse time. See [“To Recover Data from a Previous Backup Time” on page 217](#).

4. Select File>Start Backup or click the Start toolbar button to begin the recovery.

It takes the Solstice Backup server a few moments to recover your file, depending on file size, network traffic, server load, and tape positioning. During this time, messages appear for each recovered file in the `nwrecover` program, so you can monitor the progress of the recovery.

If there is a naming conflict with the file you are recovering, the Solstice Backup server prompts you to choose one of several workarounds. See [“To Resolve Recovered Data Conflicts” on page 219](#) for more information about naming conflicts.

If the recovery is successful, the `nwrecover` program will display the following (or a similar) message:

```
Successfully received 1 file

Recovery completion time: 5-30-97 2:03p
```

▼ To Recover Data from a Previous Backup Time

Using the `nwrecover` program, you can browse the entries for each backed-up file listed in the client file index created at any time you specify. Changing the browse time allows you view the entries for your backed-up files for a particular period. You might want to do this if you need to:

- Look for an item that you know you had previously, but is now gone.
- Retrieve an earlier version of an existing item.

Your browse policy determines how long older data is retained and thus, your ability to browse and recover this data. For example, if you would like to recover a lost file that you backed up six months ago, your browse policy must support it. See [“Browse and Retention Policies” on page 104](#) for more information.

To recover data from a previous backup time:

1. Select Change>Browse Time to open the Change Browse Time window.

The Change Browse Time window displays a monthly calendar from which you can select the client's backup. Click Previous or Next for earlier or later months.

2. Select the month and day that corresponds to the client data you want to recover.
3. Edit the time and date listed in the Browse Time attribute if you want a browse a time other than the one displayed.
4. Click OK.

The Change Browse Time window disappears, the `nwrecover` program now displays your client's file system as it appeared on the day and date you selected. The `nwrecover` program also displays that date in the Browse Time attribute.

5. Select the items you want to recover as described in ["To Recover Data" on page 216](#).

▼ To Relocate Recovered Data

When you recover data, the Solstice Backup server copies the data from the volume to a location on your computer. Unless you relocate the recovered data, the server tries to place the files in the original directory created when the data was last backed up. If an existing data item on your local disk has the same name as the one being recovered, the server will prompt you for instructions on how to recover the data item (rename, overwrite, or discard).

You can instruct the Solstice Backup server where to relocate the recovered data item in advance of the recovery using the Relocate feature while in the `nwrecover` program. Enter the new recovery destination path in the Relocate To attribute of the Recover Options dialog box.

To relocate recovered data:

1. Select Options>Relocate to open the Relocate window.
2. Enter the new directory location in the Relocate To attribute and click OK.

You must provide the full pathname of the relocation directory.

Perform a data recovery as described in ["To Recover Data" on page 216](#). The recovered data will automatically be relocated to the directory you entered.

▼ To Resolve Recovered Data Conflicts

By default, when the Solstice Backup server recovers data it will attempt to return them to their original location, using their 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 will occur. You can instruct the Solstice Backup server to respond to these recovery conflicts in several different ways, either on an item-by-item basis, or automatically.

You are prompted to make conflict resolution decisions each time to start a recovery using the `nwrecover` program. When you start the recovery, the Conflict Resolution window appears and asks you if you want to be prompted when a conflict occurs.

If you select *Yes* in the Conflict Resolution window, then each time there is a conflict, the Solstice Backup server will prompt you to select a resolution options.

If you select *No* in the Conflict Resolution window, then you must enter an automated conflict resolution option that the server will follow each time there is a conflict during the recovery.

[TABLE 7-1](#) lists the conflict resolution options available during a recovery.

TABLE 7-1 Recovery Options

| Option | Action |
|-----------------------------|--|
| Rename the Recovered File | Renames the recovered files by inventing a unique name that closely resembles the original name. Typically, the new name is the original name prefaced with a tilde (~). |
| Discard the Recover File | Discards the recovered data in order to prevent Solstice Backup from recovering the file if there is another file with the same name on your filesystem. |
| Overwrite the Existing File | Replaces the data on your filesystem with the recovered version. |

You can also choose to relocate your recovered data to a different directory where you know there will be no conflicts, or create a new directory using the Relocate feature. However, even if you relocate your recovered data to a new and/or empty directory, the Solstice Backup server still prompts you to enter your conflict resolution information. See [“To Relocate Recovered Data” on page 218](#) for more information on this feature.

▼ To Search for a Recoverable File

You can search for a file and directory you want to recover, but cannot locate in the `nwrecover` program, by using the Search feature. The Search feature allows you to enter the name of the item you are seeking, as well as the directory where the search will begin. This is handy if you want to speed the search by limiting the locations to be searched.

To search for a recoverable file:

1. Select View>Search to open the Search window.
2. Enter the name or a partial name of the item you are seeking.
3. Enter a path to designate where the search will begin. The default value is root (/) which will instruct the Solstice Backup server to search the entire filesystem.

▼ To View Details of Recoverable Files

You can display the contents and other details of directories listed in the left pane of the `nwrecover` window in the right pane using the File Details feature.

When the File Details feature is enabled, the `nwrecover` program displays detailed information regarding the files and subdirectories of the selected directory, including the date and time the file was last backed up, the permission of the file, and the date created.

To view details of recoverable files:

1. Select the item whose content's details you want to view.
2. Select View>File Details to display this information.

▼ To View the Version Details of Backed-Up Data

You can view the back-up details of a directory or file you select in the `nwrecover` program using the Version feature. The Versions feature provides a listing of each time the Solstice Backup server backed up that file, including the day, date, time, and size.

To view the version details of backed-up data:

1. Select the item whose backup version you want to view.
2. Select View>Versions to open the Versions window.

▼ To View Volumes Required for Data Recovery

You can view a list of the volume(s) needed to recover files and directories you have selected for recovery in the `nwrecover` program using the Volumes feature. The Volumes feature displays a windows providing the name of all volumes the Solstice Backup server needs to access in order to recovery the data you selected.

To view a list of the volume(s) required for data recovery:

1. Select the item whose volume information requirements you want to view.
2. Select View>Volumes to open the Recover Volumes Required window.

Disaster Recovery

Recovering data from Solstice Backup 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. Refer to the *Legato NetWorker Disaster Recovery Guide* for more information. Also, refer to the *Release Supplement* for the most recent disaster recovery information.

Note: You should keep a printed copy of these documents in a safe, convenient place away from the server.

Directed Recovery

Directed recovery allows you, from an administering Solstice Backup client at a remote location, to direct a Solstice Backup server to perform both of the following:

- Retrieve backed-up data from one of its Solstice Backup clients (a source client)
- Recover that data to a different Solstice Backup client (a destination client) within your enterprise.

The destination client receiving the recovered data can be the source client itself or any other Solstice Backup client within the enterprise. Because you do not have to be physically present at either the source client, the destination client, or the Solstice Backup server itself, you have a high degree of flexibility and convenience in recovering and redirecting data. Directed recovery includes all of the Solstice Backup server's standard file handling features.

You can use Solstice Backup software to administer directed recoveries between UNIX, NT, and Windows 95/98 Solstice Backup clients by using the command line recover program. The `nwrecover` program does not support directed recovery tasks.

Uses of Directed Recovery

Directed recovery is very useful when data recovery is required for a shared server or other computer where users cannot recover the data themselves. If you have a Solstice Backup client installed and the appropriate access rights, the directed recovery feature allows you to perform the recovery from your own computer remotely rather than from the source computer itself.

Directed recovery also permits the recovered data to be directed to other Solstice Backup client computers besides the requesting client computer. This allows you to share files and information with a broad range of Solstice Backup client computers.

Access Requirements for Directed Recoveries of UNIX Clients

Directed recovery is an administrator-level function and is typically not accessible for users and other nonadministrators. Performing directed recoveries requires you to have administration access for both the Solstice Backup server storing the data from the source client, and all destination clients receiving the recovered data. If you anticipate performing directed recoveries on a regular basis, you should establish the required permissions for each Solstice Backup server and client in your enterprise prior to receiving data recovery requests.

Access requirements to perform directed recoveries of UNIX-based Solstice Backup clients are as follows:

- All Solstice Backup clients and servers involved in the directed recovery must have the 6.x release of the software.
- Your username must be included in the administrator's list attribute of the Solstice Backup server that backs up the source client's data, for example, *username@administration_client*.
- The hostname of the Solstice Backup client with which you are administering the directed recovery must be listed in the servers file of the following:
 - The Solstice Backup server computer
 - Each Solstice Backup client computer designated as being a potential destination client.

- The user accounts on the Solstice Backup client administering the directed recovery must be included in the Remote Access attribute of each source Client resource of the Solstice Backup server. These are the accounts used on the Solstice Backup client that administrate the directed recovery, and the accounts owning the Solstice Backup client daemon (nsrexecd) on each destination client computer (root). For example, *user@administering_client_hostname* and *root@destination_client_hostname*.

The following sections describe how to set the various server and client administration rights required to enable directed recovery.

▼ To Grant Administrative Rights for a Solstice Backup Server

If you are using the Solstice Backup server to administer the directed recovery, your username needs to be listed as an administrator in that server.

To grant administrative rights for a Solstice Backup server:

1. Log in as *root* on the Solstice Backup server.
2. Start the Solstice Backup Administrator program.
3. Select Server>Server Setup to open the Server window.
4. Enter *user@hostname* in the Administrator attribute and click Apply.

▼ To Grant Administrative Rights to a Remote Solstice Backup Server's Client

If you are administering directed recoveries remotely rather than from the Solstice Backup server, then your user account on the remote computer needs to be listed as an administrator in the Solstice Backup server. Your remote computer must also be included in the *servers* file of that Solstice Backup server to allow your remote client computer to open a network connection to the Solstice Backup server.

To grant administrative rights to a remote Solstice Backup server:

1. Log in as *root* on the Solstice Backup server.
2. Start the Solstice Backup Administrator program.
3. Select Server>Server Setup to open the Server window.
4. Enter your *user@administering_client_hostname* in the Administrator attribute and click Apply.

Next, you must add the administering Solstice Backup client computer name to the *servers* file of that Solstice Backup server.

1. Log in as *root* on the Solstice Backup server.

2. Open the `servers` file, located in the `/nsr/res` directory, using your favorite text editor.
3. Enter the administering Solstice Backup client computer's name in the `servers` file using the following format style:

```
mercury.sun.com
```

4. Save and close the `servers` file.
5. Enter the `nsr_shutdown` command at the shell prompt to shut down the Solstice Backup daemons. A list of Solstice Backup daemons that will be shut down appears, along with a prompt that asks whether you want to continue with the `nsr_shutdown` command as follows:

```
mars# nsr_shutdown

nsr_shutdown will kill the following processes

 25768 ?      S  0:02 /usr/sbin/nsr/nsrexecd
 25770 ?      S  0:01 /usr/sbin/nsr/nsrexecd
 25771 ?      S  0:02 /usr/sbin/nsr/nsrd
 25783 ?      S  0:00 /usr/sbin/nsr/nsrmmdbd
 25784 ?      S  0:00 /usr/sbin/nsr/nsrindexd
 25785 ?      S  0:00 /usr/sbin/nsr/nsrmmmd -n 1

Do you want to continue? [Yes]? y
```

6. Start the Solstice Backup daemons by entering the following at the command prompt:

```
# /etc/init.d/networker start
```

▼ To Grant a Network Connection to a Destination Client

The administering Solstice Backup client computer must have access to the destination client in order to recover files onto it. You must add the administering Solstice Backup client computer name to the *servers* file of that destination computer.

To gain remote recovery permission to a destination client:

1. Login as *root* on the destination client computer.
2. Open the *servers* file, located in the */nsr/res* directory, using your favorite text editor.
3. Enter the administering Solstice Backup client computer's name in the *servers* file using the following format style:

```
mercury.sun.com
```

4. Save and close the *servers* file.
5. Shut down the Solstice Backup daemons by entering the following command at the prompt:

```
# nsr_shutdown
```

A list of Solstice Backup daemons that are shutting down appears. You are also prompted as to whether you want to continue with the shutdown. Confirm the shutdown in each instance.

6. Start the Solstice Backup daemons by entering the following at the command prompt:

```
# /etc/init.d/networker start
```

▼ To Gain Remote Access to a Solstice Backup Client

The administering client and the destination client must have remote administration access privileges for any source Solstice Backup client involved in a directed recovery. This access allows you to browse the source client file index and retrieve its files. You can add your *hostname* or *user@hostname* to the remote access list of any Solstice Backup client through the Solstice Backup Administrator program. Note that if you set the default client of a Solstice Backup server to include your hostname in its remote access attribute, all subsequently-added clients will also list you as having remote access.

To gain remote access to a Solstice Backup client:

1. Log in to the Solstice Backup server which includes the Solstice Backup client to which you want remote access.
2. Start the Solstice Backup Administrator program.
3. Select Clients>Client Setup to open the Clients window.
4. Select the client to which you want remote access in the Clients attribute.
5. Enter the userid and `user@administering_client_hostname` of the administering client in the Remote Access attribute.
6. Enter the Solstice Backup client daemon account in the Remote Access attribute (`root@destination_client_hostname`) for every destination client you need or anticipate needing to access this source client resource.

Directed Recovery Between Different Client Platforms and Types

You can perform directed recoveries of data between different combinations of Windows NT, Windows 95/98, NetWare, and UNIX source and destination clients. However, you must have established the proper remote access requirements for both clients and the Solstice Backup server that back ups up their data.

You must use the **recover** program at the command prompt to perform directed recoveries. If you want to perform directed recoveries using a graphical interface, you can use the Solstice Backup for Windows product's User program if it is available in your enterprise.

There are also certain restrictions on what platforms and types of Solstice Backup source and destination clients can be used for directed recoveries. The following sections specify the different client combinations possible. You should be familiar with the procedures for your specific client combination before performing a directed recovery.

Directed Recovery Between Windows NT Clients

You can perform directed recoveries of data between any source Solstice Backup for Windows NT client and any destination Solstice Backup for Windows NT client in your enterprise.

Directed Recovery Between Windows 95/98 and Windows NT Clients

You can perform only directed recoveries from a source Windows 95/98 client to a destination Windows NT client.

Directed Recovery Between Windows NT and NetWare Clients

You can perform only directed recoveries from a source NetWare client to either a NetWare or a Windows NT destination client.

Directed Recovery Between UNIX Clients

You can perform only directed recoveries from a source UNIX client to a destination UNIX client.

▼ To Perform a Directed Recovery Between UNIX Clients

You perform directed recoveries between a UNIX source and a UNIX destination client using the `recover` command and various options at the command prompt. You cannot back up data from UNIX-based clients and recover them to Windows NT or other non-UNIX clients, or the reverse. The source and destination client must be of the same platform type.

Also remember that a directed recovery is an interactive process only. Once you invoke the recovery, you still need to interactively select and recover files you choose from the source client index.

recover Command in Directed Recoveries

The `recover` command has a variety of options you can add to modify the basic browse and recover functionality. For performing directed recoveries, here are the following options:

- `-d` The name of the directory where you wish to relocate the recovered files.
- `-s` The name of the source client's backup server.
- `-c` The name of the source client.
- `-R` The name of the destination client.

- **-i (Y,N,R)** The recover option specifying what the Solstice Backup server should do if recovered files conflict with existing files of the same name. Here is an explanation of each option (these are not case sensitive):
 - **Y** Overwrite the existing file when there is a naming conflict.
 - **N** Skip restoring this file when there is a naming conflict.
 - **R** Rename this file when there is a naming conflict. The Solstice Backup server automatically appends each conflicted filename with a “.R”.
- *[directory_name]* The name of the initial directory in which to begin browsing.

For a complete listing of the `recover` command and all its options, see `recover` in the *Legato Command Reference Guide*.

Example: Directed Recovery Overwriting Existing Files

If the source client’s backup server was named *venus*, the source client itself was *mars*, the destination client was *jupiter*, and you wanted the recovered files to overwrite any existing files with the same name, you would 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 specify the items you want to mark for recovery using a set of procedures to mark files and directories for recovery.

To perform a directed recovery of files and directories you specified:

1. Enter the command to perform the directed recovery, using the following syntax style:

```
# recover -s server -c source_client -R destination_client -irecover_option  
[directory_name]
```

2. Navigate to the directory containing file or directory you want to “mark” for directed recovery (if you did not specify the initial directory to browse by including a *directory_name*). For example:

```
# cd /var/adm
```

3. Mark the item for directed recovery by entering the following command at the prompt:

```
# add file_name
```

4. Repeat this procedure for each item you want to mark for directed recovery.
5. View the items you have marked for recovery by entering the following command at the prompt:

```
# list
```

6. Start the directed recovery by entering the following command at the prompt:

```
# 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 level full backup first, then recover the other backups in level order from 1 – 9, and then recover the incremental backups.

How Save Set Recovery Works

Save set recovery allows the recovery of data from a save set rather than the regular method of selecting items from the Recover browse window. The data recovered can be individual files, directories, or the entire save set itself. When entire save sets are recovered—for example, an entire disk volume or partition on a disk—the recovery

process is accelerated because the Solstice Backup server does not have to search for and locate each individual file in the index. Instead, the server locates the save set containing all of the files needed for recovery at once. The names of each save set for the client and the number of versions of each save set is displayed in the Save Sets dialog box.

You might want to perform a save set recovery under the following circumstances:

- The files are no longer browsable in the Recover dialog box because their browse periods have expired or their entries have been removed with the **Remove Volume** command.
- A disk crash during the backup produced an incomplete backup, making it impossible to recover the files using the `recover` command.
- Using 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 Solstice Backup file handling features.

Note: To perform a save set recovery, you must be listed as an administrator in the Solstice Backup server's resource, or your *hostname* or *user@hostname* must be entered in the Solstice Backup client resource's Remote Access attribute of the client that is providing the backed-up data.

▼ To Recover Data from a Save Set

The information in this section explains how to recover data from a save set.

1. Start the Solstice Backup Administrator program.
2. Select Save Set>Recover to open the Save Set Recover window.
3. Enter or select the client whose save set data you want to recover in the Client attribute.

The save sets available for recovery are displayed in the Save Set attribute.

4. Select the specific save set(s) whose data you want to recover in the Save Set attribute. You can select as many save sets as are listed in this attribute.

The selected save set's backup instances appear in the Instances attribute.

5. Select the particular instance of a save sets backups in the Instances attribute.

6. Click Details if you want to view the details of the save set instance you selected, such as the backup date, number of backed-up files, save set ID number, and pool.
7. Click Recover to open the Save Set Recover Status window. Here you can customize the save set recovery in a variety of ways.
8. Perform any or all the tasks associated with the features in the Save Set Recover Status window:
 - Specify what save set data is to be recovered (see [“To Specify What Save Set Data Is Recovered” on page 231](#)).
 - View the volumes required to recover the specified save set data (see [“To View Volumes Required for Save Set Data Recovery” on page 232](#)).
 - Relocate recovered save set data to a non-default location (see [“To Relocate Recovered Save Set Data” on page 232](#)).
 - Resolve recovered save set data conflicts (see [“To Resolve Recovered Save Set Data Conflicts” on page 233](#)).
9. Click Start to begin the save set data recovery. The recovery status is displayed in the Status attribute of the Save Set Recover Status window.

▼ To Specify What Save Set Data Is Recovered

You can specify what save set data is recovered while in the Save Set Recover Status window. If you want to recover the entire save set, do not change the default value displayed in the Paths to Recover attribute. When you initiate the save set recovery, the entire save set is recovered.

If you want to limit the recovered save set data to specific files and/or directories, you can specify that information in the Path to Recover attribute. Enter the full path of the first item you want to recover and press [Return].

The path you entered is then displayed in the list of save set recovery paths below the Path to Recover attribute. Repeat this process for each separate item you want to recover from the save set. Once you specify individual items and their paths, the complete save set will no longer be recovered, only the particular items you entered.

When specifying directory paths for recovery, place a forward slash after a specified directory. For example, enter:

```
/usr/paul/test/
```

instead of

```
/usr/paul/test
```

This means that only a directory `test` will be restored and not `test1`, `test2`, and so forth. Save set recover will parse a directory name and restore all instances that match the leading string, unless limited by a slash and assuming the overwrite option in cases of a recovered data conflict.

Once you have specified what save set data you want recovered, click Start to begin the recovery, or use any of the other features present in the Save Set Recover Status window.

▼ To View Volumes Required for Save Set Data Recovery

You can view a list of the required volume(s) necessary to recover save set data while in the Save Set Recover Status window.

To open the Save Set Recover Volumes Required window, click Volumes. This window displays each volume that will need to be available to your backup device to complete the save set recovery.

If no volumes are displayed, that is indicative 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 Solstice Backup server. See [“Recovering Client File Indexes” on page 238](#) for more information.

▼ To Relocate Recovered Save Set Data

You can instruct the Solstice Backup server to relocate recovered save set data from the data’s original location during its backup to another place on your filesystem while in the Save Set Recover window.

To relocate recovered save set data:

1. Click Options to open the Save Set Recover Options window.
2. Enter a path in the Relocate Recovered Data To attribute.

For more information on relocated recovered data, see [“To Relocate Recovered Data” on page 218](#).

▼ To Resolve Recovered Save Set Data Conflicts

You can instruct the Solstice Backup server to perform a particular action in the case of a conflict between an item recovered during the save set recovery, and a like-named file already existing on the filesystem while in the Save Set Recover Status window.

To resolve recovered save set data conflicts:

1. Click Options to open the Save Set Recover Options window.
2. Select a choice in the Duplicate File Resolution attribute.

For more information on recovered data conflicts and their resolution, see [“To Resolve Recovered Data Conflicts” on page 219](#).

▼ To Restore a Save Set Entry in the Online Indexes

The following sections describe the process of restoring save set entries into one or both of the online indexes.

Client File Index Only

If the file is not browsable (which means that the save set’s browse policy has expired), but its save set is still tracked by the Solstice Backup 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. Run the `mminfo` command by entering the following command:

```
# mminfo -a -v volume_name
```

2. From the `mminfo` output, find the `ssid` that you believe contains the file you want. Make sure it is not the bootstrap `ssid`.
3. Restore the save set entry in the client file index using the `nsrmm` program by entering the following command:

```
# nsrmm -s ssid -t retention_time
```

Client File Index and Media Database

If the save set that contains the file is not browsable and the save set is not represented 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. Run the **scanner** program on the backup volume that you believe contains the file you want by entering the following command:

```
# scanner device_name
```

Make a guess based on the adhesive label on the volume.

See [“To Find a Volume Name” on page 176](#) for other suggestions on finding volume names.

2. Use the output from the **scanner** program to decide whether to reintroduce the contents of this volume into the client file indexes and whether the save set you want to rebuild is on this volume. (You must locate *all* the volumes that contain this save set ID.)
3. After you have determined which volumes to reintroduce into the online indexes, run the **scanner** command:

```
# scanner -i device-name
```

The **scanner** command 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*.

4. Use the **nwrecover** program to browse the file index for the file you want to recover.

▼ To Recover an Entire Save Set to the Solstice Backup Server

To recover an entire save set directly to the Solstice Backup server's filesystem, use the **scanner** program by entering the following command at the prompt:

```
# scanner -S save_set_id device_name | uasm -rv
```


This command reads all the information associated with the *ssid* from the volume and places a copy of this data on the Solstice Backup 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 Solstice Backup server's filesystem.

If you want to be sure this action is correct 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.

You could also use **rsh** (or its equivalent) in conjunction with the following command to recover the save set to the client, if the save set originated on a Solstice Backup client instead of the Solstice Backup server:

```
# scanner -S ssid device_name | rsh client (cd destdir; /path/to/uasm -rv)
```

▼ To Recover One File Directly from a Volume

To recover a single file from a volume, run one of the following commands:

```
# scanner -S save_set_id device_name | uasm -rv filename
```

or

```
# scanner -S save_set_id device_name | 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 and contrasts the various data recovery methods provided by the Solstice Backup product.

Advantages of Index-Based Recoveries

Standard recovery and directed recovery are both based on using the client file index to locate and recover the desired data.

The following key benefits are associated with index-based recoveries:

- Directed recovery is well-suited for centralized administration. A large network can be administered for direct recovery from a single administering computer. Data recovered to an administering server can then be forwarded to a shared server drive and made available to a wide number of users.
- Directed recovery is useful if a computer is inoperable, a hard disk has crashed, or a network is no longer recognizing 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 also convenient if you want to transfer files between two clients. For example, if one client's system 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 Solstice Backup server recovers data to the data's original client. Thus, both kinds of client file index-based recoveries can be contrasted with a save set recovery.

You might prefer index-based recovery over save set recovery for the following reasons:

- Because users can browse the index to mark data for recovery, neither users nor administrators need to know the exact filename, location, or date of the specific file(s) they want to recover.

The Solstice Backup 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 are relieved of the burden of remembering explicit filenames from months past.

- Users can query the client file index with a target time for which they want to rebuild the file hierarchy. The Solstice Backup server constructs for the user a view of the filesystem that was written closest to the browse target time without exceeding it.
- The Solstice Backup server automatically re-creates the data exactly as it existed at the time that the backup occurred.

For example, if a user requests recovery of a backup version that is a level 5 backup, the Solstice Backup 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.

Disadvantage of Index-Based Recoveries

The disadvantages of index-based recovery are situation-specific:

- If the data a user wants to recover has passed its browse policy period, its entries have already been removed from the index. Entries for the files can be recovered and restored to the client file index using the `scanner` command, but this can be time-consuming or inconvenient. See [“Recovering Client File Indexes” on page 238](#) for more information on `scanner`.
- An index-based recovery takes time to mark each file for recovery.
- An index-based recovery requires memory proportional to the number of files marked for recovery.

For all these reasons, if recovery is planned for a significant number of files, you might prefer to use a save set recovery.

Costs of Save Set Recoveries

Three costs are associated with save set recovery:

- 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.
- Data item renaming issues are not handled automatically as they are with client index-based recovery.

A save set recovery follows the basic steps of an index-based recovery, but without the advantage of client file index handling that the Solstice Backup server provides automatically with index-based recovery.

Caution – For this reason, 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. Solstice Backup file handling is not available for save set recovery, so file X remains on your Solstice Backup 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 force the `nwrecover` program 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.

Deciding Between Index-Based and Save Set Recovery

This section summarizes the situations in which it is better to use client file index-based recovery than to use save set recovery.

Client file index-based recovery is most appropriate under the following conditions:

- You cannot determine which save sets to recover.
- You need to recover files from many backups to restore the filesystem to an acceptable level.
- You are recovering the Solstice Backup indexes.
- Recovering extra files is not acceptable.

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 Client File Indexes

The following sections explain different procedures for recovering the online index.

▼ To Recover a Client's Current File Index

Unlike regular backed-up data, recoveries of client file index information is performed using the `nsrck` program with the `-L7` option. If you do not specify any time or date information, the current client file index is recovered.

To recover the online index of a client, enter the following command:

```
# nsrck -L7 client_name
```

▼ To Recover Client File Indexes for a Specific Time and Date

With the Solstice Backup release 6.0 and later software, it is possible to recover a client file index for a certain 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.

Client file index recovery using the Solstice Backup product is non-destructive, so when you recover the client file index as of a time and date in the past, `nsrck` will add the entire contents of the index as of that time and date, to the current contents of the index. Generally, this means your client file index will become much larger.

The benefit of recovering the client file index as of a specific time and date is that you can browse save sets that had passed their browse policy but were still recoverable. The save sets referred to by the recovered client file index will be marked as browsable. They will remain browsable for the length of time they were originally browsable.

To recover a client index for a specific date, enter the following command:

```
# nsrck -t "time mm/dd/yyyy" -L7 client_name
```

Example: Recovering a Client File Index for a Specific Time and Date

In this example, you are recovering a client file index to 4 PM, March 2 1999:

```
# nsrck -t "04 pm 03/02/1999" -L7 client_name
```

Relocating Recovered Online Indexes

You can relocate recovered online indexes depending on the release of Solstice Backup software that performed the backups.

Solstice Backup Release 6.0 Software

You can recover online indexes to a different directory only if all the backups were made using the Solstice Backup release 6.0 software.

Solstice Backup Release 5.5 and Earlier Software

You cannot relocated recovered online indexes to a different directory if they contain any backups performed using the Solstice Backup release 5.0 or earlier software. Instead, you must recover them to the directory where they originally resided. After you recover the indexes to their original location, you can move them to another directory if necessary. If you attempt to recover the indexes to a different directory, you will receive a warning message stating that the recovery for the client index was unsuccessful.

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.

Mount the volume and execute the `scanner` command with any of its applicable options. The option that is most appropriate depends entirely upon the situation, your level of knowledge about the location of the data, and the nature of the recovery.

The following are the most commonly used `scanner` options:

- Rebuild the media database entries for everything on the storage volume (by executing `scanner -m device`) and then perform a save set recovery.
- Rebuild the media database and client file index entries for everything on the storage volume (by executing `scanner -i device`) and then perform either a save set recovery or a client file index-based recovery.
- Rebuild the media database entries for a specific save set on the storage volume (by executing `scanner -s save_set_id -i device`) and then perform a save set recovery.

- Rebuild the media database and client file index entries for a specific save set on the storage volume (by executing `scanner -s save_set_id -m device`) and then perform either a save set recovery or a client file index-based recovery.
- Execute `scanner` with the `-c client_name` and `-N save_set_name` options, along with any other of the options previously described to further specify the save sets by client or name.
- Execute `scanner -s save_set_id` to send a specific save set to standard output.
- Use the `scanner` program's `-x` option to introduce further processing commands on a per-save set basis.

Detailed information about `scanner` program syntax and examples of how `scanner` options can be customized to meet the needs of a particular recovery situation are available in the *Legato Command Reference Guide* on the Documentation Suite CD-ROM.

The entries placed in the client file index receive the original browse and retention policies as of the time you scan in the save set. If a save set originally had a browse time of 1 month and retention time of 3 months, and that save set had expired, and you scanned that save set back in, it would remain browsable for 1 month and recoverable for 3 months after you scanned it in.

The entries placed in the media database also receive the original browse and retention policies as of the time you scan in the save set, and the volume mode is listed as “read-only.” See [“Storage Management Operations” on page 167](#) for more information on managing the client file index and media database entries.

If you upgrade a Solstice Backup server to 6.0 or later, the current browse and retention policies are applied to save sets that had been scanned into the client file index and the media database.

Data Archiving Methods

This chapter describes the features and functionality of the optional NetWorker Archive application.

Overview

The NetWorker Archive application provides archive services to Solstice Backup client computers for which you have enabled archiving. The archive process captures files or directories as they exist at a specific time, and writes the data to special archive storage volumes. Browse and retention policies do not apply to archived data. After the archive process completes, you can groom (delete) the original files from the disk to conserve space.

Use NetWorker Archive in addition to scheduled Solstice Backup backups to protect your data. Although backups provide short-term insurance against data loss due to hardware failure or user error, archives offer a long-term strategy for data preservation.

Some of the benefits of using the NetWorker Archive application include:

- Archived data is never subject to automatic recycling, so archived data cannot be accidentally overwritten.
- Archived volume entries are maintained in the media database indefinitely.
- Archived files can be removed from primary disk storage to make space for newer files, and retained in archive volumes for quick retrieval in the future.

Archives Versus Backups

Archive save sets are similar to backup save sets. The principal difference is that archive save sets have no expiration date. Also, archives are always full; there are no levels [1 – 9] or incremental archives.

Browse policies do not apply to archive data. This ensures that archived data is never eligible for automatic recycling.

Requirements

Before you begin licensing and using the NetWorker Archive application, make sure that you have the following:

- Solstice Backup Network Edition or Solstice Backup Power Edition installed and enabled on the Solstice Backup server.
- A device, either stand-alone or in an autochanger or silo, connected to a Solstice Backup server or storage node. To clone your archives, you must have two devices available.
- Meet the necessary licensing requirements for the Archive application.

Licensing

The NetWorker Archive application is an optional module that you must purchase and license separately from other Solstice Backup software components. The Solstice Backup distribution files include the NetWorker Archive software. Although no additional software is required to use this feature with the Solstice Backup product, the functionality only works when the Archive application is properly licensed in the Solstice Backup server.

The same evaluation conditions and procedures are the same for the Archive application as for other Solstice Backup products. See [“Licensing” on page 39](#) for more information on licensing procedures.

Caution – You can only use the evaluation enabler code once per network, or it disables all the Solstice Backup servers (even for backup) that you enabled with it.

Archive Application Permissions

The following sections describe the various server- and client-based permissions required to use the Archive application.

Server Permissions

To use the Archive application, you must have the following permissions:

- You must be listed:
 - In the Administrator list of the Solstice Backup server performing the Archive application tasks to configure archive requests and other scheduled archive functions.
 - In the Archive Users attribute of the Client resource to request archiving for that particular client.
- Read permission is also required for each file that you want to archive.

Anyone can browse archive save sets to look at the information in the media database. However, you may only retrieve files that you own, unless you are the administrator, in which case you may retrieve anyone's files.

If you want to overwrite an archive tape, first make sure nobody will ever need the data again. Then simply relabel the volume, as you would a backup volume. The relabeled volume will still be a member of the archive pool unless you change the pool association when you relabel the volume.

Client Permissions

After you enter the archive enabler code for the Solstice Backup server, all clients of that server are enabled for the NetWorker Archive application by default.

- To disable or enable archiving for individual clients, use the Archive Services attribute in the Clients resource.
- To archive data that resides on the Solstice Backup server, make sure that the Archive Services attribute is enabled in the Clients resource for the server.

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 you have a Legato NetWorker Module for a database

application and the Solstice Backup client software is installed on the same computer that backs up to the same Solstice Backup server, both client resources have the same name. The Archive Services attribute is either enabled for both or disabled for both.

Restricting Archive Access

To restrict certain client access to the NetWorker Archive application, grant read permission to specific users in the Archive Users attribute of the Clients resource. Users listed in the Archive Users attribute can archive any file for which they have read permission, and they can browse the archive save sets (view the annotation information in the media database).

By default, only the owner of the archived files or the Solstice Backup administrator can retrieve the save set. You do not have to be the creator of the archived files.

Enabling Public Archive Access

To allow all enabled archive users to retrieve any archived file, enable the Public Archives attribute in the Server resource. Whether you enable or disable the Public Archives attribute, retrieved files retain their original file ownership and access permissions.

You cannot omit specific clients from the retrieval permissions when you enable the Public Archives attribute.

How the Solstice Backup Server Archives Data

Whether you initiate the archive on the client or the server, the archive is performed by the client's `nsrarchive` program, which is initiated by the client's `nsrexecd` service. You can use the graphical `nwarchive` program for manual archives from the client's command line. You can schedule archives from the server or client using the Archive Request resource in the Solstice Backup Administrator program.

FIGURE 8-1 illustrates how the Solstice Backup server archives data.

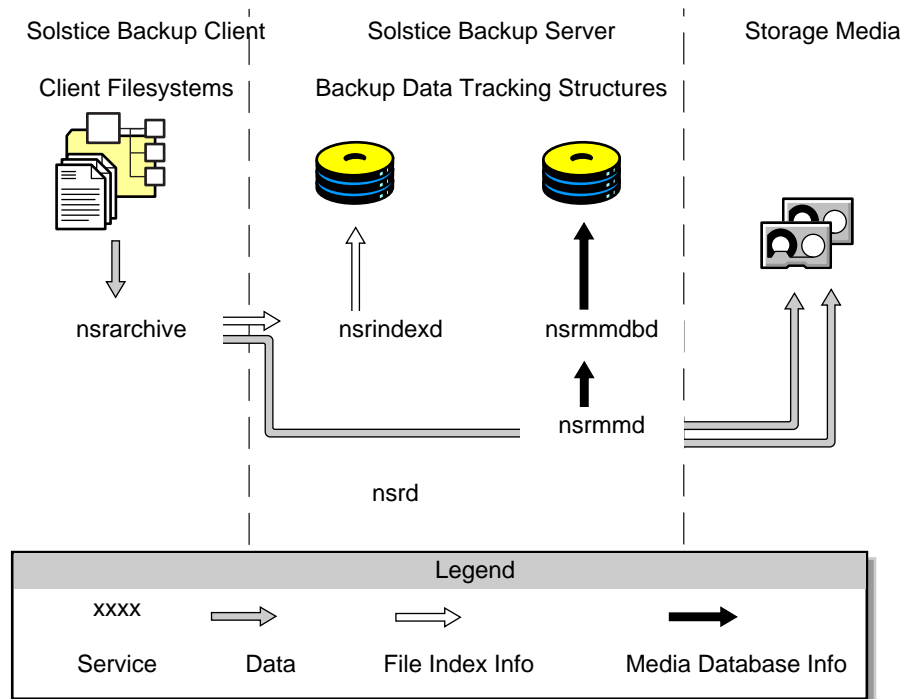


FIGURE 8-1 Overview of Archive Operation

Archiving of Non-Indexed Data

Archiving of UNIX-based data does not add entries into the client file index. Although you can select any combination of directories and files in a filesystem to be archived using the `nwarchive` program, that collection of data is archived as a single save set. Subsequent retrieval of this data can only be in the form of that single save set. For instance, if you wanted to recover a single file that was archived with other files in this manner, the entire save set would need to be recovered to get to the individual file. Because there is no client file index, this type of archive is known as *non-indexed* archive data.

When you manually archive UNIX data, or create an archive request for scheduled archiving of UNIX client data, the Solstice Backup server prompts you to provide an annotation that describes the save set. That annotation is the primary means of identifying archived UNIX client data when you want to retrieve that data.

Archive Pool

UNIX client archive data is by default written to the Archive pool. The archive volume assigned to the Archive pool can be in a device attached to the Solstice Backup server or a device attached to a storage node. Information about the archive data is tracked in the Solstice Backup server's media database. You cannot change the settings for the preconfigured Archive pool.

You cannot enable the Store Index Entries attribute in the Archive pool resource to store client file index entries for a non-indexed archive. If you select the option and apply your changes, you receive an error message.

Archiving of Indexed Data

PC-based Solstice Backup clients, such as those for Windows NT, can have entries added to their client file index as the data is archived as a save set. Subsequently, you can select individual files for retrieval, rather than the entire save set (as with non-indexed archive data). This type of archive is known as *indexed* data.

Note: Selecting individual items for retrieval requires the use of the Solstice Backup for Windows User program. Refer to your Solstice Backup for Windows products documentation for more information.

In order for PC clients to take advantage of indexed archiving, you must enable the Store Index Entries attribute of the client's Client resource. Certain client resource attribute values must be set to enable client file index entry additions during the archiving process, otherwise the data is archived as a single save set like in non-indexed client data.

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.

PC Archive Pool

Archived PC-based data is written to the PC Archive pool rather than the Archive pool, which writes non-indexed data only. This option is not supported on Solstice Backup UNIX clients. You cannot change the settings for the preconfigured PC Archive pool.

Using the PC Archive pool creates large index pools; as a result, you might run out of hard disk space. To conserve disk space, you can purge or compress the index.

Volume Tracking

The Solstice Backup server tracks the volumes used for archives separately from those used for backups. You cannot archive files to a backup volume, or back up files to an archive volume. This way, the backups are kept separate from the archives. An archive volume must be loaded and mounted in the server device to complete an archive.

The `nwarchive` Program

Archiving of backed-up client data is managed using the Solstice Backup `nwarchive` program. The `nwarchive` program provides archive services to Solstice Backup client computers for which you have enabled archiving. To view a display of the `nwarchive` program, see [“nwarchive” on page 36](#).

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

```
# nwarchive -s server
```

Use the `-s server` option to specify a particular Solstice Backup server on the network. If you enter `nwarchive` without the `-s server` option, then `nwarchive` connects to your default backup server.

Data Archiving Options

During both the manual and scheduled archiving processes, you are prompted to select one or more archiving options:

- Compress
- Clone
- Verify
- Groom

Compression

The Solstice Backup server compresses the data as it is archived, saving storage space. Compression and decompression of data typically takes more CPU power and may slow data archiving and retrieval.

Cloning

The Solstice Backup server writes a copy of each archive save set to a volume from an Archive Clone pool, similar to the process of cloning backup save sets.

Verification

The Solstice Backup server checks the integrity of the data on the storage volume against the original data on the client system. The server also checks the archive volume to ensure that it is writable and contains no bad spots.

Grooming

When an archive is complete, you are prompted to remove the files you just archived from your local disk in order to free up disk space. This procedure is known as *grooming*.

- To keep the archived files on your disk, select *No*.
- To confirm that you would like the Solstice Backup server to remove the archived files from your disk, select *Yes*.

If you elect to remove the files, the Solstice Backup server displays the Remove Archived File window. Here you can set the server to prompt you for each file that is removed, or allow the server to remove them automatically. You are prompted to confirm the removal of each file, thus giving you the opportunity to confirm or cancel the file deletion.

Your archived files are stored on special backup media labeled as an archive volume for long-term storage. If you choose to remove the old files from your disk, you will have more space on it. You can recover the old files from the archive volume in the future, if you should ever need them.

▼ To Create an Archive Client

Creating an archive client for use with the Archive application is no different from creating a regular client, except that certain archive attributes must be completed and enabled to indicate that the client is an archive client.

After you enable the Archive application, all existing Solstice Backup clients automatically become archive-enabled. If you add new clients after installing the Archive application, you need to enable their archive services.

To create an archive client:

1. Create a client resource for the client computer. (See [“To Create Clients” on page 90](#) for instructions.) Follow the steps below for setting specific client resource attributes necessary for archiving the client’s data.
2. Select *Enabled* in the Archive Services attribute.
3. Enter *Administrator* and any appropriate user names in Archive Users attribute to allow users on the client to perform manual archives.
4. Make the remaining configuration choices appropriate for an archive client. For more information on archive requests, see [“To Create an Archive Request” on page 255](#).

Enabling Archive Services for an Existing Client

You can update an existing client resource’s attributes to enable archive services at any time. See [“To Edit Clients” on page 92](#) to access the client’s attributes and edit them to enable archive services as specified in [“To Create an Archive Client” on page 250](#).

Manually Archiving Data

You can archive data at any time using the manual data archiving method. Manual archiving 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 backup, all of its nested subdirectories and files are also selected for backup. If you want to exclude certain files from manual archiving, you must deselect those items.

Remember that if the client whose data you are selecting for archiving is a UNIX client, there are no client file index entries being written during the archive. All the items you selected are archived as a single save set, and must be retrieved as a single save set, even if you want only a single item.

▼ To Manually Archive Client Data

To manually archive client data:

1. Start the `nwarchive` program as follows:

- If you want to archive to the default Solstice Backup server, enter the following command:

```
# nwarehouse
```

- If you want to archive to a Solstice Backup server other than your default server, enter the following command:

```
# nwarehouse -s server
```

The `nwarehouse` program displays a representation of the client's filesystem, similar to that of the `nwarehouse` program you use for manual backups.

2. Select each directory and/or file for manual archiving. From the pulldown menu, select Mark>Mark. You can also select the item and click the Mark button, or click the "mark" box next to each item listed. When you mark an item for archiving, a checkmark appears next to that item.

To deselect an item currently marked for backup, select Mark>Unmark. You can also select the item and click the Unmark toolbar button or click the checked "mark" box next to the item to remove the checkmark.

3. Select File>Start Archive or click the Start toolbar button to begin the manual archive.

The Archive Options window appears.

4. Enter a term or comment to help identify the contents of this archive in the Annotation attribute. Your annotation cannot exceed 1024 characters.

When archiving UNIX client data, all selected files are archived as a single save set. The annotation you enter is the primary means of locating that save set. Take care in determining this annotation.

5. Enter or select the appropriate pool to which the Solstice Backup server will write your archive data in the Archive Pool attribute. The default pool is *Archive*.

PC client data should be directed to the *PC Archive* pool.

6. Select one or more options for the archived data in the Operators attribute. See ["Data Archiving Options" on page 249](#) for information on these options.

7. Enter or select the pool to which the Solstice Backup server will write the clone of your archive data in the Clone Pool attribute. The default pool is Archive Clone pool.

8. Click Archive. The Archive Status window appears displaying the progress of the data archiving.

If you do not see the directories and files you requested for archive appearing in the Archive Status dialog box, or if you receive a message such as the following:

```
Waiting for Server
```

and there is no apparent backup activity, the Solstice Backup server might not recognize your client or there might not be a volume mounted in the storage device. You can end the backup by clicking Cancel. For more information about configuring clients, see ["" on page 89](#), and for more information about mounting volumes, see ["Solstice Backup Volume Operations" on page 171](#).

When the Solstice Backup server has successfully finished the backup, the `nwarchive` program displays the following (or a similar) message:

```
Archive completion time: 2-15-97 3:27p
```

Caution – There are certain kinds of bad files or errors on computer disk volumes that the Solstice Backup server cannot detect during an archive. It simply backs up the bad data. It is a good idea to run diagnostic programs regularly on your computers to correct disk volume errors.

Scheduled Archives

Scheduled archives are performed from the Solstice Backup Administrator program rather than the `nwarchive` program for manual archives. Archive requests are Solstice Backup resources that allow you to schedule the archiving of UNIX client data using the Archive Application. The Archive Request resource contains the attributes for scheduling and configuring a scheduled archive. You can create, edit, and delete archive requests similar to other Solstice Backup resources.

Unlike scheduled backups, scheduled archives run only once. The advantage of a scheduled archive is that you can run the archive at a time when network traffic is low and not tie up computer and network resources during business or heavy-use hours.

▼ To Set Client Archive Privileges

The Solstice Backup software allows a client to schedule its own archive requests. However, to schedule archive requests, the client must have administrator privileges. The term you use to designate administrator privileges when you create or edit clients varies depending on the type of client.

Enter the appropriate term in the client's Archive Users attribute of the Client resource:

- For Windows NT clients, enter "Administrators" (note upper case "A" and trailing "s")
- For UNIX clients, enter "root"
- For NetWare clients, enter "supervisor"

Key Archive Request Attributes

Many of the attributes located in the Archive Request dialog box are common to other Solstice Backup resources. However, some require additional explanation.

[TABLE 8-1](#) displays some of the key attributes for the Archive Request resource.

TABLE 8-1 Archive Request Attributes

| Attribute | Action |
|--------------------|---|
| Annotation | Entering an annotation is characteristic of archiving UNIX client data. Unlike Windows and other non-UNIX clients, archived UNIX client data is not recorded in the client file index. Archived UNIX client data is not browsable, and the annotation you enter is required to retrieve archived data. You can enter any single term or phrase in the Annotation attribute. |
| Save Set | An archive request archives data at the save set level. Your archive request can archive a single save set, or as many as you enter into the Save Set list box. Enter each save set and press [Return] before entering the next one. |
| Archive Pool | Because UNIX and PC desktop clients such as Windows NT or Windows 95 are archived differently, they must go to separate archive pools. UNIX client data should be directed to the Archive pool choice, while other clients should go to the PC Archive pool. You also have the option of creating your own customized archive pools and archive clone pools, prior to creating the Archive Request resource. These additional pools will appear as attribute options along with the preconfigured choices. For more information, see "To Create Archive Pools" on page 147 . |
| Archive Clone Pool | The same principals of the Archive Pool attribute apply to cloning pools. |

▼ To Create an Archive Request

To create an archive request:

1. Select **Customize Archive Request** to open the Archive Request window.
2. Click **Create**.
3. Enter a name for the archive request in the **Name** attribute. The name you enter cannot exceed 64 characters.
4. Enter a comment describing the archive request in the **Annotation** attribute. Your annotation cannot exceed 1024 characters.
5. Select one of the following options in the **Status** attribute:
 - *To have the archive starts as soon as you click Apply in the Archive Request window, select **Start Now** and proceed to [Step 7](#).*
 - *To have the archive starts at the time you specify in the Start Time attribute (see [Step 6](#)), select **Start Later**.*
6. If you selected *Start Later* in [Step 5](#), 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.
7. Enter the name of the archive client's hostname in the **Client** attribute. If you want the Solstice Backup server to archive its own files, enter the server's machine name.
8. Enter the pathname(s) of the save sets you want to archive in the **Save Set** attribute.

Note: If you want to create an archive request for all of the client's save sets, enter the value *All*. If you enter this value, be sure to select *none* in the **Grooming** attribute. If you select *remove* for the **Grooming** attribute, the Solstice Backup server will delete all of the client's archived save sets.
9. Enter or select a directive in the **Directive** attribute. For more information on directives, see ["Directives" on page 127](#).
10. Enter or select the appropriate pool in the **Archive Pool** attribute. The default pool is *Archive*.
 - For UNIX client data, archive to the *Archive* pool which stores the save set archive.
 - For PC client data, archive to the *PC Archive* pool where the client file index entries are archived in addition to the entire save set.
11. Select the appropriate response for the following data archiving options. For more information, see ["Data Archiving Options" on page 249](#).

- To instruct the Solstice Backup server to check the integrity of the data on the storage volume, select *Yes* for the *Verify* attribute.
- 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. The default pool is *Archive Clone Pool*.

For UNIX client data, archive to the *Archive Clone* pool which stores the cloned save set archive.

For PC client data, archive to the *PC Archive Clone* pool where the client file index entries in addition to the entire cloned save set archive.

- To instruct the Solstice Backup server to remove the archived files from your disk, select *Yes* for the *Grooming* attribute.

12. Enter a notification for the Solstice Backup server to take following the completion of the data archive in the *Archive Completion* attribute. For example, you can enter a notification to send mail that reports the status of the archive. See [“Notifications” on page 50](#) for more information on notifications.

13. Complete any other attributes as necessary and click *Apply*.

▼ To Edit an Archive Request

To edit an archive request:

1. Select *Customize>Archive Request* to open the *Archive Request* window.
2. Select the archive request from the list displayed in the *Archive Requests* attribute.
3. Edit the attributes of the archive request and click *Apply*.

You cannot edit an archive request that is currently in use.

▼ To Delete an Archive Request

To delete an archive request:

1. Select *Customize>Archive Request* to open the *Archive Request* window.
2. Select the archive request you want to delete from the list displayed in the *Archive Requests* attribute.

3. Click Delete.

When prompted to confirm that you really want to delete that archive request, click OK.

4. Click Apply.

Scheduled Archive Management

You can manage and monitor scheduled archive requests in a manner similar to that of scheduled group backups. This section describes how to accomplish the following tasks involving Solstice Backup backup groups:

- Immediately start a scheduled archive
- Stop a scheduled archive
- Disable a scheduled archive
- Reschedule a scheduled archive
- View the details of a scheduled archive

▼ To Start a Scheduled Archive Immediately

When you elect to start a scheduled archive immediately, the Solstice Backup 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. Select Server>Archive Request Control or click the Archive Request Control toolbar button to open the Archive Request Control window.
2. Select the name of the archive request you want to start immediately in the Archive attribute.
3. Click Start to initiate the archive request.

▼ To Stop a Scheduled Archive While In Progress

To stop a scheduled archive:

1. Select Server>Archive Request Control or click the Archive Request Control toolbar button to open the Archive Request Control window.

2. Select the name of the archive request you want to stop immediately in the Archive attribute.
3. Click Stop to stop the archive procedure in progress.

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 in order to complete and fulfill the archive request, follow the procedure [“To Start a Scheduled Archive Immediately”](#) .

▼ To Disable a Scheduled Archive

You can prevent a scheduled archive from occurring by disabling it. To disable a scheduled archive:

1. Select Server>Archive Request Control or click the Archive Request Control toolbar button to open the Archive Request Control window.
2. Select the name of the stopped archive request you want to disable in the Archive attribute.
3. Click Disable to prevent the archive request from being fulfilled.

▼ To Reschedule a Scheduled Archive

You can change the scheduled time of an existing archive request without having to modify the archive request. To reschedule an archive:

1. Select Server>Archive Request Control or click the Archive Request Control toolbar button to open the Archive Request Control window.
2. Select the name of the archive request whose archive time you want to reschedule in the Archive attribute.
3. Click Schedule to open the Archive Request Schedule window.
4. Enter a new time you want the archive request to take place. 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.
5. Click OK.

▼ To View Details of a Scheduled Archive

To view the details of a scheduled archive:

1. Select Server>Archive Request Control or click the Archive Request Control toolbar button to open the Archive Request Control window.
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 processing of an archive request, including whether or not the archive was completed successfully. You can also save or print this information from within the Archive Request Details window.

To troubleshoot failed archives, see the messages in the Solstice Backup *log* file.

Archived Data Retrieval Operations

This section describes the task of recovering archived data, a process known as *retrieving* data. You can browse and select archived save sets for retrieval using the `nwretrieve` program.

As with the Solstice Backup recovery feature, retrieved data can be relocated, renamed, or allowed to overwrite existing files of the same name. Before the retrieval begins, the Solstice Backup server prompts you to determine how to handle filename conflicts.

You or a user can begin retrieving archived data as long as the volume is mounted in the appropriate storage device on the Solstice Backup server. If the volumes are not mounted, the server sends a message according to the configurations in the Notifications resource for tape mount requests.

The distinctions made between archiving non-indexed UNIX-based data and indexed PC-based data are reflected in how that archived data can be retrieved.

The nwretrieve Program

The Solstice Backup tool for retrieving archived data is the `nwretrieve` program. You can retrieve archived save sets from the volumes managed by the Solstice Backup server or storage node, or you can recover data from the client file index entries in addition to the entire save set. To view a display of the `nwretrieve` program, see [“nwretrieve” on page 38](#).

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

```
# nwretrieve -s server
```

Use the `-s server` option to specify a particular Solstice Backup server on the network. If you enter `nwretrieve` without the `-s server` option, then `nwretrieve` connects to your default backup server.

Note: The `nwretrieve` program is capable of retrieving the archived client data of both UNIX and PC clients in save set form only. If you want to browse and select individual items from a PC client archive, you must use the Solstice Backup for Windows User program. For more information, refer to your Solstice Backup for Windows products documentation.

Permissions

In order to retrieve or recover archived data, you must have administrator or archive user privileges for that save set, or enable the Public Archives attribute in the server's Server resource.

Retrieving Archived Data to the Client

The Archive Services attribute in the Clients resource controls whether the archive feature is enabled or disabled for each client. After you have enabled the optional NetWorker Archive application on the Solstice Backup server, archive services are available for all Solstice Backup clients that connect to the server. The Archive Services attribute in the Clients resource allows you to restrict the services to selected clients. When you enable or disable the Archive Services attribute for a client, the attribute changes for all Clients resources with the same name.

To restrict the users of the client computer that can request archives services, enter user IDs in the Archive User attributes of the Clients resource.

Retrieving Archived Data to Another Client

When the Public Archives attribute in the Server resource is enabled, all client computers that have Client resources defined on the Solstice Backup server can retrieve the archives of all other client computers defined on that Solstice Backup server.

Non-Indexed Archive

Because archived data is not usually recorded in the client file index, all the data in an archive save set is retrieved as a single unit. The `nwretrieve` program displays the archived save sets on that server, listed by client name. You may retrieve a particular save set only if you have administrator or archive user privileges for that server. The annotation you gave the save set at the time of its archive is required for locating the non-indexed archive.

The `nwretrieve` program is only designed to recover archived save sets, which is the standard method of archiving for UNIX-based clients whose archived data is written to the Archive pool.

Indexed PC Archive

Unlike non-indexed archives, you can browse and select individual files and directories for recovery that were originally written to the PC Archive pool during the archiving process. The Store Index Entries attribute in the PC Archive pool resource must be enabled at the time of the archiving for individual filenames to be recorded in the client file index.

In order to browse and select archived data from PC-based clients, you must use the PC-based Solstice Backup User program (`winworkr.exe`), which is the equivalent of the UNIX `nwarchive` and `nwretrieve` programs. For information on the Solstice Backup User program, refer your Solstice Backup for Windows server and client documentation.

▼ To Retrieve Archived Save Sets

To retrieve archived save sets:

1. Start the `nwretrieve` program as follows:
 - If you want to retrieve data from the default Solstice Backup server, enter the following command:

nwretrieve

- If you want to retrieve data from a Solstice Backup server other than your default server, enter the following command:

```
# nwretrieve -s server
```

2. Enter the name of the client whose archived data you want to retrieve in the Client attribute.
3. Enter the name of the archived save set you want to retrieve in the Save Set Name attribute.

If you are not sure of the name of the archived save set, you can leave this attribute blank and search for the archived save set's annotation instead.

4. Enter all or part of the annotation assigned to the save set at the time of its archiving in the Search Annotation For attribute.

To aid the search, enter wildcard symbols in this attribute to aid the search.

5. Click Query to search for the save set(s) that match the criteria you specified in the Clients, Save Set Name, and Search Annotation For attribute. The matching archived save sets are displayed in the Archives attribute.
6. Select any archived save set listed in the Archives attribute and click Details to open the Retrieve Details window.

Here you can view details of an archived save set such as its size, save set ID, and archive date.

7. Select the archived save set you want to retrieve and select File>Start Retrieve or click the Start toolbar button to begin the retrieval.

The Retrieve Status window appears. Here you can enter information similar to when you recovered normal save sets.

8. Perform any or all the tasks associated with the features in the Save Set Recover Status window:
 - Relocate recovered save set data to a non-default location (see [“To Relocate Recovered Save Set Data” on page 232](#)).
 - Overwrite existing files (see [“To Resolve Recovered Save Set Data Conflicts” on page 233](#)).
9. Click OK to begin the archived save set data retrieval. The retrieval status is displayed in the Status attribute of the Retrieve Status window.

Cloning

The Solstice Backup server can make duplicates, or *clones*, of complete volumes of stored data as well as individual save sets. Cloning allows you to make duplicates of data for purposes of storing offsite, sending data to another location, or verifying backups.

Volume Cloning

Volume cloning is a process of reproducing complete save sets from a storage volume to a clone volume. You can clone save set data from backups, archives, or migration. You can clone save sets automatically (as part of a backup, archive, or migration operation), or manually at another time.

How Volume Cloning Works

Volume cloning is a two-part operation:

1. The Solstice Backup server recovers data from the source volume.
2. The Solstice Backup server writes the data to a clone volume (a volume from a *clone pool*).

Caution – Cloning requires at least two active devices, because one is required for reading the source volume, and one is required for writing the new, cloned data. Essentially, this means that the Active Devices attribute must be set to greater than 1 in order for cloning to succeed.

During cloning, the reproduction of data is from source volume to clone volume. Cloning does not involve data stored on the clients or server. The Solstice Backup server allows only one clone of a save set per volume. Therefore, if you specify three clones of a save set, each is written to a separate 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. The cloning operation leaves traceable information entries in both the client file index and the media database. Having the capability to track cloned data is what distinguishes cloning from an operating system or hardware device copy operation.

When you specify that a particular volume be cloned, the Solstice Backup server uses the save sets on the specified volume as the source data. When you specify a clone of a particular save set, the server determines whether the 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 requires human intervention to mount. Command line options enable you to specify the precise save set clone to use as the source, if you want.

If you execute a clone operation manually, no completion report is generated. The Solstice Backup Administrator program displays cloning information in the Messages window. Clone-related messages are also logged to the Solstice Backup message file `/nsr/logs/messages`.

Automatic versus Manual Cloning

Automatic volume cloning (cloning associated with a scheduled group backup operation) is performed after all backup operations are complete. The savegroup completion report that is issued after a scheduled backup also includes a report of the success or failure of the cloning operation for each save set.

The location of the devices where the clone data is written is determined by the list in the Storage Nodes attribute in the Clients resource for the Solstice Backup server. You can add or remove the names of storage nodes and the Solstice Backup server at any time, but you cannot have a different list of storage nodes to receive clone data than you have to receive backup data.

If you want to perform cloning long after a group has finished, you must perform the cloning manually, volume by volume, or from the command line using a script in combination with a batch file. If you execute cloning manually, no report is generated.

Volume Cloning Versus Volume Duplication

When you clone a volume, the volume is not simply duplicated. Each save set on the volume is reproduced completely, which might mean that more or less space is used on the clone volume than on the source volume.

You might prefer to just make exact copies (duplicates) of Solstice Backup volumes to provide additional disaster recovery protection. This approach is not recommended, but it can serve certain environments adequately. If you rely on an exact copy command, you must first ensure that the destination volume can hold the number of bytes that are contained in the source Solstice Backup volume. In addition, the Solstice Backup server will have no information about the duplicated volume.

If you enabled automated media management and you leave the volume in an autochanger managed by Solstice Backup software, the volume may be considered eligible for relabeling and use during a scheduled backup. This is because the volume does not have a valid Solstice Backup label.

Similarly, it is possible to make an exact copy of a Solstice Backup archive volume. However, the annotation that is associated with each archive save set is information that is stored in the Solstice Backup server's media database, not on the volume itself. Therefore, a duplicate volume of the archived save set does not include the annotation. If the entry of the original archive save set is removed from the media database, the annotation that describes it is also lost.

Volume Cloning Requirements

The requirements for cloning save sets with the Solstice Backup server are listed below:

- Multiple storage devices, one to read the existing data and one to write the cloned data.
- Source and destination volumes.

Volume Cloning and Data Tracking Information

The clone operation does not insert entries into the client file index. Cloned save sets are only tracked through the media database. During cloning operation, 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 identification (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 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

Caution – You might decide to manually change the mode of a cloned volume to recycle with the intent of reusing a particular 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 the mode of the volume changes to recycle, you effectively change the status of all save sets on the clone volume to recycle. 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 instances of a save set that exist on other volumes, first change the mode of the volumes that you want to protect to manual recycle. This way, the Solstice Backup server cannot automatically recycle the volume. Then, you can safely change the volume that you intend for reuse to *recyc*.

Similarly, if you *purge* a clone volume, you effectively remove from the client file index all the file entries associated with all save sets that reside there (in whole or in part).

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. From the entry, the `nsrim` program marks for deletion the information about the location of one of the save set clones 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 Performance

In general, a volume write that occurs as part of a backup operation and one that occurs as part of a cloning operation proceed at the same speed. However, if a clone operation is automatically requested as part of a scheduled backup, you may experience a performance degradation in other scheduled backups that follow.

The Solstice Backup server generally attempts to complete one group's scheduled backup before a scheduled backup is initiated for another group. However, the server considers that a group backup is 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, you may experience contention for `nsrmmmd` resources or specific volumes.

To avoid this problem, you may decide to refrain from automatic cloning and instead initiate a single clone operation by passing a set of `ssids` to `nsrclone` as part of a job that runs at a nonpeak time after all backups are complete.

Volume Cloning and Recovery

A clone volume is used for recovery any time the Solstice Backup 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.

You can always 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, traditional recovery is available. For more information on the `scanner` program, see [“Recovering Client File Indexes” on page 238](#). For information on how to recover data with the `scanner` program, refer to the *Disaster Recovery Guide*.

▼ To Clone a Volume

You can manually clone all of the data on a volume in the Clone Volumes window. This window displays information about the volumes for the server currently in use and is a starting point for beginning the cloning process.

To clone a volume:

1. Select `Media>Clone Volumes` to open the Clone Volume window.
2. Select the volume you want to clone from the Volumes attribute and click `Details` to open the Clone Volume Status window.
3. Enter the name of the clone pool to which the clone volume will belong. Cloned volumes must belong to a clone pool, either one you created or one of the preconfigured clone pools. If you do not select a clone pool, the Solstice Backup server automatically selects the Default Clone Pool.
4. Click `Start` to begin the cloning.

Once the cloning process begins, its progress is displayed in the Status attribute of the Clone Volume Status window.

▼ To View Clone Volume Details

You can view the details of any clone volume, including its capacity, amount of space used, number of files, and mode, by opening the Volume Details window.

To view clone volume details:

1. Select Media>Clone Volume to open the Clone Volume window.
2. 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 provides you with the ability 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 a single save set or configure the Solstice Backup server to clone only full backups for an entire group.

Overview

The Solstice Backup server ensures the safety of your data by making a copy of each selected save set and then sending 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 to two clone volumes simply because the destination volumes may be a different size or have a different capacity. However, the opposite is also possible: the original save sets might require two volumes, while the cloned save sets require only one volume.

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. When you apply the manual recycle mode to a volume, the Solstice Backup server will not overwrite the volume unless you manually change the mode again to recyclable.

Performance Levels

To maintain backup and cloning performance levels, the Solstice Backup server clones save sets after completing a backup. You should allow time to complete the cloning process before scheduling another backup. By providing adequate time to complete either backing up or cloning the data, you avoid overloading the server.

Online Index Entries

The Solstice Backup server creates index entries for the cloned save sets and volumes in the online media index. However, both the original and cloned volumes use the same set of entries in the online client file index. If you eliminate either the cloned or original volume from the media database, the Solstice Backup server does not let you remove the client file index entries, as long as one of the volumes still exists.

Volumes

The original volume is the *source* volume, and the volume you clone the data to is the *destination* volume. The destination volume must be a different volume than the source volume, and it must belong to a clone pool.

Devices

You must have two storage devices to clone data. Use one storage device for the source volume and use the other for the destination volume:

- If you have an autochanger with two devices, the Solstice Backup server mounts the volumes required for cloning automatically.
- If you have stand-alone devices, you need to mount the volumes manually. The Solstice Backup server determines which volumes are required, both destination and source, and displays a message in the Pending Messages window to inform you which volumes to mount. If you want to keep the cloned data separate, you might need to premount a blank destination volume.

Save Sets Cloning Requirements

Below are the requirements for cloning save sets:

- Multiple storage devices, one to read the existing data and one to write the cloned data.
- Source and destination volumes

Save Sets Cloning Methods

You have two different methods to clone save set data:

- Automated cloning of save sets can be set in the Groups resource. For general information on configuring groups, including cloning, see [“Backup Groups” on page 182](#).
- Manual cloning of save sets is performed by entering the appropriate information in the Save Set Clone window. The next section describes this process.

Manual Save Set Cloning

Manual save set cloning is performed by the Solstice Backup administrator on an as-needed basis. You can clone an unlimited number of save sets at the same time. Because manual save set cloning requires much more labor than automated cloning, preference should be given to automated cloning whenever you know in advance you will require save set cloning.

The next section provides an example of an unplanned save set cloning request, and the steps to manually clone save set data.

Example: Preparing to Manually Clone Save Sets

In this scenario, a user has requested the Solstice Backup administrator to manually clone several save sets that are not included in an automatic cloning schedule.

The user needs to send the data to another company location out of state that also uses the Solstice Backup software to back up its network data. The Solstice Backup 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` command to retrieve the data. For more information about the `scanner` command, see [“Recovering Client File Indexes” on page 238](#).

To clone the save set for the user, the Solstice Backup administrator must have the following information:

- Solstice Backup client name
- Name of the save set
- Date the data was backed up

▼ To Manually Clone Save Sets

To clone data successfully, make sure you set the Active devices attribute, located in the Setup Server dialog box, to two or higher. Cloning requires more than one storage device.

To manually clone save sets:

1. Select Save Set>Clone to open the Save Set Clone window.
2. Click More to display additional attributes in the Save Set Clone window.
3. Enter the client name whose save sets you want to browse in the Client attribute. If you do not enter a client name, the Solstice Backup Administrator locates and displays all of the save sets for all clients.
4. Optional. Enter the name of the save set in the Save Set attribute.
5. Optional. Enter the name of the save sets' volume in the Volume attribute.
6. Optional. Enter the name of the Volume pool in the Pool attribute.
7. Optional. Click one or more choices for the sorting criteria to determine the order the save sets will appear in the Save Sets list.
8. Enter the range of dates the save set was created in the Start and End Date attributes. By default, yesterday's date appears as the Start Date and today's date as the End Date. Enter the dates as one of the following:
 - Written out completely (for example, November 1, 2000)
 - Numerically as mm/dd/yy
 - As a phrase (for example, two weeks ago, three days ago)
9. Optional. Select the appropriate Status attributes for the save sets. Browsable, Recyclable, Scanned-in, and Recoverable are already chosen by default. Solstice Backup software uses seven save set status variations:

- If the save set still has an entry in the online file index, select **Browsable**.
 - If all save sets have passed both the browse and retention policy time periods and the volume is available for automatic relabeling and overwriting, select **Recyclable**.
 - If the save set was scanned in using the scanner command, select **Scanned-in**.
 - If the entry for the save set has been removed from the client file index, but is still available for recovery from the media, select **Recoverable**.
 - If the save set is in the process of being backed up, select **In-progress**.
 - If the save set was aborted either manually by the administrator during backup or because the computer crashed, select **Aborted**.
 - If a previous recovery attempt failed for the save set, select **Suspect**.
10. Select the appropriate **Maximum Level** attribute(s) to determine the levels of the save sets to browse (default is **incremental**). The Solstice Backup server searches for and displays all the levels dependent upon your level choice. For example:
- If you select level 5, the Solstice Backup Administrator program browses and displays the following save set levels: full, one, two, three, four, and five.
 - If you select full, the Solstice Backup Administrator program searches for and displays only level full save sets.

You must always have at least one level selected, but you may not select more than one level at a time.

11. Click **Query** to display the save sets that fit your filtering criteria in the **Save Sets** attribute.
12. Select the save set(s) to clone from the **Save Set** attribute.
13. Click **Clone** to begin the cloning.

After clicking **Clone**, the **Save Set Clone** status window appears, displaying information on the status of the cloning operation.

▼ To View Details of Save Sets Prior to Cloning

You can view details of a selected save set you plan to clone in the **Save Set Clone Details** window. Information includes the name and date of the save set, the number of files it contains, along with its size and status.

To view details of save sets for cloning:

1. Select **Save Set>Clone** to open the **Save Set Clone Window**.
2. Click **Query** to display a list of save sets in the **Save Sets** attribute.

3. Select the save set whose details you want to view.
4. Click Details to open the Save Set Clone Details window.

Automated Save Set Cloning

The cloning of save set data can be performed automatically by the Solstice Backup server at the time of a scheduled group backup. The backed-up save sets you identify for automated cloning are cloned without any direct involvement by the administrator. Whenever you know of a save set cloning requirement in advance, you should try to use automated save set cloning.

The next sections provide examples of preplanned save set cloning request for a single client or multiple clients and the steps to automatically clone save set data.

Example: Automatically Cloning One Client

In this scenario, a system administrator wants to clone data automatically on a daily basis for the new computer *shakespeare*, where the company product documentation resides. *shakespeare* performs an incremental backup every night of the week except Friday, when it performs a full backup.

The system administrator needs to accomplish the following tasks before the cloning can take place:

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

▼ To Automatically Clone a Single Client

1. Create the group resource. For more information on creating groups, see [“To Create Groups” on page 185](#).
2. Select the appropriate attribute values in the *Documents* Groups resource. The required choices are:
 - Autostart = Enabled
 - Clones = Yes
 - Clone Pool = Default Clone

3. Create the client resource. For more information on creating clients, see [“To Create Clients”](#) on page 90.
4. Select the appropriate attribute values for the client resource. The required choices are:
 - Schedule = Full Every Friday
 - Group = Documents

FIGURE 9-1 illustrates the level of control you can achieve when cloning data with the Solstice Backup server. You can create an unlimited number of automatic cloning configurations using the Groups, Schedules, and Client resources.

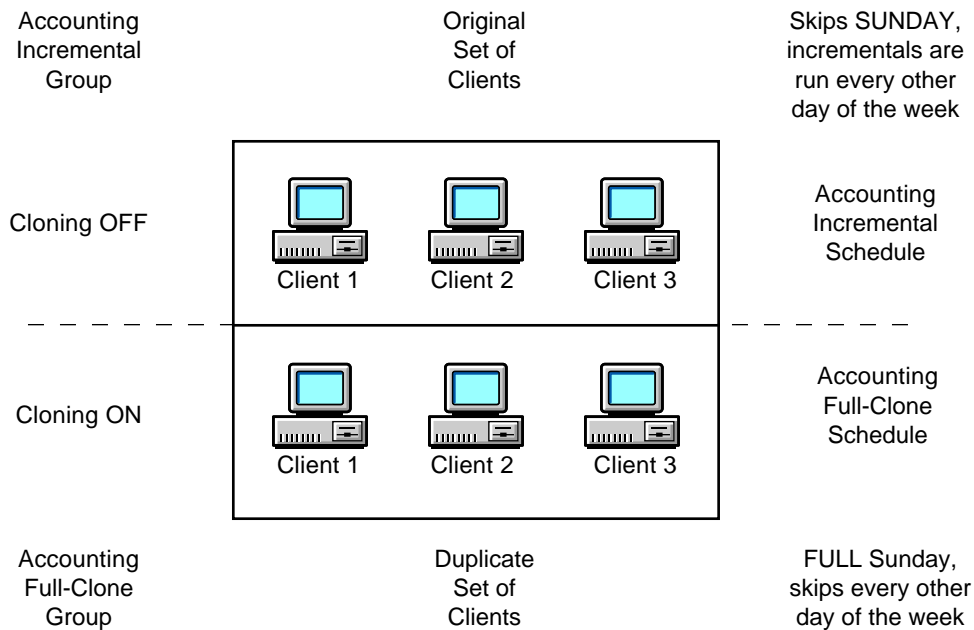


FIGURE 9-1 Automatic Cloning Scheme

Example: Automatically Cloning Several Clients

In this example, the system administrator wants to automatically clone the data for a small group of Solstice Backup clients that have valuable accounting information. However, he or she only wants to clone full backups to store in an offsite vault.

To do this, two different groups and schedules need to be created and applied to two identical sets of individual clients in the Client resource.

Both groups belong to the same pool and use the browse and retention policies specified for the Client resource.

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

The following is an overview of the tasks the system administrator needs to accomplish:

- Create two schedules, one to perform incremental backups for the clients and one to perform full backups and clones:
 - [TABLE 9-1](#) displays values for creating one schedule named *Accounting Incremental*.

TABLE 9-1 Accounting Incremental Schedule

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

- [TABLE 9-2](#) displays values for creating another schedule named *Accounting Full-Clone*.

TABLE 9-2 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 a group. Therefore, two groups are necessary in this situation:

- The Accounting Incremental group is used for performing incremental backups.
- The Accounting Full-Clone group is used for performing full backups and clones.

For more information on creating groups, see [“To Create Groups” on page 185](#).

- Create two sets of identical clients. Two sets are needed because only one schedule can be applied to a Solstice Backup client at a time:
 - One set of clients uses the Accounting Incremental schedule

- The other (identical) set of clients uses the Accounting Full-Clone schedule.
- Select the appropriate group and schedule for each set of clients in the Client resource.

▼ To Automatically Clone Multiple Clients

These instructions are just one example of how a system administrator might fulfill automatic cloning requirements based on the previous example.

1. Select Customize>Schedules to open the Schedules window.
2. Create the schedule *Accounting Incremental* to perform incremental backups every day except Sunday, when it skips the backup (see [TABLE 9-1](#)).
3. Create the schedule *Accounting Full-Clone* to perform full backups every Sunday and skip backups the rest of the week (see [TABLE 9-2](#)).
4. Create the group *Accounting Incremental*. Set the following attribute values:
 - Autostart = Enabled
 - Clones = No
5. Create the group *Accounting Full-Clone*. Set the following attribute values:
 - Autostart = Enabled
 - Clones = Yes
 - Clone = Backup Clone
6. Create two sets of identical clients and select the appropriate schedule and group for each.
7. Select each client in the first set of clients. Assign the client to the Accounting Incremental group. Set the following attribute values:
 - Schedule = Accounting Incremental
 - Group = Accounting Incremental
8. Click Apply after making these choices for each client.
9. Select each client in the duplicate set of clients and assign them to the group Accounting Full-Clone. Set the following attribute values:
 - Schedule = Accounting Full-Clone
 - Group = Accounting Full-Clone
10. Click Apply after making these choices for each client.

When the group Accounting Full-Clone 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.

Cloning Archived Data

If you already made an archive volume and now want to clone it, follow these steps:

1. Select Save Set>Clone to open the Save Set Clone window.
 2. Make your selections for browsing save sets. For instructions, see [“To Manually Clone Save Sets” on page 271](#).
 3. Enter *Archive* in the Pool attribute.
 4. Click Query to display the save sets matching your criteria.
 5. Select the save set you want to clone and click Clone.
-

Recovering Cloned Save Sets

When recovering save sets, the Solstice Backup server always looks for the original volumes by default. For more information on opening and viewing the Recover Volume Required window, see [“To View Volumes Required for Data Recovery” on page 221](#).

In order to recover save set files from a cloned volume rather than the original volume, you must first designate the original save set as suspect, and thus unable to be recovered. Once the Solstice Backup server realizes that the original save set is unusable, it will look for a cloned version. The clone volume containing the cloned save sets will then be displayed in the Recover Volumes Required Volumes window.

The Solstice Backup server does not display more than one set of volumes needed for recovery. Therefore, although there may be more than one set of clones needing recovery, only one set will display at a time.

Note: The process for recovering save set data from a cloned volume is the same as for a regular volume. See [“To Recover Data from a Save Set” on page 230](#).

▼ To Change the Original Save Set to Suspect

To change the original save set to suspect:

1. Select Media>Volumes or click the Volumes toolbar button to open the Volumes window.
2. Select the save set whose status you want to change from the Save Set attribute.
3. Select Save Set>Change Status>Suspect to change the status to *suspect*.
4. Close the Volumes window.

Now that the original save set has *suspect* status, the Solstice Backup server will not use that save set for recovery unless its status is changed back to *normal*. Until that occurs, all subsequent recoveries of that save set will require the cloned volume containing the same cloned save set. From this point on, you can proceed as you would for a normal save set recovery.

Note: If you cannot determine which specific save set on the volume is the one you want to recover in order to make it *suspect*, you need to change the status of all of the save sets on the volume to *suspect*. This ensures that the Solstice Backup server will not find the save set it needs on the original volume during the recovery process.

Clone Storage Node Affinity

The link between a client's resource of a storage node and a list of available storage nodes to receive cloned save sets from the storage node client is called *clone storage node affinity*. Data is cloned from media that contains the original save sets to media on the specified clone storage node. You define clone storage node affinity in the Clone Storage Nodes attribute, which is found in the Clients resource of a storage node. When you make a change, the changed value is propagated to any additional Clients resources configured for that storage node client.

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 utilizes the exact hostname you specify in the attribute, instead of using the hostname prefix for the remote device name configured in the Devices resource.

When a volume is being cloned, the Solstice Backup server checks the value of the Clone Storage Nodes attribute for that storage node client:

- If the Clone Storage Nodes attribute has a null value, then the value listed in the server's Clone Storage Nodes attribute is used.

- If that list also contains 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 Clients resources you configure for the storage nodes, and configure only the Solstice Backup server's Clone Storage Nodes attribute. This tactic 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 Solstice Backup 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 configuring and operating the Solstice Backup server 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 the Solstice Backup software with an autochanger requires you to purchase and enable the Autochanger Software Module. Contact your Solstice Backup sales representative for more information. For general information on licensing Solstice Backup products, see [“Licensing” on page 39](#).

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 Solstice Backup software. Stackers access tapes sequentially, while autochangers can access media in any order.

To back up to an autochanger, you need to accomplish the following tasks:

- Install and enable the Solstice Backup software.
- Install and enable the Autochanger Software Module.
- Load and label your 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, bar code choices, auto media management, and cleaning devices.

The Solstice Backup 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.

▼ To Add Autochangers

You cannot add or create autochangers using the Solstice Backup Administrator program; instead you need to use the `jbconfig` program at the command line. However, once the autochanger has been successfully added and configured, you can subsequently modify its configuration using the Solstice Backup Administrator program.

Caution – Prior to running the `jbconfig` program, check what existing devices are defined in the Devices panel of the Solstice Backup Administrator program. By default, the Solstice Backup software will configure a single tape device as a 8mm media type. If this is not the correct media type, then delete this device. For information on deleting devices, see [“To Delete Devices” on page 158](#).

If you do not perform this task, the Solstice Backup server will backup to your device using the wrong blocking factor. While adding the autochanger using the `jbconfig` program, you will be able to specify the correct devices your autochanger contains.

Follow this procedure to add autochangers:

1. Navigate to the directory containing your Solstice Backup software.
2. Enter the following command at the prompt:

```
# jbconfig
```

3. The `jbconfig` program displays the following installation script and asks for your choices. Press [Return] to accept the suggestions.

The following example shows the required responses to configure a SCSI autochanger on a Solaris system. The driver software detects and displays the information for all SCSI autochangers attached to the system.

```
# jbconfig

1) Install a GEMS SmartMedia Jukebox.

2) Install an Autodetected SCSI Jukebox.

3) Install an SJI Jukebox.

4) Install an RLM Jukebox.

5) Install an STL Silo.

What kind of Jukebox are you installing? 2

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

Which one do you want to install? 2

Installing a 'Quantum DLT/Compaq DLT' jukebox.

Name you would like to assign to the jukebox device? QuantumDLT_1

Pathname of the control port for the jukebox device?
[scsidev@0.2.1] [Return]

Do you want automated device cleaning support enabled? (yes/no) n

Enter pathname of media drive 1 [/dev/nrst8]:? [Return]

This media device has not been configured yet. Please select a
media device type for /dev/nrst8.

a) hint
```

```
b) qic
c) 4mm
d) 8mm
e) 8mm 5GB
f) 3480
g) dlt
h) vhs
i) optical

Choice? g

Jukebox has been added successfully
```

The following example shows the responses to configure an SJI autochanger.

```
# jbconfig

1) Install a GEMS SmartMedia Jukebox.
2) Install an Autodetected SCSI Jukebox.
3) Install an SJI Jukebox.
4) Install an RLM Jukebox.
5) Install an STL Silo.

What kind of Jukebox are you installing? 3

Enter the number corresponding to the type of jukebox you are
installing:

1) ADIC-1200c/ADIC-1200d          22) Qualstar
```

- | | |
|--------------------------------|--|
| 2) ADIC-VLS | 23) Spectralogic |
| 3) ARC DiamondBack | 24) STK 9704/Lago 340 |
| 4) Sun 20Gb 4mm Tape Loader | 25) STK 9708/Lago 380 (SCSI) Datawheel |
| 5) Breece Hill | 26) StorageTek 9730 |
| 6) Breece Hill Saguaro | 27) StorageTek 9738 |
| 7) Philips Blackjack | 28) Dell PowerVault 130T |
| 8) DLI Libra Series | 29) Hewlett-Packard A4853A |
| 9) Quantum DLT/Compaq DLT | 30) IBM 3570 |
| 10) Exabyte 10e or 10h | 31) IBM 7331/IBM 9427 |
| 11) Exabyte 10i | 32) ATL/Odetics SCSI |
| 12) Exabyte 18D | 33) HP-Optical |
| 13) Exabyte 60 | 34) Sony TSL-7000 |
| 14) Exabyte 120 | 35) Compaq 4mm DAT TLZ9L |
| 15) Exabyte 210 | 36) Compaq 4mm DAT (TLZxx) |
| 16) Exabyte 220 | 37) Compaq TL800 Series |
| 17) Exabyte 218 | 38) Compaq TL810 Series |
| 18) Exabyte 400 Series | 39) Compaq TL820 Series |
| 19) HP-C1553A/Surestore 12000e | 40) Compaq Optical |
| 20) HP-C1557A/Surestore 12000e | 41) Compaq TK Series |
| 21) Metrum (SCSI) | 42) Standard SCSI Jukebox |

Choice? **9**

Installing a 'Quantum DLT/Compaq DLT' jukebox.

Name you would like to assign to the jukebox device? **dlt**

Pathname of the control port for the jukebox device? scsidev@1.2.0
[**Return**]

Do you want automated device cleaning support enabled? (yes/no) **n**

Enter pathname of media drive 1 [/dev/nrst8]:? [**Return**]

This media device has not been configured yet. Please select a media device type for /dev/nrst8.

a) himt

b) qic

c) 4mm

d) 8mm

e) 8mm 5GB

f) 3480

g) dlt

h) vhs

i) optical

Choice? **c**

Jukebox has been added successfully.

▼ To Test the Autochanger Connection

To test the autochanger connection:

1. Become root on the Solstice Backup 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
```

You need to 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` command 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 `jbxercise` program on the Quantum DLT/Compaq DLT autochanger detected by the `jbconfig` script:

```
# jbxercise -c scsidev@0.2.1 -m "Quantum DLT/Compaq DLT"
```

For more information, see **jbxercise** in [the Legato Command Reference Guide](#) or refer to the `jbxercise(1m)` man page.

Key Autochanger Attributes

TABLE 10-1 lists some of the key autochanger attributes and their functions.

TABLE 10-1 Autochanger Attributes

| Attribute | Action |
|-----------------------|--|
| Control Port | Enter the path name of the autochanger's command control port. |
| Devices | Enter the names of each device in the autochanger. |
| Bar Code Reader | Select Yes to enable this function for an autochanger with bar code reading capability. |
| Match Bar Code Labels | Select Yes for your bar code-enabled autochanger to match bar code 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 Solstice Backup server sends you a notification to replace the cartridge. |
| Auto Clean | Select Yes to enable autocleaning. |
| 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 Solstice Backup server to use. |

▼ To Add Devices to an Autochanger

To add more devices to an autochanger, you must first configure the Solstice Backup server to recognize the added devices. The Device resource contains attributes for you to enter information about the additional autochanger device. Then you must enter the pathname for the autochanger device in the Jukebox resource so the server can recognize that the device belongs to the autochanger. For an example of pathname naming, see [FIGURE 10-1](#).

To add devices to an autochanger:

1. Select Media>Devices to open the Devices window.
2. Click Create.
3. Enter the autochanger device pathname in the Name attribute.

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
- Become root on the Solstice Backup server or storage node computer and enter the `inquire` command at the shell prompt. The `inquire` command returns a list of the SCSI devices attached to your system.

The `inquire` command is only supported for AIX, Solaris, IRIX, Tru64 UNIX and DYNIX/ptx systems. To view a list of the SCSI devices attached to an HP-UX system, enter the `ioscan -f` command. For instructions on how to configure support for the Autochanger Module with an HP-UX system, refer to the *Legato NetWorker Installation Guide, HP-UX Version*.

4. Select the appropriate media type in the Media Type attribute.
5. Complete the other attributes as necessary and click Apply. See [“Devices” on page 155](#) for more details on configuring devices.
6. Select Media>Jukeboxes to open the Jukeboxes resource.
7. Select the appropriate jukebox in the Jukeboxes attribute.
8. Enter the storage device pathname in the Devices attribute and click Add.

If you like, you can make other modifications to your autochanger configurations at this time. For example, you might select bar code labels, or decide to change the range of available slots.

If an autochanger has more than one device, the device pathnames must be listed in the same order as their physical location in the autochanger. If you are unsure of their order, refer to the autochanger hardware manual.

[FIGURE 10-1](#) shows an example of pathname naming.

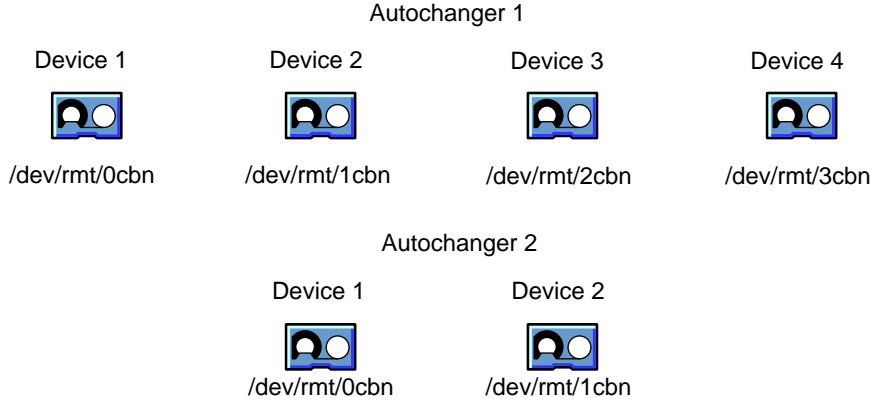


FIGURE 10-1 Autochanger Device Pathnames

Caution – There is a need to reset an autochanger every time changes or modifications are made to the device. This should be done with the `sjieelm` utility and/or preferably with the `nsrjb -HE` command-line program. For example, the autochanger inventory must be correct after adding drives to a SJI compliant autochanger, such as adding DLT7000 drives to an ETL 7/3500 device.

To make the Solstice Backup software aware of these new drives, execute `nsrjb -HE` to reset the autochanger. This must be done after running `jbconfig`. The `-E` option re-initializes the autochanger's element status. Some autochangers have the ability to keep track of whether or not there is media in a component in the autochanger. This feature is known as an *element status* capability.

In some special cases, there is no need for this hardware reset if the autochanger is capable of reinventory by itself.

▼ To Edit an Autochanger

To edit an autochanger:

1. Select Media>Jukeboxes to open the Jukeboxes window.
2. Select the autochanger from the list displayed in the Jukeboxes attribute.
3. Edit the attributes of the autochanger and click Apply.

▼ To Delete an Autochanger

To delete an autochanger:

1. Select Media>Jukeboxes to open the Jukeboxes window.
2. Select the autochanger from the list displayed in the Jukeboxes attribute.
3. Click Delete.

When prompted to confirm that you really want to delete that notification, Click OK.

4. Click Apply.

If you decide that you want to reinstall the autochanger you removed, or to add a new autochanger to your Solstice Backup server, remember you need to run the `jbconfig` program.

Remote Autochanger Management

Autochangers that are connected to storage nodes require a few additional configuration and management steps.

You can control most operations on remote autochanger devices from the Solstice Backup Administration program. But for some remote autochanger operations, such as reset, you must use the `nsrjb` or `jbconfig` commands 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. The method for defining devices is described in [Chapter 5](#). An overview of how to add autochanger devices is provided here.

When you add a remote autochanger device, first add the storage node's hostname to the Administrator attribute in the Server resource by using the following syntax:

```
root@storage_node_hostname
```

Then, run the `jbconfig` program on the storage node computer to define each device in the autochanger. For more information, see ["Autochanger Installation and Configuration" on page 281](#), and refer to the *Legato Command Reference Guide* or the `jbconfig` man page.

The device names for remote devices begin with a prefix of “rd=” and the name of the storage node computer. For example, rd=omega:/dev/rmt/1mbn is a device called /dev/rmt/1mbn on a storage node computer called *omega*.

▼ To Adjust Sleep Periods for Autochanger Tasks

The Jukebox resource includes attributes that specify the number of seconds that the autochanger is inactive after volume loading, unloading, and/or ejecting. This function is known as *sleeping*, and during a sleep period, 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 in effect, the loading or unloading of volumes might fail.

When you configure a Jukebox resource using the `jbconfig` program, the Solstice Backup software automatically configures default sleeping periods. You should only change these values if you are troubleshooting your autochanger’s performance, or if a Solstice Backup technical support specialist instructs you. Typically, the higher the sleep values specified in the attributes, the longer it will take the autochanger to perform the task. Be cautious when changing these values.

[TABLE 10-2](#) lists the attributes related to the autochanger sleep function.

TABLE 10-2 Jukebox Resource Sleep Attributes

| Attribute | Description |
|--------------|---|
| Eject Sleep | Specifies the number of seconds the Solstice Backup software sleeps after ejecting a volume. the Solstice Backup software uses this attribute if the autochanger is not configured for volume auto-ejection. Autochangers configured for volume auto-ejection 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 Solstice Backup software uses this attribute for all autochangers when a volume is unloaded. |
| Load Sleep | Specifies the number of seconds the Solstice Backup software sleeps between loading a volume and trying to access that volume’s data. |

To access and set the sleep-related Jukebox resource attributes, follow these steps:

1. Select Media>Jukeboxes to open the Jukeboxes window.

2. Select View>Details to display hidden attributes.
3. Adjust one or more of the sleep-related attributes.

Sharing Autochangers Between Solstice Backup Hosts

The Solstice Backup software permits the control and usage of individual devices within an autochanger to different Solstice Backup hosts (a Solstice Backup server or storage node) within a datazone, a feature known as *autochanger sharing*.

The presence of a storage area network (SAN) within your datazone is not required for using autochanger sharing.

Autochanger Sharing Licensing and Availability

Licensing and availability of the Solstice Backup autochanger sharing feature varies depending of the type of Solstice Backup license you have.

- Solstice Backup Power Edition – Autochanger sharing is automatically enabled as part of Solstice Backup Power Edition. However, you must still have a valid Solstice Backup autochanger license for each autochanger.
- Solstice Backup Network Edition – An enabler license for Solstice Backup autochanger sharing is all that is required to use the autochanger sharing feature within a single Network Edition datazone. Only a single autochanger sharing enabler is required no matter how many autochangers you have within the datazone. However, you must still have a valid Solstice Backup autochanger license for each autochanger.
- Solstice Backup Workgroup Edition – Autochanger sharing is not available for Solstice Backup Workgroup Edition. If your datazone environment requires autochanger sharing, you should upgrade your Solstice Backup 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 your Solstice Backup datazone, regardless of the number of autochangers present. However, you will still have to license each autochanger to enable the autochanger functionality itself. For general information on licensing of Solstice Backup products, see [“Licensing” on page 39](#).

Note: Only Solstice Backup Network Edition requires licensing of autochanger sharing, since this feature is already enabled in Power Edition, and is not available with Workgroup Edition.

How Autochanger Sharing Works

Autochanger sharing allows one Solstice Backup host to control the autochanger's robotic arm, while other Solstice Backup 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 Solstice Backup host.

FIGURE 10-2 shows how multiple Solstice Backup hosts can share autochanger devices.




FIGURE 10-2 How Autochanger Sharing Works

Autochanger Sharing Configuration

You can set an autochanger for sharing during the initial autochanger setup using the `jbconfig` command. The `jbconfig` command must be run from the Solstice Backup host that controls the autochanger (controls the robotic arm). When running `jbconfig`, you will be prompted to enter the Solstice Backup hostname controlling the first autochanger device. The default hostname value is the local Solstice Backup host on which `jbconfig` is running. However, you can change the default hostname value to another Solstice Backup host within the datazone. You will continue to be prompted to enter the controlling host for each device in the autochanger.

Example: `jbconfig` Command Additions

Here is an example of the additions to the `jbconfig` command for the Solstice Backup 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/lc1bn

using 'rd=pluto:/dev/rmt/lc1bn' as device name
```

Device Naming During Configuration

The `jbconfig` command automatically generates the device names for a shared autochanger accordingly under the following conditions:

- If a device is controlled by the Solstice Backup server, only local device syntax (`device_path`) is used.
- 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 Solstice Backup server contains an attribute in the Jukebox resource for designating a host for use in recoveries when the host controlling a required device is not already listed in the Client resource's Storage Nodes attribute. This new attribute is the Read Hostname attribute, and its default value is the Solstice Backup 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 new functionality designates an alternative device to use for recoveries and cloning if a storage node's preferred hosts are not available. Normally, when a recovery is requested, the Solstice Backup server will match the host that is controlling the device required for the recovery with the hosts listed in the Storage Nodes list located 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 displays an example of a Solstice Backup datazone environment that consists of a server `venus`, client `andromeda`, and storage node `mars`. There is a shared autochanger with the first device connected to `venus`, and a second device connected to `mars`.

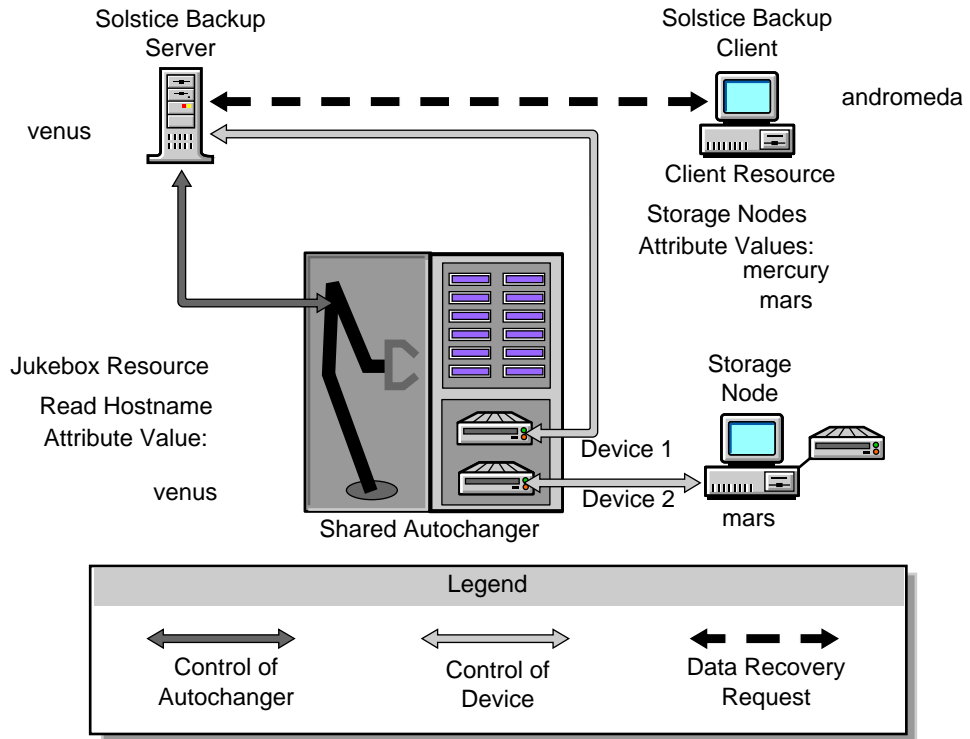


FIGURE 10-3 Data Recovery Using a Shared Autochanger

The client *andromeda* 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 *andromeda*'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 Solstice Backup 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 Solstice Backup 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 Solstice Backup 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 Solstice Backup host controlling the autochanger (such as having control of the robotic arm) regardless of whether the autochanger is shared or not.

Managing Volumes on Server Devices

For devices controlled by the Solstice Backup server, you manage your autochanger volumes using the `nsrjb -f` command with the device path, even if the `jbconfig` command for the shared autochanger was run from another host. For example, if you wanted to load a volume into a shared autochanger device that is controlled by the Solstice Backup server, you would enter the following command at the prompt:

```
#nsrjb -l -f device_path volume_name
```

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.

Note: In a datazone where no shared autochangers reside, the Solstice Backup server automatically adds the remote device host (`rd=host`) information, resulting in only the need to input the local device syntax (`device_path`). However, if there is one or more shared autochangers in the datazone, then you must always enter remote device syntax (`rd=host:device_path`) information with the device path.

Backups and Recoveries on a Shared Autochanger Device

Autochanger sharing also requires that the host of a volume being selected for a backup or recovery be further specified to reflect a particular host. Normally, the Solstice Backup server automatically performs this activity and requires no user input. However in cases where you need to perform this task at the command line, the `-J` option has been added to `nsrjb`. The `-J` option is only available when specified with one of the following:

- The `nsrjb -l` option for loading a volume.
- The `nsrjb -L` option for labeling a volume.

It is not available for use in conjunction with the `nsrjb -f` option for specifying a device path.

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 Solstice Backup resource for that autochanger. The Idle Device Timeout attribute specifies the number of minutes a mounted volume can remain idle before it is automatically unmounted from the device and returned to its slot, where it can then be accessed by another device.

Setting Compatible Block Size for UNIX and Windows NT

There are different SCSI hardware limitations between UNIX and Windows NT which can lead to compatibility problems. For example, if you have a device defined in UNIX and it is physically attached to a Windows NT host bus adapter (HBA), it is

possible to define a block size greater than that allowed by the Windows NT hardware. This could lead to I/O errors in both write and read states on this device. If you are going to use both platforms, you must determine a common block size acceptable to both.

Problem Determination

To determine the allowed block size, run the following command at the prompt:

Solaris:

```
# /usr/bin/mt -f \\.\Tape0 status
```

Windows NT:

```
# mt.exe -f \\.\Tape0 status
```

Solutions

There are two separate methods for solving this problem.

- On Solaris, 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_DEVICE_NAME
```

where *device_name* is the backup device type available to your Solstice Backup server. Note that the syntax for the device name to be all upper case, with underscores (_) replacing blank spaces. Therefore, a device displayed in the Solstice Backup software as “4mm 20gb” would be listed as follows:

```
4MM_20GB
```

- On Windows NT, you can install a later model HBA or upgrade to drivers that can support up to 128 KB blocks for the above Solaris example. Windows NT also accepts the same environmental variable format for setting block size as listed for UNIX.

Autochangers Operations

This section includes information for performing autochanger operations. 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 your autochanger.
- How often you need to mount volumes for backup manually.

▼ To Designate Available Autochanger Slots

The available slots feature allows you to control exactly which volumes The Solstice Backup 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 your Solstice Backup 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 Solstice Backup server can proceed uninterrupted with an automatic backup.

To enter the number of available slots:

1. Select Media>Jukeboxes to open the Jukeboxes window.
2. Select the autochanger whose available slots you want to change in the Jukeboxes attribute.
3. Enter the range of available slots into the Available Slots attribute and click Apply.

Note that with two-sided media, the number of available slots is effectively doubled. For example, if you have 32 optical disks, labeled “atlas.001.a” to “atlas.032.b,” you have a total of 64 sides, and therefore 64 slots from which to choose.

▼ To Clean Autochanger Devices

Maintaining a storage device in good working order requires periodic cleaning. The Solstice Backup 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 you to configure attributes in both the Jukebox resource and the Device resource of each appropriate device.

To clean autochanger devices:

1. Select Media>Jukeboxes to open the Jukeboxes window.
2. Select the autochanger for which you want to configure cleaning.
3. Enter the slot number of the autochanger that contains the cleaning cartridge in the Cleaning Slots attribute. If more than one slot contains a cleaning cartridge, specify the range with a hyphen, "8-9."
4. Enter the number of times the cleaning cartridge can be used before it needs to be replaced in the Default Cleanings attribute.
5. Select Yes in the Auto Clean attribute to enable autocleaning.
6. Click Apply.
7. Select Media>Devices to open the Devices window.
8. Select the appropriate device in the Devices attribute.
9. Select Yes in the Cleaning Required attribute.
10. Enter a value in the Cleaning Interval attribute. This attribute indicates the amount of time from the last cleaning until the next cleaning. The time may be specified in days, weeks, or months.
11. Click Apply.

Note that the last time the device was cleaned is listed in the Date Last Cleaned attribute.

12. Repeat these steps for each device included in the autochanger.

To ensure that there will not be any interference with other autochanger operations, the Solstice Backup server only cleans devices before mounting or after unmounting a volume in a device. For recommendations on the frequency and cleaning methods for its autochanger devices, check the autochanger manufacturer's documentation.

TABLE 10-3 lists the notification messages the Solstice Backup software displays to inform you about cleaning cartridge operations:

TABLE 10-3 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. Select **Media>Jukeboxes** to open the Jukeboxes window.
3. Select **View>Details** to display hidden attributes.
4. Select *Yes* in the Auto Clean attribute.
5. Specify the slot number you want to use for the cleaning cartridge in the Cleaning Slots attribute.
6. The Default Cleanings attribute lists the number of uses left on the cleaning cartridge. Record this number for use in [Step 8](#).
7. Change the Available Slots attribute to reflect the range of slots available for data volumes.

If the slot for your 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
```

8. Enter the following command at the prompt:

```
#nsrjb -U uses -j autochanger -s slot
```

- Replace *uses* with the number of uses left on the Default Cleaning attribute (the number you wrote down in [Step 6](#)).
- Replace *slot* with 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 your autochanger does not support the barcode labeling option, you must issue the command shown in [Step 8](#) to tell the autochanger to add the cleaning cartridge to its inventory.

▼ To Replace Cleaning Cartridges

When a cleaning cartridge expires, the Solstice Backup software displays the following message:

```
Date server_name: Solstice Backup Cleaning cartridge expired:  
(notice) Cleaning cartridge required in Jukebox 'autochanger_name'
```

Replace the expired cartridge and enter the following command at the prompt:

```
#nsrjb -U number of cleanings
```

If you do not enter the `nsr jib -U` command along with the number of cleanings, the Solstice Backup software does not recognize the new cartridge and continues to send the “cleaning cartridge required” notification.

If your autochanger does not support bar code labeling, you must enter the following command at the prompt to tell the autochanger to add a cleaning cartridge to its inventory:

```
# nsr jib -U number_of_uses -j autochanger_name -s slot
```

- Replace *number_of_uses* with the number of cleanings the cartridge supports.
- Replace *autochanger_name* with the name shown in the Jukebox attribute.
- Replace *slot* with the slot number where you inserted the cleaning cartridge.

Auto Media Management with Autochangers

Auto media management further automates the backup procedure by freeing administrators from the task of prelabeling volumes before backup. When auto media management is enabled in the Jukebox resource, the Solstice Backup software assumes that all the volumes loaded in the autochanger may be handled exclusively as Solstice Backup media. Volumes that appear unlabeled are considered blank and are automatically labeled, mounted, and overwritten with new data. The Solstice Backup software only uses volumes that it considers unlabeled if it cannot locate a writable volume or a volume eligible for automatic recycling.

The Solstice Backup software does not recognize the following volume labels and considers these volumes unlabeled:

- Volumes without a label.
- Volumes labeled with something other than a Solstice Backup label.
- Volumes with Solstice Backup labels that were written in a density different from the device in which the volume is currently loaded.

For example, if you performed Solstice Backup backups on a new 8mm tape drive, an older 8mm 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 Solstice Backup volumes between 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 Autochanger and Device resources. For autochangers, you only need to enable auto media management in the Autochangers resource. The Solstice Backup server does not allow you to enable auto media management in the Devices 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 [“Auto Media Management” on page 162](#).

When auto media management is enabled, the Solstice Backup server assumes that the autochanger is exclusive to the server, and the volumes loaded in the device are handled as Solstice Backup media. Do not enable this feature if the autochanger or device is a shared resource. The Solstice Backup 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 Solstice Backup Server Selects a Volume for Writing

If auto media management is enabled and the Solstice Backup software is connected to an autochanger, the autochanger searches for a mounted volume for writing. If no mounted volume is a candidate for writing, the Solstice Backup software looks for an unmounted, appendable volume from the appropriate pool that can be mounted.

If none is found, the Solstice Backup software looks for a recyclable volume from the required pool that can be mounted. If a recyclable volume is found, the Solstice Backup software mounts the volume and relabels the volume. As part of the relabeling operation, the Solstice Backup software removes all entries for the save sets and volume that remain in the media database. A new entry for the relabeled volume is added to the media database.

The Solstice Backup software formats the label by obtaining the next volume name from the pool's associated label template. It 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 a plain language label on the outside of the volume as well for easier identification.

When a volume is recycled, the expiration date is automatically reset to the value of the Volume Expiration attribute of the Device resource (two years by default). If an administrator uses the `nsrjb` program to relabel a volume in an autochanger, an expiration date can be specified from the command line.

If no recyclable volumes from the required pool are found, the Solstice Backup software looks for an unlabeled volume. If any unlabeled volume is found, it mounts the new volume and labels it. An entry for the newly labeled volume is added to the media database.

Automatic Mounting of 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 Solstice Backup software uses the following hierarchy 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 Autochanger Volumes

To mount autochanger volumes:

1. Select the appropriate autochanger in the Devices attribute of the Solstice Backup Administrator program.
2. Select Media>Mount to open the Jukebox Mounting window.
3. Select a volume from the Slot attribute and click OK.

The Solstice Backup server automatically mounts the volume without any further prompting. The Devices attribute in the Solstice Backup Administrator program reflects the device's mounted status.

It is a good idea to place an adhesive label on the autochanger identifying its device pathname(s). This reminds you which device pathnames belong to the autochanger. This is especially helpful when you back up to more than one autochanger.

▼ To Mount Uninventoried Volumes

This section describes how to mount a volume that is not included in the autochanger inventory, but is a valid (properly labeled) Solstice Backup volume.

To mount an uninventoried volume:

1. Select the appropriate autochanger device in the Devices attribute of the Solstice Backup Administrator program.
2. Select Media>Inventory to open the Jukebox Inventory window.
3. 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 313](#).

The Solstice Backup software will not inventory the slot you specified. Now you can mount the inventoried volume as described in [“To Mount Autochanger Volumes” on page 307](#).

Note that unlabeled media is not considered “valid” for mounting in this case because media of this type can only be mounted for labeling, not inventorying. An I/O error will occur and the volume will be ejected if you try to mount uninventoried volumes using unlabeled media.

▼ To Unmount Volumes

To unmount volumes:

1. Select the appropriate autochanger device in the Devices attribute of the Solstice Backup Administrator program.
2. Select Media>Unmount to unmount the volume.
3. Select the device with the volume you want to unmount from the Device attribute and click Unmount.

The Solstice Backup server automatically unmounts the volume without any further prompting and return the volume to its appropriate slot. The Devices attribute in the Solstice Backup Administrator program reflects the appropriate device’s unmounted status.

▼ To Label Autochanger Volumes

The Solstice Backup 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 your own. For more information on label templates, see [“How the Solstice Backup Server Uses Label Templates” on page 148](#).

It will take the Solstice Backup server some time to label the volumes in the 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. Select the appropriate autochanger device in the Devices attribute of the Solstice Backup Administrator program.
2. Select Media>Label to open the Jukebox Labeling window.
3. Enter a label in the Starting With attribute.

You can enter your own label only if you are labeling a single volume in the autochanger. If you are labeling a range of volumes, the label will automatically be assigned based on the pool you select.

4. Complete the First Slot and Last Slot attributes:
 - To label a range of slots, enter the range in the First Slot and Last Slot attributes.
 - To label only one volume, enter the same number in both the First Slot and Last Slot attributes.
5. Select a pool in the Pool attribute for the volumes being labeled.

The Solstice Backup server uses the label template associated with the selected pool to label the volumes. Remember, only pools that have been enabled in the Pool resource will appear in the Jukebox Labeling window.

6. Select the Manual Recycle attribute only if you want to recycle the volume manually.

Note: Use this option 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.

When a valid Solstice Backup label already exists on the media you are trying to label, a confirmation box appears asking you if you really want to relabel the volume. The confirmation is important because it keeps you from accidentally relabeling the media. When a volume is relabeled, its contents under the previous label cannot be recovered.

Caution – If you label a specific range of volumes, the name that appears in the Volume Label 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.

Bar Code Labeling of Volumes

Labeling media with bar code labels provides two distinct advantages over non-bar code labeled volumes. Bar code labeling:

- Significantly speeds up volume inventory.
- Provides improved labeling accuracy.

Inventorying with bar code labels is fast and efficient because you do not have to mount the volumes into the device. Instead, the autochanger scans the external bar code labels with an infrared light while the volumes remain in their slots.

Inventorying with bar code labels greatly reduces the time it takes to locate a volume or determine the contents of an autochanger.

Bar code 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 bar code, the Solstice Backup server records and tracks the label in the media database.

The Solstice Backup server uses bar code 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 bar code labels) is impossible because the Solstice Backup server does not know which slot contains the needed volume.

A volume must have a label, but it does not require a bar code label.

Caution – The autochanger includes hardware to read the bar code labels. The bar code information is then forwarded to the Solstice Backup server. Problems reading bar code labels indicate hardware problems. In case of a bar code-related problem, consult your autochanger documentation or hardware vendor.

▼ To Label a Volume with Bar Codes

First, you will enable the necessary bar code attributes in the Jukebox resource.

1. Apply the bar code labels on your volumes.
2. Place the volumes with the bar code labels in the autochanger.
3. Select the appropriate autochanger device in the Devices attribute of the Solstice Backup Administrator program.

4. Select Media>Jukeboxes to open the Jukeboxes window.
5. Select the autochanger whose volumes you want label with bar codes in the Jukeboxes attribute.
6. Select *Yes* in the Bar Code Reader attribute.

If you want the volume label to match the bar code label, select *Yes* in the Match Bar Code Labels attribute.

In either case, you will have two labels: the volume label (which may or may not match the bar code label depending upon your bar code choices) and the bar code label.

7. Click Apply.
8. Select Media>Label to open the Jukebox Labeling window.

If you enabled the Match Bar Code Labels attribute in the Jukebox resource, the Volume Label attribute displays the following label:

| |
|----------------|
| Bar code label |
|----------------|

Otherwise, the Volume Label attribute displays a label template name.

After you complete selecting your bar code label choices and labeling the volume, both the volume label name and bar code label for the volume appear in the Volumes window.

Bar Code Labeling Tips

The volume label and bar code label are both pieces of information that the Solstice Backup server uses to identify volumes. Both labels are recorded in the media index. However, the volume label is also recorded internally on the media (internal volume label). The Solstice Backup server uses bar code labels to inventory volumes and uses the volume labels to identify the volumes needed for backup and recovery. You can require that the volume label matches the bar code label by making the appropriate selections for bar code labeling in the Jukebox resource.

You should not use identical bar code labels for any of your Solstice Backup volumes. Using identical labels defeats the purpose of using bar code labels, which is to speed up the inventory process and ensure labeling accuracy.

Every volume must have a unique volume label. If you try to label a second volume with an identical bar code label and you have enabled Match Bar Code Labels attribute in the Jukebox resource, the Solstice Backup 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 bar code 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 Match bar code labels in the Jukebox resource while labeling the second volume.

It is not necessary to label existing volumes with bar code labels if they are stored in a vault or offsite for long periods at a time. This is because you will not inventory these volumes often, if ever.

If you decide to use bar code labels on your existing volumes, you must first adhere the bar code labels to the existing volumes. Then, load and mount each volume individually, so the Solstice Backup server can match the bar code label with the existing volume label.

You may purchase a variety of bar code labels from a third-party vendor. You can choose numeric labels, alphanumeric labels, or a special combination of numbers and characters to meet your labeling needs. Furthermore, you can order bar code labels that match your current volume labeling scheme.

If you label your 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 bar code labels should be provided with the autochanger hardware documentation. If you have questions about bar code labels, call the hardware manufacturer.

Using a consistent labeling scheme helps you better organize and track your volumes. It also aids the inventory process if all the volumes use bar code labels rather than just a limited number of them.

When the Solstice Backup 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 Solstice Backup software scans the bar code label during the labeling process and updates the media index with the new volume name and its associated bar code label.

▼ To Use Unmatching Volume and Bar Code Labels

If you choose not to match the volume label to the bar code label, you should attach labels to the outside of your media.

1. Apply the bar code labels to your volumes.
2. Place the volumes with the bar code labels in the autochanger.

3. Select the appropriate autochanger device in the Devices attribute of the Solstice Backup Administrator program.
4. Select Media>Jukeboxes to open the Jukeboxes window.
5. Select the autochanger whose volumes you want label with bar codes in the Jukeboxes attribute.
6. Select *Yes* in the Barcode Reader attribute.
7. Select *No* in the Match Bar Code Labels attribute.
8. Click OK.

The Solstice Backup 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 index.

9. Select Media>Inventory to open the Jukebox Inventory window.
10. Click OK.
11. Select Media>Volumes or click the Volumes toolbar button to open the Volumes window.
12. Determine the volume label and bar code label for each volume. Use the information in the Volumes window to match the correct volume labels to the bar code labels. This is the easiest way to make sure that you adhere the correct volume labels to the volumes with the existing bar code labels.

If you enabled the bar code functionality and have forgotten to adhere a bar code label, you receive an error message informing you that a bar code label does not exist.

Inventorying Autochanger Volumes

When the Solstice Backup software labels the contents of an autochanger, it registers the location of the volumes in the autochanger slots when it assigns the volume label. If you do not change the volumes in the autochanger after labeling them, the Solstice Backup 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 Solstice Backup software that the autochanger now holds a different set of labeled volumes or that the volumes are in a different order. *This is called taking inventory.* When you inventory the volumes in the autochanger, the Solstice Backup software

reads the label of each volume and records its slot number. 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.

If you need to label the volumes in the new magazine, you do not need to inventory them because the Solstice Backup software automatically records the slot number in which the newly labeled volume is located.

The Solstice Backup software provides bar code labels for speeding up the inventory process. We strongly suggest that you use bar code labels if you have a large number of volumes or change your autochanger contents often. For more information, see [“Sharing Autochangers Between Solstice Backup Hosts”](#) on page 293.

▼ To Inventory Autochanger Volumes

1. Load the new volumes in the autochanger.
2. Select the appropriate autochanger device in the Devices attribute of the Solstice Backup Administrator program.
3. Select Media>Jukeboxes to open the Jukeboxes window.
4. Enter the range of the slots you want to inventory in the First Slot and Last Slot attributes, or leave the default values as is to inventory all the slots.
5. Click OK.

If your volumes do not have bar code labels, the Solstice Backup 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 Solstice Backup Administrator program displays the results in the Volumes window.

I/O Errors

If you do an inventory of volumes in an autochanger, and a volume is not labeled, the following message can appear:

```
volume tape label read, I/O error
```

The process by which an inventory attempts to verify the label for each volume in an autochanger cannot distinguish between an I/O error caused by reading an unlabeled volume and an I/O error caused by reading a bad volume. Once the volume is loaded into the device, the `nsrjb` program makes a request to `nsrd` to have the volume loaded and its label verified. `nsrd` allocates a `nsrmmmd` session to perform the verify label operation. `nsrmmmd` then attempts to position the volume to read the label and generates the error.

We recommend you do not run an inventory command on an autochanger filled with new unlabeled volumes. Instead, run the following command at the prompt:

```
# nsrjb -HE
```

The `nsrjb -HE` command (hardware reset with initialize element status) synchronizes the state of the library, after which you can label the volumes in the autochanger. At this point, you can run the inventory on the volumes in the autochanger, using either the steps described in this section, or by running `nsrjb -I`.

Using Autochangers with a Volume Import/Export Capability

The Solstice Backup server 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, loading port, and so forth. The import/export feature is the mechanism in an autochanger that deposits and withdraws volumes from slots in the autochanger. 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 media, the element status of the autoloader is invalidated, requiring the long operation of reinitializing the autochanger. However, the Solstice Backup server will not automatically inventory the volume after a deposit.

Note: the procedures in this section do not apply to silos. For more information, see [“Using Silos with a Volume Import/Export Capability”](#) on page 326.

▼ To Deposit Volumes

To deposit volumes:

1. Become root on the Solstice Backup server or storage node computer.
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 following command at the prompt:

```
#nsrjb -d -s slot
```

Slot is the slot number for the volume.

If you do not specify a volume name with the `nsrjb` command, the Solstice Backup 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 Volumes

To withdraw volumes:

1. Become root on the Solstice Backup server or storage node computer.
2. Enter following command at the prompt:

```
#nsrjb -w -s slot -P port
```

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 and Suggestions for Using Autochangers

This section provides additional suggestions to help you use your autochanger and backup media effectively and reliably.

Autochanger Notifications

The Solstice Backup server uses notifications to send a variety of messages about Solstice Backup events. A number of preconfigured notifications are included, which 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.

The Notifications resource displays all of the notifications, including:

- Three tape mount requests.
- Two cleaning cartridge notifications.
- Two device cleaning notifications.

For more information, see [“Notifications” on page 50](#).

The Solstice Backup 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 Solstice Backup Administrator program, which informs you to mount a specific volume.

After you correct an autochanger problem, you might need to mount a volume so the Solstice Backup server can continue to back up or recover files. Check the Solstice Backup Administrator program for messages.

Recycling Versus Adding More Volumes

The Solstice Backup 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 Solstice Backup server will overwrite the data with new backups, yet maintain 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 is automatically changed 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 your site requires that you maintain the backups in storage for a certain 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, we recommend that you estimate the number of volumes needed for your full backups and assign them to the Full pool. This ensures that your 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

Check the autochanger manufacturer's documentation for information about the frequency and method for calibrating the loading mechanism for the device.

▼ To Disable the Element Status

Some autochangers support the element status feature. The Solstice Backup software uses the element status feature to determine whether a slot contains a volume. If volumes have bar codes, the element status feature speeds up the inventory process by eliminating the need to load and read volume labels. If the volume has a bar code label, the Solstice Backup software reads it without loading the volume.

To disable the element status option for your autochanger if your volumes do not have bar codes:

1. Select Media>Jukeboxes to open the Jukeboxes window.
2. Select View>Details to display hidden attributes.
3. Deselect the Elements_Status option in the Jukebox Features attribute.

4. Click Apply. The element status option is now disabled.

▼ To Reset the 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 Solstice Backup server acts as if a slot is empty whereas in reality a volume is loaded in the device. The opposite might be also be true: the autochanger acts as if the slot is full, but in reality the slot is empty. Using the reset feature can return an autochanger with these characteristics back to normal functionality.

To reset the 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, consult your autochanger hardware documentation.
3. Manually reinsert the volume into the cartridge access port, flipping it over so that the “A” side faces up.
4. Reload the volume from the cartridge access port 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 will now reset itself.

Optical Autochangers

If your optical autochanger does not work with the Solstice Backup software, your operating system might not currently include support for the optical autochanger media device.

When you create the Solstice Backup 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 Solstice Backup software manages silos and silo media similarly to autochangers. For best results, you should 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 Solstice Backup 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 Solstice Backup Software Interacts with a Silo

A Solstice Backup server acts as a client of the silo management software, which resides on the silo server. The Solstice Backup server communicates with the silo through the Silo Tape Library Interface (STLI), which must be installed on the Solstice Backup server that uses the silo.

To access the volumes and devices in a silo, the Solstice Backup 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 Solstice Backup 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 actually mounts the volume in the requested device.

The silo management software controls many of the operations that the Solstice Backup 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 Silos

To install a silo for use with the Solstice Backup software:

1. Install the silo management software on the silo server.
2. Install the Silo Tape Library Interface (STLI) on the Solstice Backup server (if necessary).

Do not install the STLI library on the following models, because all the necessary software was installed when you installed the Solstice Backup software:

- StorageTek[®] on Solaris, AIX, HP-UX, and DYNIX/ptx
- Advanced Digital Information Corporation (Adic) EMASS[®]/Grau on Solaris, AIX, HP-UX, and Windows NT
- IBM 3494 on Solaris and AIX

For other silo models, install the STLI library on the Solstice Backup server or storage node that uses the silo. Follow the instructions provided by the silo vendor.

3. Ensure that the Solstice Backup 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 STL Silo).
- Type of silo.
- Name of the silo. Enter any alphanumeric string.
- Host name of silo server.
- Whether to enable automatic device cleaning. Enter `no`. The silo server controls device cleaning in silos.
- Pathname of the STL library software. Enter the full path on the Solstice Backup server where the silo interface library software resides.
- Number of devices to configure.

- For each device you configure, the program prompts you for the pathname of the device, media type of the device (if the device is not already configured), and silo name of the device.

For additional information, see [“To Add Autochangers” on page 282](#).

5. License the Silo Support Module. For general information on licensing, see [“Licensing” on page 39](#).

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 Storage Tek (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 an MVS system. These program 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 Solstice Backup server can use. There is no method to find this information from the Solstice Backup 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 Solstice Backup Administrator program to change the order of the device names in the STL Device Names attribute of the Jukeboxes 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 gets 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 Solstice Backup software, or test to decide 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. 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 Solstice Backup Software

To set up DAS to work with Solstice Backup software, follow these steps:

1. Ask the silo administrator to configure DAS to accept commands from your Solstice Backup server or storage node computer.
2. Ask the silo administrator to either:
 - Use the `dasadmin allocd` command to allocate one or more devices to the Solstice Backup server or storage node.
 - Configure your Solstice Backup server or storage node as an administrator, so you can enter the `dasadmin allocd` command to allocate devices from your Solstice Backup server or storage node computer.

To find the names assigned to the devices in the silo, you can use a utility called `dasadmin` that is supplied with the Solstice Backup software. Take the following steps 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 Solstice Backup 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. Issue 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 Solstice Backup server to recognize and use the device, you must first delete the existing Jukebox resource, and then reconfigure your silo's Jukebox resource to access the new device.

To define a silo device:

1. Log on as Administrator to the Solstice Backup server that uses the silo.
2. Select Media>Jukeboxes to open the Jukeboxes window.
3. Select the silo listing in the Jukeboxes attribute.
4. Click Delete.

When prompted to confirm the deletion of the silo resource, click OK.

5. Click Apply.
6. Run the `jbconfig` program to reconfigure the silo. For more information, see [“To Add Autochangers” on page 282](#). 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 Solstice Backup 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 Solstice Backup server, you have the option of restricting silo access only to the Solstice Backup 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. Select the appropriate device in the Devices attribute in the Solstice Backup Administrator program.
2. Select Media>Release Shared Device.
3. Repeat these steps for each device you want to release.

Shared Devices

The basic elements of device sharing have been implemented in the Solstice Backup software. However, due to current limitations in the SCSI hardware on Solstice Backup platforms, none of the device-sharing functions have been implemented. If you issue device reservation commands, there will be no errors, but the devices are not reserved. The STL Device Sharing and STL Device Reservation attributes in the Solstice Backup Administrator program have no effect on the behavior of the Solstice Backup server.

Silo Device Cleaning

Do not enable automated cleaning for the silo in the Solstice Backup Administrator program. To find out how to clean devices in your silo, refer to the documentation for your silo management software. The automated device cleaning feature cannot be used in a silo because it depends on fixed slot numbers.

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 protect the volumes used by other programs from being overwritten by the Solstice Backup software.

This section addresses how volumes are specified for the Solstice Backup software to use, how volumes are mounted in the devices, and how volumes are identified and counted in a silo.

Silo Slot Numbers

In an autochanger, the Solstice Backup software specifies many of the functions by a slot number. A silo uses the same idea. In an autochanger, there is a fixed number of slots, and the Solstice Backup software uses the slot number to refer to the physical location of a volume.

A silo, however, has a variable number of slots, starting at zero when it is first configured and limited by the silo license you purchased. The fundamental identifier of a silo volume is its bar code, or *volser* (volume serial number). The *volser* never changes over the life of a particular volume.

When the `nsr jb` command displays a list of the contents of a silo, it also displays 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 bar codes that have been allocated. If you allocate additional bar codes that fall earlier in the sort sequence, the slot numbers of all the volumes later in the sequence change. For more information, refer to the *Legato Command Reference Guide* or the `nsr jb` man page.

Because the slot number is not a perfect identifier for 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 Solstice Backup 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 cartridge access port (CAP), mail slot, loading port, and so forth. The import/export feature is the mechanism in a silo that deposits and withdraws volumes from slots in the autochanger.

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 the Solstice Backup software will not automatically inventory the volume after a deposit.

You can use Solstice Backup 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 Solstice Backup Administrator program. Rather, you need to enter the appropriate commands at the command prompt:

- The Solstice Backup command to allocate and deposit volumes is

```
#nsrjb -a -T tags -d
```

- The Solstice Backup command to deallocate and withdraw volumes is

```
#nsrjb -x -T tags -w
```

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 Solstice Backup software to insert volumes.

On StorageTek silos, due to differences between the internal operations of the Solstice Backup software and the StorageTek silo management software, the Solstice Backup software can only withdraw one volume at a time. You must physically remove the volume from the silo's import/export area before any more volumes can be withdrawn.

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 Solstice Backup pool.

For details on mounting and unmounting volumes, see [“Remote Autochanger Management” on page 291](#).

To specify a bar code identifier or template for the volumes, you can use the `-T` option with either `nsrjb` command. For more information, refer to the *Legato Command Reference Guide* or the `nsrjb` man page.

Labeling a Silo Volume

The Solstice Backup labels for volumes in a silo include both a regular Solstice Backup volume label (written on the media of the volume) and a silo bar code identifier. The volume label is usually based on the volume pool's label template. The bar code identifier is written on a physical label on the outside of the volume, which the bar code reader in the silo can scan during inventory. For instructions on how to label silo volumes, see [“To Label Autochanger Volumes”](#) on page 308.

The use of bar codes, and matching bar code labels and Solstice Backup volume labels, are both enabled for a silo. You can change the Match Bar Code Labels attribute, but do not set the Bar Code Reader attribute to *No*. When both attributes are both set to *Yes*, the internal volume label that the Solstice Backup software writes on the media of each volume will match the bar code label on the outside of the volume. When the labels match, it is easier to track volumes, but the Solstice Backup 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” bar code label to those volumes. Although you can use volumes without bar codes, it is very difficult to maintain integrity, since after you remove the volume from the silo, the information about the virtual bar code is lost. Any volume without a bar code can be reinserted into the silo under a virtual bar code that the Solstice Backup software (or another application) associates with some of your data.

Allocating a Silo Volume

When you allocate volumes, you direct the Solstice Backup 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 Solstice Backup software. To prevent this problem, most silo management software includes methods to limit access to volumes based on the host name of the computer on which the Solstice Backup software and the other programs run. The Solstice Backup 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 Solstice Backup 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 Solstice Backup software.

- If the volume does not exist, the following message is displayed:

```
barcode xyz 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 Solstice Backup error.

▼ To Allocate a Silo Volume

To allocate silo volumes:

1. Select the appropriate device in the Devices attribute in the Solstice Backup Administrator program.
2. Select Media>Add Library Volumes to open the Add Library Volumes window.
3. Select the method (Template or List) of entering the names of the volumes to allocate, and then enter their names in the Barcodes attribute.
4. 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 issue the following command:

```
#nsrjb -a -T tags -d
```

On StorageTek silos, the silo management software deposits volumes automatically.

For more information on the silo tape library, see [“How Solstice Backup Software Interacts with a Silo”](#) on page 320.

Deallocating a Silo Volume

If you no longer need an STL volume in a silo for the Solstice Backup 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 Solstice Backup media database remain intact. If you allocate the volume again later, the Solstice Backup software can retrieve the data from it.

Use deallocation when your silo license limits the number of slots you can use, or when you move data offsite for safer storage. If your 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 Solstice Backup server access the data. If you remove the volume from the silo for offsite storage, you must deallocate it with the Solstice Backup software and then use the silo management software to eject it from the silo.

Caution – StorageTek silos currently can eject only one volume at a time. The silo operator must remove each volume before another `nsr jb -x -w` command can be issued. If you deallocate and withdraw multiple volumes, they are all deallocated, but only the first is ejected.

▼ To Deallocate a Silo Volume

To deallocate silo volumes:

1. Unmount the volume from the device. For instructions on unmounting volumes, see [“To Unmount Volumes” on page 308](#).
2. Select the appropriate device in the Devices attribute in the Solstice Backup Administrator program.
3. Select Media>Remove Library Volumes to open the Remove Library Volumes window.
4. Select the method (Template or List) of entering the names of the volumes to deallocate, and then enter their names in the Barcodes attribute.
5. Click OK.

For more information on the silo tape library, see [“How Solstice Backup Software Interacts with a Silo” on page 320](#).

Solstice Backup Cluster Support

This section discusses Solstice Backup support for backing up data that resides on a cluster.

Solstice Backup software supports the Sun Cluster environment.

The main sections in this chapter are:

- [“What Is Clustering?”](#)
- [“Data Backup in a Cluster Environment” on page 334](#)
- [“Data Recovery in a Cluster Environment” on page 336](#)
- [“Solstice Backup Cluster Licensing” on page 336](#)

What Is 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 servers*, each with its own IP address. A virtual *server* can have zero or more shared disks.

Example: General Cluster Configuration

FIGURE 11-1 displays a general cluster configuration consisting of two or more nodes and at least one Solstice Backup 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
- Virtual Server C owns Shared Disks 3




FIGURE 11-1 Basic Cluster Configuration

In this example, the Solstice Backup virtual server can failover between Node A and Node B, however, the Solstice Backup server only runs on one node at a time.

The Solstice Backup server carries various data services and has its own IP address. Although the Solstice Backup 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 Solstice Backup virtual server is running, the Solstice Backup virtual server moves to another node in the cluster.

Cluster Terminology

This chapter uses the term *node* to refer to a physical host. This chapter also uses the term *virtual server* in place of the Sun Cluster term *Agent*.

Types of Cluster Support

A *virtual server has failover* (relocate) capability. In the event of system failure, the virtual server moves from one node to another in the same cluster. This failover process allows another node in the cluster to take over the operations of a failed node and continue to provide service. It picks up where the failed node left off.

There are two types of cluster support in the Solstice Backup environment:

- *Cluster server*

A Solstice Backup cluster server is configured as a highly available application within a cluster.

- *Cluster client*

A Solstice Backup cluster client is either a physical node or a virtual client. A physical client backs up data belonging to a physical node within a cluster. A *virtual* client backs up data belonging to a highly available service within a cluster.

To enable a cluster client, see [“Solstice Backup Cluster Licensing” on page 336](#) for details.

Cluster-aware and Cluster-unaware Solstice Backup Servers

To perform scheduled backups within a cluster, you must have a cluster-aware Solstice Backup server and client.

- If a Solstice Backup server does not have cluster-client support, it cannot recognize cluster clients. When a scheduled backup is performed, the Solstice Backup 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 Solstice Backup client does not have cluster-client support, it backs up all data that belongs to 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

can be used for backups within a Sun Cluster environment. By configuring as an agent within a Sun Cluster environment, it becomes a highly available application. For the complete configuration procedure, refer to the *Solstice Backup Installation Guide*.

A Solstice Backup 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 scheduled backup, the interrupted save operation restarts when the failover has completed, beginning with the interrupted save set.

Refer to the *Solstice Backup Installation Guide* for complete details on:

- How to install in a cluster environment.
- How to migrate an existing Solstice Backup server running Solaris to a cluster server.

Data Backup in a Cluster Environment

This section describes how to backup data from the following:

- Physical or virtual node
- Virtual cluster client

Backing up data from a physical or a virtual node within a cluster is similar to backing up data on a computer that is not part of a cluster. For detailed instructions regarding backing up data with the Solstice Backup server, see [Chapter 3](#).

To prepare for regular backups, install the Solstice Backup client software on all physical nodes within the cluster. For detailed installation and configuration instructions, refer to the *Solstice Backup Installation Guide*.

▼ To Back Up Data From a Virtual Cluster Client

To back up data from a shared disk of a virtual client:

1. Make sure that the root user on each node has administrative privileges for that client.

2. Make sure that the virtual client is defined for the Solstice Backup server. For example:
 - a. Open the nwadmin GUI and select Clients>Clients Setup.
 - b. Click Create.
 - c. In the Name field, enter the name of the virtual client that you want to backup.

Note: With Compaq TruCluster Server 5.0A, you must create a Solstice Backup client using the cluster default alias as the name for this client.

- d. In the Remote Access field, enter the *user@hostname* of all the physical nodes in the cluster. For example:

```
root@clust_phys1
.
.
root@clust_phys4
```

3. Perform a manual save:

Enter the **save** command with the **-c** option, followed by the name of the virtual service whose data you are attempting to backup up. For example:

```
# save -c hostname_name path_name
```

For further information regarding the *save* command, refer to the *Legato Command Reference Guide* or the *save* man page.

Caution – If you manually stop a scheduled backup, the Manual restart option is set automatically, and overrides the Autorestart setting. For a Solstice Backup server to have failover capabilities within a cluster, you must enable the Autorestart setting. Disable the Manual restart option in the Solstice Backup Administration GUI to ensure that the scheduled backup restarts automatically.

The Manual restart option is enabled each time you stop a scheduled backup manually. Consequently, you must disable the Manual restart option each time you have stopped a scheduled back up manually.

Data Recovery in a Cluster Environment

Recovering data backed up from *private disks* 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 backed up from a *shared disk* belonging to a virtual client:

1. Make sure that:
 - You log in to the node which is running the service associated with the virtual client.
 - You are using an account with administrator privileges.
 - The Remote Access attribute of the virtual client that you want to recover the data for, contains the name of the physical host you are using to recover the data. For example:

```
root@...example.com
```

2. 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 client for which you are attempting the recovery. For example:

```
# recover -s server_name -c client_name
```

For further information regarding the `recover` command, refer to the *Legato Command Reference Guide* or the `recover` man page.

Solstice Backup Cluster Licensing

This section discusses Solstice Backup cluster licensing issues. To register and authorize your Solstice Backup product, contact the Sun License Center using one of the following methods:

- World Wide Web: <http://www.sun.com/licensing>
- E-mail: license@Sun.com
- Voice: 1-801-431-1595

- Fax: 1-801-431-3657

If you lose your license, you can request another via e-mail at:

lost-license@sun.com.

Enabling Support for Solstice Backup Cluster Clients

You can enable support for Solstice Backup cluster clients using one of the following modes of editions of the Solstice Backup software:

- Evaluation Mode
- Network Edition
- Power Edition (includes two cluster client licenses)

You can purchase cluster client licenses based on the total number of nodes. You can add more cluster clients using a Cluster Client Enabler, available from your Solstice Backup sales representative.

For specific licensing and installation requirements, refer to the *Solstice Backup Installation Guide*.

Licensing

Several changes have been made to Solstice Backup licensing within a cluster environment:

- You must license all physical nodes in the cluster.
- Existing cluster clients that are using the previous cluster licensing scheme will continue to be licensed, until the next time you upgrade your Solstice Backup client.
- Existing virtual clients that are already licensed will have their current license removed when these virtual clients are licensed under the physical client. This removal process of removing the license occurs during the first backup after the client software has been upgraded.
- Any storage nodes that are configured locally—using a hostname that matches the NetWorker server's physical hostname—do *not* require a separate storage node enabler.

For example, on a server with the following physical hostname: *server_1*, a remote device defined as follows, *rd=server_1:dev/* would not require a separate storage node enabler.

- Under the new Power Edition licensing convention, a Solstice Backup server will have two cluster client connections for the given platform.

If you upgrade a Solstice Backup server with the new Power Edition license, the two cluster client connections are UNIX-specific. Despite the new license, you can still have Windows NT clients within this cluster. The backups of these clients will still function as long as the version of Solstice Backup software on the clients is not updated.

In such a scenario, you will, however, periodically receive a message stating the following:

```
There are more NT cluster clients connected than the number of NT
cluster licenses
```

Once you upgrade the version of Solstice Backup software on the Windows NT clients, the license will be reallocated. Additional cluster client connections will be required for the Windows NT clients.

NDMP Support for the Solstice Backup Server

This chapter provides instructions on how to enable and use the Solstice Backup NDMP Connection.

Introduction

Network Data Management Protocol (NDMP) is a standards-based, storage management client/server protocol. It allows the Solstice Backup server to provide connections to computers with NDMP data modules for tape operations. It also allows the Solstice Backup server software to maintain the client files and media indexes, as well as to control backup operations on a computer with the NDMP service installed.

NDMP Connection Support

The Solstice Backup NDMP Connection allows the Solstice Backup server to communicate with other computers on the network that support NDMP. With the NDMP Connection, you can connect to Solaris computers that have an active NDMP service and an NDMP data module installed. The computers, except for the server, do not need to have the Solstice Backup software installed.

[TABLE 12-1](#) lists the terms used with NDMP.

TABLE 12-1 NDMP Terms

| Term | Definition |
|--------------------|--|
| NDMP | A client/server protocol that allows any storage management application, such as Solstice Backup software, to control backup operations on a computer. |
| NDMP service | A network communication service that uses NDMP. This service can include a variety of additional features, such as backup and retrieval utilities. |
| NDMP data module | An agent used to drive an NDMP service. The Legato Celestra [®] product is an example of an NDMP data module or agent. |
| Application server | The computer where the NDMP service and the NDMP data module reside. This computer controls the data to be backed up. |
| NDMP tape server | The application server with the tape device attached. |
| NDMP computer | Any computer with an NDMP data module or agent installed. |

The NDMP Connection, part of the Solstice Backup server software, is distributed on CD-ROM.

To install an NDMP data module or enable NDMP support on an application server, refer to the respective vendor's documentation.

NDMP Support for the Solstice Backup Server

To add NDMP support to your network, you must:

- Install the Solstice Backup server software on a Solaris computer.
- If you use a tape drive in an autochanger, install the SCSI storage device on the computer where the autochanger software will reside.

NDMP Connection Software and Installation Requirements

Before you enable the NDMP Connection, you must install Solstice Backup server software on a Solaris computer designated for this purpose.

For more information on installing the Solstice Backup software, refer to the *Solstice Backup Installation Guide*.

TABLE 12-2 displays the operating systems supported by this release of the Solstice Backup software with NDMP Connection

TABLE 12-2 Supported Operating Systems

| Platform | Supported as a Solstice Backup Server | Supported as an NDMP Tape or Application Server | Version |
|--|---------------------------------------|---|--|
| Solaris (SPARC) | Yes | Yes | 2.5.1, 2.6, 2.7 (32-and 64-bit) (UFS/VxFS) |
| NetApp [®] (Data ONTAP [™]) | No | Yes | 5.3.5 with NDMP enabled |
| EMC [®] | No | Yes | Celerra v2.1 |

NDMP Client Support

To request backup and recovery services from the Solstice Backup server, Solstice Backup clients must be able to access the Solstice Backup software.

You do not need to install Solstice Backup client or storage node software if you plan to back up application servers through NDMP only and those computers are not Solstice Backup servers. If the application servers are also Solstice Backup servers, the Solstice Backup client and storage node software is installed as part of that installation.

If you plan to back up the computers with standard Solstice Backup backups in addition to the NDMP backups, you must install Solstice Backup client software on each computer and Solstice Backup storage node software on each computer with a tape device.

Silo Support

Silos controlled by either ACSLS or DAS controllers with the NDMP-enabled Solstice Backup server are supported. 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 Solstice Backup server.

Limitations of NDMP

The Solstice Backup software has the following limitations when using NDMP:

- To back up the bootstrap file, you must have a non-NDMP tape device attached to the Solstice Backup server. The bootstrap includes the server index, media index, and configuration files needed for recovering the Solstice Backup server after a disk crash.
- The *incremental* setting is not supported.
- Direct Access Restores are not supported.
- File type devices are not supported as NDMP devices.

Legato Celestra Software and NDMP

Keep in mind the following limitations when you use Legato Celestra software and NDMP:

- Only *full* and *level 1* backup types are supported.
- You must be using Solstice Backup Power Edition.

NetApp Computers and NDMP

Keep in mind the following limitations when you use NDMP and NetApp:

- You cannot recover a NetApp NDMP backup to another NetApp. You can only recover data to the original host.
- You must be using Solstice Backup Network Edition or Solstice Backup Power Edition.

Important Notes on Swap Space

If you want to back up or retrieve a large number of files, you probably need to increase the amount of swap space. Backups and retrievals fail if there is an adequate amount of swap space.

- For backups, the Solstice Backup server uses 100MB of swap space for every one million files.
- For retrievals, the application server uses 350MB of swap space for every one million files retrieved.

Configuring the Solstice Backup Server for NDMP Operations

This section assumes that you have completed installation of your Solstice Backup software.

After installation, you need to perform these tasks:

- Enable and register the NDMP Connection.
- Install and configure an NDMP data module or NDMP-enabled application server to work with the Solstice Backup server.
- Configure your autochanger or tape device.
- Configure the Solstice Backup server for NDMP operations.
- Run a test backup and retrieval to be sure your configuration is correct.

[FIGURE 12-1](#) contains a setup and configuration roadmap.

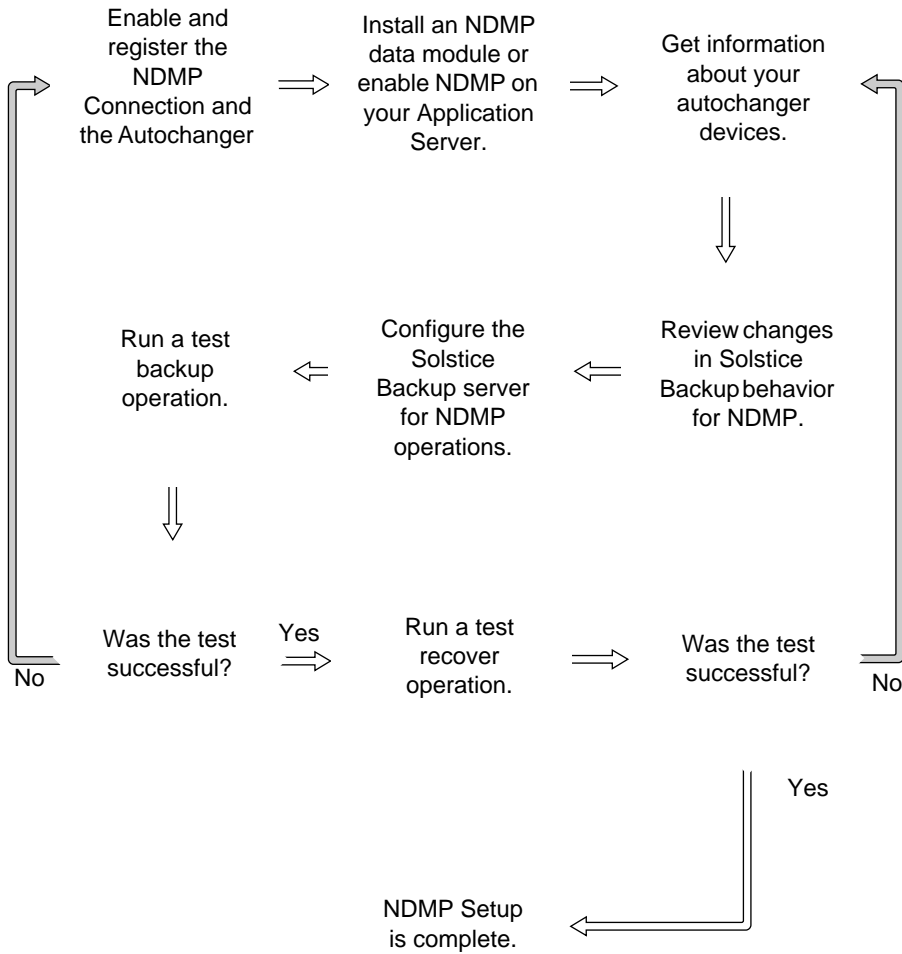


FIGURE 12-1 Setup and Configuration Roadmap

Licensing the Solstice Backup NDMP Connection

Solstice Backup NDMP Connection is an optional product that must be licensed separately from the Solstice Backup base product. Licensing of Solstice Backup NDMP Connection is the same as for other Solstice Backup products. For more information, see [“Licensing” on page 39](#).

Configuring Autochanger Support for NDMP Operations

If you back up data to a tape device attached to the NDMP server, no additional configuration is required. If you back up data to tape devices in an autochanger, you must configure the autochanger for NDMP operations.

You can create a file for NDMP devices, called *ndmpjbconf_NDMPhostname*, that can later be searched by the Solstice Backup `inquire` command.

If you do not create the file, you cannot use the `inquire` command for NDMP devices. You do not need to create the file to set up your autochangers or run the `jbconfig` command.

Obtaining Information About NDMP Autochangers

Before you can configure your autochanger for NDMP operations, you must tell the Solstice Backup server the *NDMP autochanger handle* and the pathname of the media drive.

Before you set up your autochanger or configure the Solstice Backup server for NDMP operations, install Legato Celestra 2.0 software or enable NDMP support on your application server.

- For Legato Celestra devices, review [“To Find the Autochanger Handle and Pathname for Celestra Devices”](#) to use the `ls -l` command to get your autochanger handle names.

- For NetApp devices, review [“To Find the Autochanger Handle and Pathname for NetApp Computers” on page 347](#) to use the `sysconfig` command to get your autochanger handle names.

▼ To Find the Autochanger Handle and Pathname for Celestra Devices

The autochanger handle and device pathname can be identified if you use the `ls -l` command to match the SCSI IDs of the two.

To identify the autochanger handle and media device pathname for a Solaris computer:

1. Enter 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 the autochanger robotics can be identified by the `:8r` at the end.
- The SCSI ID field for the *rsjb* device in this example is `/sjb@1, 0`.
- The autochanger handle is `/dev/rsjb2`.

2. Determine the device filename 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           87 Nov 17 12:03 /dev/rmt/
0mbn -> ../../devices/pci@1f.0/scsi@2/st@1,0:mbn
```

- The SCSI ID field for the media device in this example is `/st@1, 0`.
- The media device pathname is `/dev/rmt/0mbn`.

With the autochanger handle and media device pathname information, you can configure your NDMP autochanger and create the `ndmpjbconf_NDMP` hostname file.

- To create the `ndmpjbconf_NDMPhostname` file, which allows you to use the `inquire` utility to find NDMP devices, continue with [“To Create the ndmpjbconf_NDMPhostname File”](#) on page 348.
- To configure the NDMP autochanger before you create the `ndmpjbconf_NDMPhostname` file, skip to [“To Configure an NDMP Autochanger”](#) on page 350.

▼ To Find the Autochanger Handle and Pathname for NetApp Computers

The autochanger handle and device pathname can be identified with the `sysconfig` command.

The procedure that follows is only an example. Check with your Network Appliance documentation for the correct syntax and usage.

To identify the autochanger handle and media device pathname:

1. Enter the following command at the prompt:

```
# sysconfig -m
```

The computer responds with the devices on your computer:

```
Media changer (6.4)  QUANTUM  Powerstor L200  
  
mc0 - media changer device
```

The autochanger handle is `mc0`.

2. Find the device name by entering the following command at the prompt:

```
# sysconfig -t
```

The computer 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.
 - The media device name is `nrst0`.
3. Verify the information is correct by entering the following command at the prompt:

```
# sysconfig -v

The computers responds with:

NetApp Release 5.3.5: Sun Mar 12 13:29:37 PST 2000

    System ID: 0016784666 (hawaii)

        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
```

With the autochanger handle and media device pathname information, you can configure your NDMP autochanger and create the `ndmpjbconf_NDMPhostname` file.

- To create the `ndmpjbconf_NDMPhostname` file, which allows you to use the `inquire` utility later to find NDMP devices, continue with [“To Create the ndmpjbconf_NDMPhostname File” on page 348](#).
- To configure the NDMP autochanger before you create the `ndmpjbconf_NDMPhostname` file, skip to [“To Configure an NDMP Autochanger” on page 350](#).

▼ To Create the `ndmpjbconf_NDMPhostname` File

The `ndmpjbconf_NDMPhostname` file allows you to use the `inquire` command to get information about the NDMP autochanger. The file is not required to initially set up your NDMP autochanger.

The `ndmpjbconf_NDMPhostname` file is a one-line text file that contains:

- NDMP hostname

- User name
- User password
- NDMP autochanger handle

The format of the file is:

```
server user password autochanger_handle
```

At the end of the configuration, you are notified of the location of the created file.

The user name and password must match the client and device/autochanger resources that you will set up later as you configure the Solstice Backup server for NDMP backups.

- The user name is the NDMP user account name for the computer where the autochanger robotics are installed.
- The password is the password for the NDMP user account for the computer where the autochanger robotics are installed.

To create the *ndmpjbconf_NDMPhostname* file:

1. Make sure you have the autochanger handle and device pathname information for the NDMP autochanger. To get more information, see [“Obtaining Information About NDMP Autochangers” on page 345](#).
2. Become administrator on the Solstice Backup server.
3. Enter the following command at the prompt:

```
# ndmpjbconf
```

The script uses the term *jukebox* instead of *autochanger*. The terms are interchangeable.

4. Provide the information as requested:

```
Enter NDMP Server name: hawaii [Return]

Enter NDMP user name: root [Return]

Enter NDMP user password (characters will not be echoed):

Enter NDMP jukebox handle: /dev/rsjb2 [Return]
```

The configuration files are created in `/usr/lib/nsr`. You are notified of the location when `ndmpjbconf` is finished.

▼ To Use the `inquire` Command with NDMP Devices

Once you have created the `ndmpjbconf_NDMPhostname` file, you can use the `inquire` command to search for available devices on the NDMP server. Two new options have been added to `inquire` to support NDMP devices.

Note: The `inquire` command can retrieve only information in the `ndmpjbconf_NDMPhostname` file. It cannot retrieve device filenames for NDMP devices.

- To get NDMP autochanger information from the location on the local server, enter the following command at the prompt:

```
# inquire -N NDMP_type
```

Where `NDMP_type` is either `standard`, `celestra`, `celestra-hp`, or `netapp`.

- To get NDMP autochanger information on a specific NDMP host when the file is not in the default location, enter the following command at the prompt:

```
# inquire -N NDMP_type -f absolute_path_of_ndmpjbconf_NDMPhostname
```

The `-f` option is only valid with the `-N` option. It cannot be used alone.

If neither `NDMP_type` nor the absolute path of the `ndmpjbconf_NDMPhostname` file are provided, the `inquire` command shows devices on the local Solstice Backup server.

For more information on the `inquire` command, refer to the `inquire` man page.

▼ To Configure an NDMP Autochanger

The Solstice Backup autochanger resource configuration tool, `jbconfig`, helps you configure your autochangers for NDMP operations.

To configure your autochanger:

1. Make sure you have the autochanger handle and device pathname information for the NDMP autochanger. To get this information, see [“Obtaining Information About NDMP Autochangers”](#) on page 345.
2. Become root on the Solstice Backup server.

3. Enter the following command at the prompt:

```
# jbconfig
```

4. Select *Install an Autodetected NDMP SCSI Autochanger* when prompted and respond to the questions. An example that shows how to configure an NDMP autochanger for a NetApp Filer follows.

Note: The script uses the term *jukebox* instead of *autochanger*. The terms are interchangeable.

Note: The example might not be valid for your NDMP server. Refer to the documentation that came with your NDMP server to identify valid device names.

Example: How to Configure an NDMP Autochanger for NetApp Computers

```
# jbconfig
```

```
1) Install a SmartMedia Jukebox.
2) Install an Autodetected SCSI Jukebox.
3) Install an Autodetected NDMP SCSI Jukebox
4) Install an SJI Jukebox.
5) Install an STL Silo.

What kind of Jukebox are you installing? [1] 3 [Return]

Enter NDMP Server name: hawaii [Return]

Enter NDMP user name: root [Return]

Enter NDMP user password (characters will not be echoed):

Enter NDMP Jukebox handle:? mc0[Return]

What is the NDMP type of 'hawaii'?
```

1) One of the Standard NDMP Servers.

2) NetApp or Celestra HP.

Choice? 2 [Return]

Communicating to devices on NDMP Server 'hawaii', this may take a while...

These are the SCSI Jukeboxes currently attached to your system:

1) scsidev@1027.0.0: Standard SCSI Jukebox, Vendor,<ATL>, Product <L500 6320000>:

Which one do you want to install? 1 [Return]

Installing an 'Standard SCSI Jukebox' jukebox

Name you would like to assign to the jukebox device? *atl* [Return]

A brief description of the jukebox device? [scsidev@1027.0.0: Standard SCSI Jukebox [Return]

Pathname of the control port for the jukebox device? [scsidev@1027.0.0] [Return]

Do you want automated device cleaning support enabled? (yes/no) **yes** [Return]

Enter hostname that owns media drive 1: ? *hawaii* [Return]

An example of a valid device name for this OS is 'nrst0'

Enter pathname of media drive 1: ? *nrst0* [Return]

using 'rd=hawaii:nrst0' as device name

Should the drive be configured as a NDMP device? (yes/no) **yes** [Return]

This media device has not been configured yet. Please

```
select a media device type for rd=hawaii:/dev/rmt/0mbn.
```

- | | |
|-------------------|--------------|
| 1) 3480 | 20) dlt8000 |
| 2) 3570 | 21) tz85 |
| 3) 3590 | 22) tz86 |
| 4) 4890 | 23) tz87 |
| 5) 4mm | 24) tz88 |
| 6) 4mm 4GB | 25) tz89 |
| 7) 4mm 8GB | 26) tzs20 |
| 8) 4mm 12GB | 27) tkz90 |
| 9) 4mm 20GB | 28) dst (NT) |
| 10) 8mm | 29) dst |
| 11) 8mm 5GB | 30) dtf |
| 12) 8mm 20GB | 31) himt |
| 13) 8mm AIT | 32) qic |
| 14) 8mm AIT-2 | 33) SD3 |
| 15) 8mm Mammoth-2 | 34) vhs |
| 16) 9490 | 35) file |
| 17) 9840 | 36) logical |
| 18) dlt | 37) optical |
| 19) dlt7000 | |

```
Choice? 19
```

```
Your jukebox does not report that it has a bar code reader, but it
does report that it can handle volume tags. Some jukeboxes that
have barcode readers report this way. If your jukebox has a barcode
reader, do you want bar code reader support enabled? (yes/no) no
```

```
Jukebox has been added successfully
```

```
Would you like to configure another jukebox? (yes/no) no
```

For more information on the `jbconfig` and `nsr_jukebox(5)` commands, refer to the *Legato Command Reference Guide* or the `jbconfig` and `nsr_jukebox` man pages.

NDMP Effects on Solstice Backup Functionality

Adding NDMP to your Solstice Backup datazone environment renders the following changes:

- NDMP operations require two new utilities and changes to a third:
 - `nsrndmp_save`, which controls backups on an NDMP computer, supports most of the same command line arguments as `save(1m)`. Additional options have been added for NDMP support.
 - `nsrndmp_recover`, which controls recover operations on an NDMP computer, supports most of the same command line operations as **recover(1m)**. Additional options have been added for NDMP support.
 - `nsrmmmd`, which adds the `-r hostname` option, indicates the hostname of the NDMP computer.
- You can browse other NDMP computers for recoveries through both the command line and the Solstice Backup Administrator program. You cannot browse NDMP computers for backups.
- Savegroup probes are not supported in the NDMP client resource. You must manually list the filesystems to be backed up with NDMP. You cannot use the default option, *All*, to choose filesystems to be backed up.
- You must create a separate client definition for each tape drive you want to do parallel backups.
- Cloning is not supported.

- Archiving is not supported.
- Multiplexing is not supported.
- Save set consolidation is not supported.
- Storage node failover is not supported.
- The `jbexercise` utility is not supported. To verify that you can find the autochanger on the NDMP server, you can either reset the autochanger with the `nsrjb -HE` command or inventory the autochanger with the `nsrjb -I` command.
- The `scanner` command generates media database entries only for data backed up with NDMP. You cannot recover data with the `scanner` utility or regenerate the file index information. Instead, you must use special tape readers provided by your NDMP vendor.
- NSR resource attributes (`nsrmmmd polling interval`, `nsrmmmd restart interval`, and `nsrmmmd control timeout`) do not apply to NDMP computers.
- The Save Set Recover options *Rename Recovered File* and *Discard Recovered File* are not supported by NDMP. Current files are always overwritten.

Configuring the Solstice Backup Server for NDMP Operations

Before you can back up data with NDMP, you must configure the Solstice Backup server for NDMP operations. Use the Solstice Backup Administrator program to configure resources on the Solstice Backup server before you back up data:

- Server resource
- Client resource
- Group resource
- Schedule resource
- Policy resource
- Pool and a Label Template resource, if applicable

When you configure NDMP, note that certain attributes of the resources are completed differently than for non-NDMP computers.

Client Resource Configuration

For information on creating a Client resource, see [“To Create Clients” on page 90](#). The following section covers issues relating to specific attributes of the Client resource that are affected by NDMP.

- **Save Set:** List the filesystems to be backed up. You must specify the filesystem mount points or the raw partition names here. Savegroup probes are not supported, and the *All* default option will not work.
- **Remote Access:** List all computers that will store and retrieve entries for the client. In addition, list any computer that is used to browse entries for the NDMP computer. Since the Solstice Backup server stores the index entries for the NDMP computer, be sure that *root@server* is listed.
- **Remote User:** Enter the NDMP account user name for the computer where the NDMP data module is installed. If this is *root*, be sure that *root* has a password.
- **Password:** Enter the password for the NDMP user account. This password should match the password for the NDMP user account for the computer where the NDMP data module is installed.
- **Backup Command:** 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 backup:

- Legato Celestra backups – enter **nsrndmp_save -T celestra**
- NetApp backups – enter **nsrndmp_save -T dump**
- Auspex backups – enter **nsrndmp_save -T gtar**
- EMC backups – enter **nsrndmp_save -T pax**
or **nsrndmp_save -T dump**

Application Information

List environment variables that the Solstice Backup server uses for an NDMP backup. For the backups to be successful:

- Separate the environment variable name from its value with an equals (=) sign. For example:

```
Application Information:  
DDIMAGE=n
```

- Put each entry in the Application Information attribute on a separate line.

- Click *Add* after each entry.

The values you enter depend on which NDMP data module or NDMP-enabled application server you use.

TABLE 12-3 lists required values for currently supported NDMP data modules and NDMP-enabled application servers. These values were correct at the time of publication. Check with the vendor for the latest information.

TABLE 12-3 Required Application Information Values

| NDMP Type | Required Values* | Means |
|--|------------------|---|
| EMC [®] | HIST=y | y/n Default is n. |
| | JB=cxtylz | The autochanger device handle. c=controller x=controller number t=target y=target number l=lun z=lun number |
| | TD=cxtylz | The tape device filename. c=controller x=controller number t=target y=target number l=lun z=lun number |
| | FILES=./ | The files that will be backed up. Default is ./. |
| | OPTION=NT | Required if the volumes contain NT data. |
| NetApp [®] | HIST=y | y/n Default is n. |
| *These values were required by the NDMP data module or the NDMP-enabled application server at the time of publication. Check with the vendor 's web site for the most up-to-date values. | | |

- In the NDMP attribute, select *Yes*.
- In the Clone Storage Nodes attribute, do not enter a value. Cloning is not supported.

Device Resource Configuration

For details on how to configure the Device resource in the Solstice Backup software, see “Devices” on page 155. The following section provides Device resource information that pertains to using NDMP.

NDMP devices are created using the following syntax:

```
rd=NDMP_computer:tape_device_pathname
```

Use the following information to complete the fields in the Device Resource window:

- Target Sessions: Set this value to 1. Multiplexing is not supported on NDMP computers.
- NDMP: Select yes if you use NDMP.

This option can be set only when you create the device. You cannot change the NDMP attribute after the device has been created. To change the device configuration, you must first delete the configuration, then recreate it.

- Remote User: Enter the NDMP account user name for the computer where the NDMP autochanger robotics or tape device reside.
- Password: Enter the password for the NDMP user account. This password should match the password for the NDMP user account for the computer where the NDMP autochanger robotics or tape device is installed.

The NDMP Workflow

The workflow for NDMP backup and recover operations is similar to the workflow for traditional Solstice Backup backup and recover operations. Several utilities have been added and several options have been added to current utilities to enable NDMP operations.

The Solstice Backup server with NDMP Connection supports two configurations:

- Local – the backup is done from a tape device attached to the same server where the data resides (FIGURE 12-2)
- Three party remote – the backup is done from one NDMP server to another NDMP server (FIGURE 12-3)

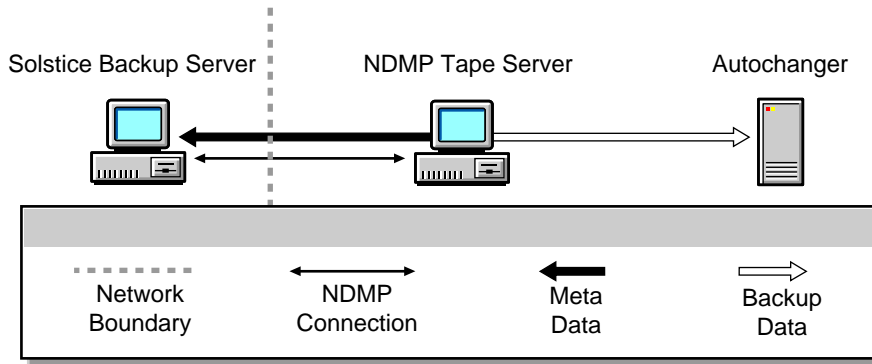


FIGURE 12-2 NDMP Local Configuration

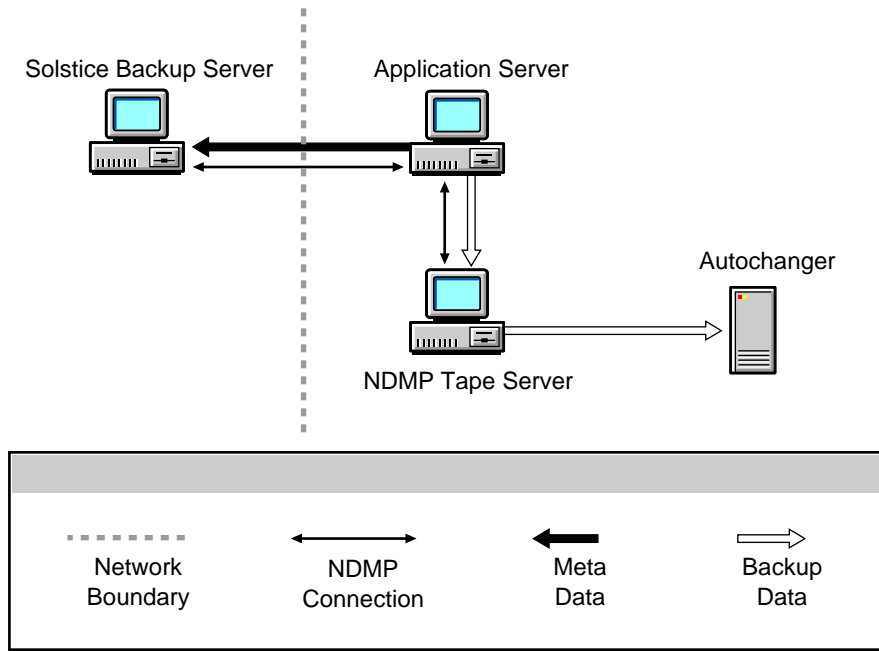


FIGURE 12-3 NDMP Three Party Remote Configuration

NDMP Backup

A new command, `nsrndmp_save`, is used in NDMP operations. This command supports most of the same command line arguments as `save(1m)`. An option within `nsrndmp_save`, `-T ndmp_type`, identifies the NDMP backup type.

The following sections describe how to perform an NDMP backup using the Solstice Backup Administrator program. You can also perform this task from the command line using the `nsrndmp_save` command. For information on this command, refer to the *Legato Command Reference Guide* or the `nsrndmp_save` man page.

Prerequisites

You can perform a test backup after the following prerequisites are met:

- You have enabled the NDMP Connection.
- You have added the NDMP computer to a backup group.
- You have selected a backup schedule with browse and retention policies.
- You have selected a pool for your data.
- You have installed an NDMP data module or enabled your application server for NDMP operations. Filesystems on computers without an NDMP data module cannot be backed up through the Solstice Backup NDMP Connection.

Next you will perform a test backup by starting a scheduled group backup.

▼ To Start a Scheduled Backup

Instructions for a stand-alone device are provided here. To use a device in an autochanger or silo, see [Chapter 10](#).

Media Preparation

The first task is to prepare the media for the backup:

1. Insert a volume into the device you configured for Solstice Backup backups.
2. Select Media>Label or click the Label speedbar button to open the label window. The Label window displays the proposed label name, based on the preconfigured label template you preconfigured for NDMP operations.
3. Select the Mount After Labeling attribute to automatically mount the volume following its labeling and click OK.

The Solstice Backup software will label the volume and mount it in preparation for the test backup of your NDMP data.

NDMP Client Configuration

Next, to configure your NDMP client resource:

1. Select Clients>Client Setup to open the Clients window.
2. Select the NDMP client you created in the Clients attribute.

The attribute values for your NDMP client are displayed.

3. Enter the save sets you want to back up in the Save Sets attribute and click Apply.
To back up all the data on the NDMP computer, list all the filesystems you want to back up.

Savegroup probes do not work with NDMP, and the *All* option for backups cannot be used. You must manually list all the filesystems you want to back up.

Starting the Group Backup

To start the group backup:

1. Select Customize>Groups to open the Groups window.
2. Select the group that contains the NDMP client whose data you want to back up.
3. Select Start Now in the Autostart attribute. and click Apply.

The Solstice Backup server will not begin the group backup that includes your NDMP client.

Next, you need to view the backup details.

Viewing the Group Backup Details

Follow these steps while the backup you just started is in progress.

1. Select Server>Group Control or click the Group Control toolbar button to open the Group Control window.
2. Select the Default entry in the Name attribute. This entry's status should indicate that it is in the process of being backed up.
3. Click Details to open the Group Control Details window. Here the Solstice Backup server displays more detailed messages about the progress of the backup.

When a backup is complete, lists of the successful, unsuccessful, and incomplete save sets are displayed.

Viewing Client File Index Information

Select Clients>Indexes or click the Indexes toolbar button to open the Indexes window after the backup has completed. Here you can view details of the NDMP client backup, and perform other tasks on your client file indexes as well. For general information regarding the management of the client file indexes, see [“Online Index Structures and Management” on page 70](#).

If the backup did not run to completion successfully, refer to the troubleshooting information in [Chapter A](#) to determine the cause.

NDMP Recovery

A new utility, `nsrndmp_recover`, is used for NDMP operations, and it supports most of the same command line operations as `recover(1m)`, in addition to several others that specifically support NDMP.

The `nsrndmp_recover` command performs the following tasks:

- Connects to the appropriate Solstice Backup daemons
- Drives the recover operation on the NDMP computer

If you retrieve data from the Solstice Backup Administrator program, the `nsrndmp_recover` command is selected automatically.

Caution – The Save Set Recover options *Rename Recovered File* and *Discard Recovered File* are not supported by NDMP. Current files are always overwritten.

▼ To Perform a Destructive Save Set Recovery

The following section describes how to perform a destructive save set recovery from the Solstice Backup Administrator program. A destructive save set recovery overwrites the entire contents of the filesystem.

You can also perform this task from the command line using the `nsrndmp_recover` command. For information on this command, refer to the *Legato Command Reference Guide* or the `nsrndmp_recover` man page.

To perform destructive save set recovery:

1. Select Save Set>Recover to open the Save Set Recover window.
2. Enter or select the NDMP client whose save set data you want to recover in the Client attribute.

The save sets available for recovery are displayed in the Save Set attribute.

3. Select the specific save set(s) whose data you want to recover in the Save Set attribute. You can select as many save sets as are listed in this attribute.

When you select a save set, that save set's backup instances are then displayed in the Instances attribute.

4. Select the particular instance of a save sets backups in the Instances attribute.
5. Click Details if you want to view the details of the save set instance you selected, such as the backup date, number of backed-up files, save set ID number, and pool.
6. Click Recover to open the Save Set Recover Status window. Here you can customize the save set recovery in a variety of ways.
7. Specify the UNIX mount point in the Mount Point attribute using the following syntax:

```
file_system_mount_point
```

8. Specify the raw device name of the destination in the Raw Device attribute using the following syntax:

```
destination_raw_device_name
```

Note: If you do not specify the raw device name of the destination, the save set is restored and the files are overwritten, but the current filesystem structure is retained.

9. Click Start.

For additional information on save set recoveries, including recovery options, see ["Save Set Recovery" on page 229](#).

NDMP and NetApp Backups to a Remote Tape Drive

In order to back up a NetApp Filer to a remotely attached tape drive connected to a device other than another NetApp Filer, you can use Celestra Inline Image to act as the NDMP tape server.

Celestra Inline Image supports NDMP three party copying in a non-Celestra NDMP environment, by allowing you to back up data from a NetApp Filer to a Solstice Backup server or a Celestra Inline Image server with tape drives attached. Celestra is

not actually performing the backup. The NetApp Filer is performing the backup while Celestra is functioning as a tape server, accepting the NDMP data stream and writing it to tape.

Note: Backing up data *from* Celestra Inline Image *to* a tape drive attached to a NetApp Filer is not currently supported.

FIGURE 12-4 shows a possible configuration for backing up data on a NetApp Filer to a Solstice Backup server. In this configuration, the data flows from the NetApp Filer, over the LAN, and into the tape library directly attached to the Solstice Backup server. Celestra Inline Image is installed on the Solstice Backup server, acting as the NDMP tape server.




FIGURE 12-4 Performing a NetApp Backup to a Solstice Backup Server

FIGURE 12-5 shows a configuration where the data is backed up to a Celestra Inline Image server. The data flows from the NetApp Filer, over the LAN, and into the tape library directly attached to the Celestra Inline Image server. The Solstice Backup server controls the NetApp backup, while the Celestra Inline Image server acts as a tape server.




FIGURE 12-5 Performing NetApp Backup to a Celestra Inline Image Server

When using Celestra Inline Image, the same values are specified in the Solstice Backup client configurations as for other NDMP backups.

SNMP Module

The Legato NetWorker SNMP (Simple Network Management Protocol) Module provides the ability for Solstice Backup 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 Solstice Backup 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.

SNMP Traps

The Legato NetWorker SNMP Module uses *traps* to communicate Solstice Backup event notifications to SNMP management stations. A trap is an unsolicited notification sent from the SNMP agent (in this case, the Solstice Backup server) to the SNMP event manager.

The types of traps that are sent by the Solstice Backup server are determined when the Legato NetWorker SNMP notification is configured within your Solstice Backup server. Typical traps include warnings, critical errors, and other messages from the Solstice Backup server. For instructions on configuring the Legato NetWorker SNMP notification, see [“Configuring Legato NetWorker SNMP Notifications” on page 368](#).

Additionally, your SNMP-enabled network management software must be configured to accept traps from the Solstice Backup server. Depending on the configuration of your network management software, traps received from the Solstice Backup server may be displayed in the management software’s graphical interface, or may be logged into a file or processed in another manner.

Configuring Legato NetWorker SNMP Notifications

Solstice Backup software provides notifications to a variety of resources about Solstice Backup server events. The Legato NetWorker SNMP Module is one of those resources. The module then forwards the notification to your SNMP management software using the `nsrtrap` program. When you configure the SNMP notification, you include the IP address or host name of your 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 Legato NetWorker SNMP notification, you must first enable the Legato NetWorker SNMP module. For information about enabling the Legato NetWorker SNMP module, see [“Licensing” on page 39](#).

Command-Line Options for `nsrtrap`

The `nsrtrap` program is used by the Legato NetWorker SNMP Module to communicate SNMP traps from the Solstice Backup server to your SNMP management software. The following command line options can be included in Action attribute when the SNMP notification is configured:

TABLE 13-1 `nsrtrap` Command Line Options

| Option | Description |
|---------------------------|--|
| <code>-c community</code> | <p>This option allows you to specify the SNMP community that is authorized to receive traps from the Solstice Backup server. SNMP communities are configured at your SNMP server. The default setting for this option is “public,” meaning that the public community can receive traps from the Solstice Backup server.</p> <p>For security purposes, system administrators often customize their SNMP servers to limit the communities from which their server will accept traps. If a community other than “public” has been configured for your SNMP server, you should include the appropriate community name using this option when you configure your SNMP notification.</p> |
| <code>-t trap_type</code> | <p>This option is used to set the type of trap sent from the Legato NetWorker SNMP Module to your SNMP server. The default setting is 6, which means that this is an “enterprise-specific” trap. Because traps sent from the Solstice Backup server 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 Solstice Backup notification.</p> |

TABLE 13-1 nsrtrap Command Line Options (*Continued*)

| Option | Description |
|-------------------------|---|
| -s <i>specific_type</i> | <p>This option is a generic setting that can be used to identify the type of trap being sent from the Solstice Backup server. 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 Solstice Backup 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 your 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 data-bbox="422 579 836 604">/usr/sbin/nsr/nsrtrap -s 1 host</pre> <p>With this setting, your SNMP management software can be configured to recognize that Solstice Backup traps with the specific trap type of 1 are critical messages. Additional SNMP notifications can have other settings for the -s option to further differentiate various traps from the Solstice Backup server.</p> |
| -v | <p>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.</p> |

▼ To Modify the Preconfigured Legato NetWorker SNMP Notification

The Solstice Backup server has a preconfigured SNMP notification which can be modified if necessary. The only modify 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. Start the Solstice Backup Administrator program.
2. Select Customize>Notifications to open the Notifications window.
3. Select View>Details to display the event and priorities for this notification.
4. Highlight the SNMP Notification.
5. To add command line options for nsrtrap, such as the SNMP community, enter any necessary options to the command listed in the Action attribute. For more information about command line options, see ["Command-Line Options for nsrtrap" on page 368](#).

The events and priorities associated with the preconfigured SNMP notification cannot be modified. If you want to set different events and priorities for your SNMP notification, see [“To Create a Legato NetWorker SNMP Notification” on page 370](#).

▼ To Create a Legato NetWorker SNMP Notification

To create additional Legato NetWorker SNMP notifications:

1. Start the Solstice Backup Administrator program.
2. Select Customize>Notifications to open the Notifications window.
3. Select View>Details to display the event and priorities for this notification.
4. Click Create.
5. Enter “SNMP notification request” for the Name attribute.
6. Enter the path to your Solstice Backup executables, along with `nsrtrap` *network_management_station*, for the Action attribute, where *network_management_station* is the DNS name or IP address of the host on which your SNMP management software is running.

Command line options for `nsrtrap`, such as the SNMP community, should be included in this attribute if necessary. For more information about command line options, see [“Command-Line Options for nsrtrap” on page 368](#).

For example, the Action attribute might read:

```
/usr/sbin/nsr/nsrtrap SNMPhost -c networker
```

7. Select the events and priorities that to be communicated by the notification.

Note: The events and priorities cannot be modified after the notification has been created.

8. Click Apply.

Configuring SNMP Management Software

In order for your SNMP management software to accept traps sent from Solstice Backup servers, it must be configured to recognize the traps. Instructions for configuring SNMP management software to accept specific traps vary depending on the type of management software you are using.

For specific instructions on configuring the specific types of traps that the software will accept, refer to your SNMP management software documentation.

Solstice Backup SMI Network Management Private Enterprise Code

When configuring your management software to accept traps, you must tell it the specific type of the trap that you want it to accept. This is done using the 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 Solstice Backup server is 160. (The complete code is actually .1.3.6.1.4.1.160).

Receiving Traps in Your SNMP Network Management Software

Typically, once your network management software is configured to accept traps from Solstice Backup servers, icons representing each Solstice Backup server are displayed on your network management console. You can then configure your management software to determine what it should do when it receives a trap from a Solstice Backup server. Using your network management software, you may want to:

- Configure the manner that the management software notifies indicates that a trap was received (for example, the Solstice Backup 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. See your network management software documentation for information on setting up SNMP trap templates.

You may also want to create additional SNMP notification schemes in the Solstice Backup Administrator program that have different priorities and events. You can use the `-s specific-type` command line option for `nsrtrap` so that your 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 368](#).

Notes and Troubleshooting Tips

If you have a problem with Solstice Backup software, or if the product does not work the way you expect, use the information in this appendix to diagnose your problem.

Information to Gather 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 version of the Solstice Backup component.
- The version of operating system that you are running. You can determine this with the `uname -a` command.
- Your hardware configuration.
- Information on your devices and other SCSI IDs. Use the `/etc/LGTOuscsci/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 SUNWsbu2`.

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 Solstice Backup command was successful before you made any changes and, if so, the changes you made

Solstice Backup Backup and Recover

This section explains how to troubleshoot various problems you might encounter with Solstice Backup backup and recover operations.

Multiple nsrexec Processes

You might notice multiple instances of the `nsrexec` process running on the Solstice Backup server during a backup. This is normal behavior and happens because one `nsrexec` process is spawned for each active save stream.

Checking the Solstice Backup Daemons

If you have trouble starting Solstice Backup programs, the daemons might not be running properly. To determine whether the required daemons are running, enter one of the following commands at the prompt:

```
# ps -ef |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/nsrmmmd -n 1
12234 ?      S   0:00 /usr/sbin/nsr/nsrmmmd -n 2
```

```
12235 ?          S  0:00 /usr/sbin/nsr/nsrmmmd -n 3
12236 ?          S  0:00 /usr/sbin/nsr/nsrmmmd -n 4
12410 pts/8      S  0:00 grep nsr
```

If the response indicates that the daemons are not present, start the Solstice Backup daemons with the following command:

```
# /etc/init.d/networker start
```

Unsynchronized Client Clock Errors

This situation affects computers with Solstice Backup client software, release 5.0 or later. If the setting for the system clock on a Solstice Backup client differs from that of the Solstice Backup 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 errors 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

During a backup, attempt to stop the process by clicking Stop in the Group Control window. This should stop the process for all clients in the selected group, but sometimes a client is missed. You then see messages that indicate the server is still busy.

To resolve the problem, perform these tasks:

- Determine the group that is currently being backed up by looking at the messages displayed in the Solstice Backup Administrator program.

- Determine the process ID for the currently running **save** process using one of the following commands:

```
# ps -ef | grep save
```

If the group status shows that the **save** processes are running, but the associated **savegrp** process is not running, perform one of these options:

- Stop the conflicting group from running by clicking Stop in the Group Control window. See [“To Stop Groups” on page 189](#) for more information.
- Shutdown and restart the Solstice Backup daemons by performing the following task:
 1. Shutdown the Solstice Backup daemons using the **nsr_shutdown -a** command.
 2. Restart the Solstice Backup daemons using the **nsrd** command.

Or

- Stop the **save** process for each process ID, using the following command:

```
# kill -9 pid
```

Save Set Cutoffs for 5.x Clients

The save set cutoff option is no longer supported in the Solstice Backup 6.x release. When Solstice Backup 5.x or earlier clients back up to a Solstice Backup 6.x server, the server will respect the save set cut-off value. If both the server and client are at the Solstice Backup 6.x release, then the save set cutoff feature is not present.

Licensing-Based Errors with Pre-5.0 Release Clients

Clients running Solstice Backup 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 pre-5.0 client types when **nsrlic -v** is run. Since the server does not know what OS the client is running, it will always use a workstation license, even if the client is on a Windows NT or UNIX server. However, such non-ID clients have no effect on which ClientPak licenses are required to allow Solstice Backup clients to function; the appropriate licenses are still required.

In the following cases, Solstice Backup 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 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.
```

- 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
```

Solstice Backup clients released before Solstice Backup 5.0 identify themselves to the Solstice Backup server at backup time.

Client File Index Check Failure

Each time the Solstice Backup server starts up, it checks the client file indexes residing on the server. If an index check fails for a client, the Solstice Backup server displays the following message:

```
check failed for client_name
```

Should this message appear, follow these steps:

1. Navigate to the directory containing your Solstice Backup software.
2. Enter the following command at the prompt:

```
# nsrck client_name
```

If the client file indexes are undamaged, the Solstice Backup server displays the following message:

```
client_name checked
```

However, if the indexes are corrupt, the Solstice Backup server once again displays the original message indicating that the index check failed for that client. In this case, you need to recover the index for each failed client. See [“Recovering Client File Indexes” on page 238](#) for more information.

No Notification of Client File Index Size Growth

The Solstice Backup 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. See [“Online Index Structures and Management” on page 70](#) for information on how to manage the Solstice Backup client file indexes. Refer to `nsr1s`, `nsrck`, and `nsrim` in the *Legato Command Reference Guide* or the `nsr1s`, `nsrck`, and `nsrim` man pages for more information.

Media Position Errors Encountered When Auto Media Verify Is Enabled

When you enable Auto Media Verify for a pool, the Solstice Backup 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, `nsrmmmd` 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 display in the Solstice Backup 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 Solstice Backup 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" volid 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, attempt 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 your 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 Solstice Backup server.

Then restart the Solstice Backup server manually. See [“Checking the Solstice Backup Daemons” on page 374](#) for commands to restart manually.

The scanner Program Marks a Volume Read-Only

When you use the `scanner` program to rebuild the index of a backup volume, the `scanner` program marks the volume 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

Suppose you attempt to recover indexes to a directory other than the one where they were originally located, and receive the following error message:

```
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 have been recovered to their original location, you can move them to another directory.

Because the indexes are holey files, using the UNIX `cp` command creates a file that consumes more disk space than the original file. To move the indexes, invoke the following command as root from within the `/nsr/index` directory:

```
# uasm -s -i client_index_directory_name | \  
(cd target_directory; uasm -r)
```

Potential Cause for Client Alias Problems

If you encounter any of the following situations, a client alias problem might be the cause:

- You receive the following error message:

```
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 according to the browse and retention policies does not occur. This is indicated by the filesystem containing the indexes continuously increasing 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, `jupiter` and `jupiter.oak.com`

- Sites using both YP (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 put aliases that are shared by other hosts on this command line.

Illegal Characters to Avoid in Configurations

When you upgrade from earlier versions of Solstice Backup software, the configuration names of label templates, directives, groups, policies, and schedules that include the following special characters are no longer allowed:

```
/\ \ * ? [ ] ( ) $ ! ^ ; ' \ " ~ > < & | { }
```

In addition, you cannot include any of the characters reserved for use as field separators, which include the following:

- colon (:)
- dash (-)
- period (.)
- underscore (_)

These characters cannot be used because volume labels, directives, groups, policies, and schedules are often passed as command line options to various Solstice Backup programs.

During installation of Solstice Backup software, these characters in your current configuration names are replaced with an underscore (_) in the resources where they were originally created.

However, in the Clients resource where these configurations are applied, the Solstice Backup server automatically replaces the selected configuration with the preexisting Default configuration. You need to reselect the configurations whose names have changed and reapply them to the individual clients.

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 message

```
please enter record size for this volume ('q' to quit) [xx]
```

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, you see a message prompting 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, you receive the following message:

```
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 the Solstice Backup software for the first time on your system, you receive the error message `nwrecover: Program not found`.

To save disk space, the Solstice Backup software delays the creation of the client index until the first backup is completed. The `nwrecover` program cannot recover data until the client index has entries for browsing. To avoid the problem, perform a Solstice Backup backup on the client.

Client File Index Is Missing Message

If you attempt to recover the client file indexes with the `scanner -i` without first using `nsrck -c` to create a new index, you might encounter a message similar to the following example:

```
scanner: File index error, file index is missing.  
  
Please contact your system administrator to recover or recreate  
the index.  
  
(severity 5, number 8)  
  
scanner: write failed, Broken pipe  
  
scanner: ssid 25312: scan complete  
  
scanner: ssid 25312: 91 KB, 13 file(s)  
  
scanner: done with file disk default.001
```

The scanner program must have a file index to rebuild from before it can proceed. To resolve the problem, always run the `nsrck` program with a `-c` option to create an index before you run `scanner -i` to rebuild them from the backup media.

Recovering Files from an Interrupted Backup

If you terminate a backup by killing the Solstice Backup daemons, you cannot recover the files because the media database is not updated when the daemons die. Consequently, the Solstice Backup 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, you receive the following message:

```
client: save point: There are no save sets in the media database;  
performing a full backup
```

This message indicates that the specified save set(s) have not been previously backed up. If the save set had been backed up, the save set would appear in the media database.

Before an incremental or level backup can be performed on a save set, it must first be backed up in its entirety.

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 Solstice Backup 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 previous backup data under the new client name:

1. Delete the Client resource configured for the old client name.
2. Create a new Client resource for the new client name.
3. Shut down the Solstice Backup daemons by entering the following command at the prompt:

```
# nsr_shutdown
```

4. Delete the index directory that was automatically created for the new client. (If you simply copy the new client index over the old client index directory, the result is a nesting of the new client index inside the old client index directory.)
5. Enter the following command to rename the old client's file index directory:

```
# mv /nsr/index/old_client /nsr/index/new_client
```

Disk Label Errors

If you receive the error message:

```
No disk label
```

you might have incorrectly configured a nonoptical device as an optical device within the Solstice Backup server. Verify that the Media Type attribute in the Devices resource matches the expected media for your device, and make corrections if needed.

Cannot Print Bootstrap Information

If your server bootstraps are not printed, you might need to enter your printer's name as a hidden attribute in the Groups resource.

1. Select View>Details to display hidden attributes.
2. Enter the name of the printer where you want the bootstraps to be printed in the Printer attribute of the Groups resource.

Server Index Not Forced

If your Solstice Backup server belongs to a group that is not enabled or 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 your 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 Solstice Backup software on more than one server and used the same Solstice Backup enabler code for them all, you receive messages similar to the following in your save group completion mail:

```
--- Unsuccessful Save Sets ---

* mars:/var save: error, copyviolation - servers 'jupiter' and 'pluto'
have the same software enabler code, 'alb2c3d4f5g6h7j8' (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, you must issue the `nsr_shutdown` command on each server, remove the Solstice Backup software from the extra servers, and then restart the Solstice Backup 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 Solstice Backup Administrator program from a client machine, it means that the client is not authorized to display Solstice Backup GUIs:

```
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 Solstice Backup GUIs:

1. From the client machine, invoke the `xhost` command:

```
# xhost server_name
```

2. Remotely log in to the Solstice Backup server and issue the `setenv` command at the prompt.

```
# setenv DISPLAY client_name:0.0
```

For command shells other than `csch` enter:

```
# DISPLAY=client_name:0.0
```

```
# export DISPLAY
```

Converting Sparse Files to Fully-Allocated Files

The Solstice Backup 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.

At this time, 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 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.

Client/Server Communications

Many of the problems that Solstice Backup users report when they set up and configure Solstice Backup software are problems with the communications in their networks. This section contains a procedure for testing the communications in a network.

The information provided here is for UNIX platforms only. If your Solstice Backup setup includes other platforms, refer to *Legato Technical Bulletin 299*, which is included in the `bulletins.pdf` file that shipped with your Solstice Backup binaries.

▼ To Troubleshoot IP Errors

If one of the following error conditions occur, you probably have an IP naming problem:

- RPC errors
- Unknown host messages
- Failure with contacting the portmapper
- Connection failures or timeouts
- Program unexpectedly exits
- Refused connections
- RCMD to active client fails
- Name to address translation fails
- Program not registered messages
- Solstice Backup services or daemons fail to start
- Solstice Backup services or daemons fail to remain active
- Messages about an invalid path

To troubleshoot IP errors:

1. Read this section. 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 hosts table for each Solstice Backup client and Solstice Backup server. See [“To Set Up a Hosts Table” on page 389](#) for more information.

3. Disable other name servers to simplify testing. See [“Troubleshooting Hostname Resolution Problems” on page 390](#) for more information.
4. Use **ping** to establish basic connectivity. See [“To Use ping to Verify Network Connections” on page 391](#) for more information.
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 392](#) for more information.

▼ To Set Up a Hosts Table

It is recommended that you troubleshoot IP problems using only the hosts table. Troubleshooting using only the hosts table does not mean you cannot use your name service, for example, DNS, with Solstice Backup software. Run tests using only the hosts table to determine whether you have the software installed correctly. After you know the software works with the hosts table, you can enable the named server you are using.

1. On the Solstice Backup client, list the client and the Solstice Backup 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 Solstice Backup server, list the Solstice Backup server itself 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 391](#) to ensure the highest success rate for parsing a hosts table within any operating system.

Recommendations for hosts table configuration:

- Do not use blank lines in the body of your hosts table.
- The end of the hosts 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 Steps 1 and 2.
- The last character of each unremarked line should be a space, not a carriage return.

On UNIX platforms, the hosts table resides in the `/etc/hosts` directory.

You can use the hosts table in addition to DNS where necessary, but it is simplest to temporarily disable DNS for troubleshooting.

Troubleshooting Hostname Resolution Problems

To simplify the troubleshooting of name resolution problems, try disabling services like DNS, WINS, and DHCP. If you have name resolution problems, first configure the hosts table for each computer, and then test your backups.

Some common problems you can encounter with DNS, WINS, and DHCP services include the following:

- 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 your entire network, just for the initial setup of the Solstice Backup clients and the Solstice Backup 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 hosts 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 search order, with DNS second and NIS last, for example:

```
hosts: files [NOTFOUND=continue] DNS [NOTFOUND=continue] nis
```

You can also set the NSORDER environment variable. Refer to Info Explorer for directions specific to your version of AIX.

▼ To Use ping to Verify Network Connections

After you create the hosts table, test it using **ping**. Use just the steps marked with an asterisk (*) if the server is the only client.

On the Solstice Backup client run the following tests with the **ping** command:

- **ping** the client short name (hostname) from the client.
- **ping** the client long name (hostname plus domain information) from the client.
- **ping** the client IP address from the client.
- **ping** the server short name from the client.
- **ping** the server long name from the client.
- **ping** the server IP address from the client.

The following example shows how to **ping** the client short name and client long name from a Solstice Backup client called *mars* in the *oak* domain:

```
# ping mars  
  
# ping mars.oak.com
```

On the Solstice Backup server run the following tests with the **ping** command:

- **ping** the server short name from the server.
- **ping** the server long name from the server.
- **ping** the server IP address from the server.
- **ping** the client short name from the server.
- **ping** the client long name from the server.
- **ping** the 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 Solstice Backup 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 OSI model; `rpcinfo` checks for communication up to the session layer.

Use the same tests with `rpcinfo` as with `ping`. Run just the steps marked with an asterisk (*) if the server is the only client.

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 Solstice Backup software works with the rest of your 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 typing `rpcinfo` at the command 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 parent nsrexecd spawns a child nsrexecd that handles the portmapper functions. As a result, you will 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 parent and the instance of nsrexecd running under the PID 20416 is the child:

```
jupiter# ps -ef | grep nsr

root 20429 204200 11:23:59 ? 0:00 /usr/sbin/nsr/nsrmmmd -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 Solstice Backup processes to function in a more secure manner that does not utilize ports or protocols that are commonly used by other programs.

▼ To Verify Firmware for Switches and Routers

If you are using switches or routers from any vendor, make sure that the switch or router firmware is dated after August 1995 (wherever they exist on your network) to ensure that RPC (Remote Procedure Call) traffic is handled properly. Most of the switch and router vendors have significantly improved their handling of RPC traffic since August 1995.

Naming Requirements

Solstice Backup UNIX clients, release 4.2 and later, use the `servers` file in the `/nsr/res` subdirectory to determine whether a Solstice Backup server is authorized to back up the client's data. If you don't have the `servers` file, you can create it in `/nsr/res` using your preferred editor.

Make sure the `servers` file on a client contains both the short name and long name of the server you want to use to back up that client's data. For example, the `servers` file on a Solstice Backup client would contain the following names for a Solstice Backup server named `mars` in the `oak.com` domain:

```
mars
mars.oak.com
```

In the Clients resource, list both the short name and the long name, plus any other applicable aliases for each client, in the Alias attribute.

Binding to Server Errors

Solstice Backup 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.

Solstice Backup 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 Solstice Backup daemons are not running and a Solstice Backup service is requested, you receive the following messages in your save group completion mail:

```
Server not available
RPC error, no remote program registered
```

These messages indicate that the Solstice Backup daemons `nsrd`, `nsrexecd`, `nsrindexd`, `nsrmmmd`, and `nsrmmdbd` might not be running.

To restart the daemons, become root and enter the following command at the prompt:

```
# /etc/init.d/networker start
```

Saving Remote Filesystems

You might receive the following error messages in your 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. Refer to the `nsrexecd(1m)` man page for detailed information about the `nsrexecd` daemon.

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 by configuring the Client resource. The Remote Access list displays the usernames that have recover access to the client's save sets. You can add or remove usernames depending on the level of security the files require.

The following users have permission to recover any files on any client, regardless of the contents of the Remote Access list:

- Root
- Operator
- Member of the operator group

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.

Networker Archive and Retrieve

This section explains how to troubleshoot various problems you might encounter with Networker archive and retrieve.

Remote Archive Request from Server Fails

If you cannot perform a remote archive request of a workstation from the Solstice Backup server, the archive client's user name (for example, `root`) might not be listed in that client's Archive Users attribute in the Clients resource.

You can also grant Solstice Backup administrator privileges for `root@client-system` in the Administrator attribute in the Server resource. Granting administrator privileges creates a potential security issue, because Solstice Backup 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 retrieve, archive them separately.

Cloned Archives Do Not Appear in Solstice Backup 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 `mm:info` 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 the one 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 the problem, do not create two archive requests with the same name; the newer one will never be executed.

Command Line Archive Does Not Start Immediately

If you run `nsrchive` from the command line, the archive does not start immediately after you type 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.

Empty Annotations in Retrieve List

You might encounter empty annotations in the retrieve list when you search for annotations using a search string.

The UNIX Networker Archive application does not allow you to enter a null annotation string. By contrast, older versions of the Networker Archive application software installed on DOS, Windows, and NetWare lack an annotation feature. As a consequence, the annotations for save sets archived with the older software are empty strings in the retrieve list.

Storage Nodes

The following sections describe errors associated with storage nodes.

Storage Node Affinity Errors

If a backup fails, and the following message appears:

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

the problem may be related to storage node affinity.

Here are some possible reasons:

- No devices are enabled on the storage nodes in the Storage Nodes list.
- 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 in this example, do one of the following:

- Enable devices on one of the storage nodes on the client's list.
- Correct the pool restrictions for the devices on 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.
- 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, you will receive the following message:

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

If you receive this message, follow these steps:

1. Shut down the Solstice Backup daemons.
2. Remove the environment setting for `NSR_MMDCONTROL`.
3. Restart the Solstice Backup daemons.
4. Start the Solstice Backup administration program.
5. Adjust the value of `Nsrmmmd Control Timeout` to the setting that was previously assigned to the `NSR_MMDCONTROL` variable, or one that best meets your current requirements.

Devices and Autochangers

This section explains how to resolve problems encountered with the use of an autochanger with Solstice Backup software.

New Attributes in the Jukebox Resource

The new Jukebox resource attributes provide a detailed view of options that the `nsrjb` program uses. These are hidden attributes; select `View>Details` from the Jukeboxes window to display them.

The new attributes are:

- `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 Solstice Backup software on the server, depending on whether you updated your storage nodes to the same Solstice Backup release as the server.

- Jukebox Features

Displays the features that your 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.

- Eject Sleep

Defines the number of seconds an autochanger should remain idle (“sleep”) after an eject operation is completed.

- Cleaning Delay

Defines the number of seconds an autochanger should remain idle before attempting to unload a cleaning cartridge.

- Unload Sleep

Defines the number of seconds an autochanger should remain idle after an unload operation is completed.

- Load Sleep

Defines the number of seconds a jukebox should remain idle after a load operation is completed.

- Deposit Timeout

Defines the number of seconds a jukebox should wait for a deposit in the mail slot before the `nsrjb` 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

Solstice Backup 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 their corresponding man pages.

TABLE A-2 Autochangers Maintenance Commands

| Command | Description |
|----------|---|
| lusbinf | Prints out SCSI information |
| lusdebug | Sets the library debugging level |
| lrescan | Rescans for devices |
| lreset | Resets the SCSI bus |
| changers | Lists the SCSI autochangers attached to the system |
| hpflip | Flips the device type of HP optical disk drives |
| 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 |
| sjiielm | Tests the standard jukebox interface (SJI) SJIIELEM command |
| sjiinq | Tests the SJI SJIINQ 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 Solstice Backup device driver software, use the `lusdebug` program to verify the server connection and the `jbexercise` program to test the autochanger. Use the value of the control port assigned to your autochanger (for example, `scsidev@0.6.0`) for *control-port* in the following commands:

```
# lusdebug control_port 0  
  
# jbexercise -c control_port -m model
```

If these commands fail or if you receive error messages, see the following sections for information on the possible cause and solution.

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. You should receive a message similar to the following:

```
scsidev@0.6.0:<EXABYTE EXB-10i EXB-10i >
```

Verify that the information supplied by the message is correct.

If the vendor and model names are incorrect, you supplied the wrong SCSI ID as 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 as it is 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 workstation 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 AAP driver shipped with earlier versions of Solstice Backup 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 AAP driver are available in Technical Bulletin 142, which is in the `bulletins.pdf` file included with the Solstice Backup software.

- 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, you can see the following symptoms: SCSI bus reset errors appear in system log files, the machine does not boot, and 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, follow the instructions provided with your autochanger hardware to determine the problem, or contact your hardware vendor.
- If the autochanger is in sequential mode, change the setting to random mode.

If none of these suggestions resolve the problem, contact Sun Technical Support. Provide the information described in [“Information to Gather before You Call Technical Support” on page 373](#) and the captured output of the `jbexercise`, `sjiinq`, and `sjirjc` programs. Refer to the *Legato Command Reference Guide* or the `jbexercise`, `sjiinq`, and `sjirjc` man pages for more information.

The jbxercise Command Fails

If the `jbxercise` command fails, review the following list of suggestions to identify potential problems and their solutions:

- The `jbxercise` 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 you receive the following error message, you did not enter a nonrewinding device name:

```
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. Contact your hardware vendor for further information on how to configure your tape drive to work with your system.

If none of these suggestions resolve the problem, contact Sun Technical Support. Provide the information described in [“Information to Gather before You Call Technical Support” on page 373](#) and the captured output of the `jbxercise`, `sjiinq`, and `sjirjc` programs. Refer to the *Legato Command Reference Guide* or the `jbxercise`, `sjiinq`, and `sjirjc` man pages for more information.

Auto-detected SCSI Jukebox Option Causes Server to Hang

If you install an Auto-detected SCSI jukebox using `jbconfig` and the server hangs, the following workaround is recommended:

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 you receive the following message:

```
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 is displayed:

```
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

If any of the following situations occur:

- The media is manually ejected from the autochanger drive
- The media is removed from the autochanger
- The autochanger door is opened

The autochanger inventory becomes outdated, which means that Solstice Backup software cannot use the autochanger.

To make the autochanger usable again, perform the following steps:

1. Verify that the media cartridge is correctly installed in the autochanger and that the autochanger door is closed.
2. Become root on the Solstice Backup 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 Solstice Backup 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

The message “Destination component full” usually is the result of a manual operation performed on the autochanger, for example, physically unloading the tape drive by means of the buttons on the autochanger rather than using the Solstice Backup server to unmount the volume. This operation causes the server to lose track of the status of the media in the autochanger.

To resolve the problem, use the Solstice Backup command `nsrjb -H` to reset the autochanger.

Tapes Are Not Filled to Capacity

You might encounter situations where the Solstice Backup software does not fill tapes to capacity. For example, a tape with an advertised capacity of 4000 MB can be marked full by the Solstice Backup server after only 3000 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 your device. When a tape is labeled, the server writes to it at the highest density supported by your device.

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 Solstice Backup server marks the tape full.

To prevent tape write errors, clean your 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.
- Solstice Backup filemarks take up space on the tape.

The Solstice Backup 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 MB. 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 your 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.

Refer to Technical Bulletin 176, available in the `bulletins.pdf` file included with your Solstice Backup software distribution, for more detailed information.

For Solaris, if your tape devices are not directly supported by Sun Microsystems, you will need to recreate your 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. Your autochanger's hardware installation manual should have instructions on how to verify whether the control port is properly connected. If you cannot determine whether the control port is working, contact the autochanger vendor for assistance.

Nonrewinding Device Requirement

You must use a nonrewinding device for Solstice Backup backups. The Solstice Backup server writes a file mark 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 file mark. If the device automatically rewinds the data, the file mark 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 pre-5.5 release of Solstice Backup software, the following must be done to allow the ATL 2640 library to be used:

▼ To Set Jukebox Attributes

The *polled_import* and *stacker_export* jukebox attributes must be set. Use the Solstice Backup Administrator program to set these attributes.

To set jukebox attributes:

1. Select Media>Jukeboxes to open the Jukeboxes window.
2. Select View>Details to display hidden attributes.
3. Select the *Polled_import* and *Stacker_export* options in the Jukebox Features attribute.
4. Click Apply.

▼ To Configure the ATL Library

The ATL library must be configured to enable User Op IOD. This must be done to allow the IOD door to open to insert tapes.

To configure an ATL library:

1. Connect a DOS/Windows computer to the library DIAG serial port.
2. Load the ATL diagnostic program *tapelib.exe* on the computer.
3. Select Config>User Op IOD, 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 Solstice Backup 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** now enables you to specify a range of volume names already contained within the media database.

To deposit a labeled volume that is a member of the Solstice Backup media database into the autochanger, specify only the names of the volumes to be deposited and not the associated port or slot number. The Solstice Backup server will deposit 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 perform an inventory after depositing labeled volumes that are also members of the Solstice Backup media database to a jukebox, make sure the following requirements are met:

- 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 slot(s) -P port(s) volume_name(s)
```

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 will be ignored. For example, in the following example, the `-o notreadonly` task will not be executed:

```
# nsrjb -o notreadonly -o notfull -S 2
```

To perform multiple `nsrjb` tasks with the `-o` option, enter the commands separately. For example, the commands in the previous example must be entered as follows:

```
# nsrjb -Y -o notreadonly -S n  
  
# nsrjb -Y -o notfull -S n
```

Environment Variables Needed for Emass/Grau and StorageTek Silo

For an Emass/Grau and StorageTek silo, edit your startup files for Solaris, HP-UX, and AIX. [TABLE A-3](#) describes where the startup files are located.

TABLE A-3 Startup Files

| Operating System | Startup File Location |
|------------------|------------------------|
| Solaris | /etc/init.d/networker |
| HP-UX | /sbin/init.d/NetWorker |
| AIX | /etc/rc.nsr |

Locate the following lines in the startup file for your operating system:

```
'start')  
(echo 'starting NetWorker daemons:') > /dev/console
```

Enter the lines appropriate for your particular silo, as shown in [TABLE A-4](#).

TABLE A-4 Environment Variables

| Silo Model | Lines to Enter |
|------------|--|
| Emass/Grau | <pre> DAS_SERVER = <i>name_of_DAS_server</i> export DAS_SERVER DAS_CLIENT = <i>name_of_system_as_defined_to_DAS_server</i> export DAS_CLIENT DAS_MEDIUM = <i>type_of_tape_drive_used</i> export DAS_MEDIUM ACI_MEDIA_TYPE = <i>type_of_tape_drive_used</i> export ACI_MEDIA_TYPE </pre> <p>Note: For DAS_MEDIUM and ACI_MEDIA_TYPE, use one of the following values:</p> <ul style="list-style-type: none"> • 3480 • OD-Thick • OD-Thin • DECDLT • 8MM • 4MM • D2 • VHS • 3590 • CD • TRAVAN • DTF • BETACAM • AUDIOTAPE |
| StorageTek | <pre> CSI_HOSTNAME = <i>name_of_ACCLS_system</i> export CSI_HOSTNAME </pre> <p><i>/solstice_backup_binaries_path/mini_el</i> & <i>/solstice_backup_binaries_path/ssi</i> &</p> |

Diagnostic Tools

A variety of diagnostic tools are available as operating system services and as part of the Solstice Backup product. This section describes some diagnostic tools that are useful with Solstice Backup software.

Diagnostic Report

Solstice Backup software includes a script called `nsr_support` that generates an exhaustive diagnostic report. Typically, you run `nsr_support` only at the request of Sun Technical Support. Redirect the output of the script to a file and then e-mail the file for analysis.

To run the script and redirect the output, log in as root on the system and enter the `nsr_support` command at the prompt:

```
# nsr_support > /tmp/filename
```

Refer to the *Legato Command Reference Guide* and the `nsr_support` man page for more information.

Communications Tests

To verify that communications sessions can be established, test with **ping** and **rpcinfo**, which are tools provided with the operating system software.

Because the Solstice Backup 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 388](#).

Contact Sun Technical Support for more tools on testing communications.

UNIX Platform-Specific Notes

This appendix address Solstice Backup notes regarding features and functionality specific to one or more UNIX platforms.

Solaris

This section provides information specific to Solstice Backup software running on the Solaris platform.

Solstice Backup Executables Not Found in Expected Location for Solaris Client

On Solaris, Solstice Backup executables are installed by default in `/usr/sbin/nsr`. If you start a group backup on a Solstice Backup server that does not have `/usr/sbin/nsr` in the search path for root, the backup fails on a client that has its Solstice Backup executables in `/usr/sbin/nsr`. This is because the `savefs` command is not in the search path.

As the best solution, set the Executable Path hidden attribute for a client that has this problem. To set the Executable Path, display the Clients attribute in details view and enter the path of the executables, `/usr/sbin/nsr`, in the Executable Path attribute.

As another solution, modify the search path for root on the Solstice Backup server to include `/usr/sbin/nsr` even if it does not exist locally.

HP-UX

This section provides information specific to Solstice Backup software storage node software running on the HP-UX platform.

Autochanger Installation on an HP-UX System

The following section explains how to install and configure the drivers supplied by Hewlett-Packard.

Select 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 your autochanger, become root on the Solstice Backup server or storage node and enter the `iocan -f` command.

SCSI addresses range from 0 to 7. However, SCSI address 7 is reserved for the SCSI interface board in your computer, so you cannot use 7. 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 `iocan` 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 Pass-Through Driver

There are two procedures for installing a pass-through driver, each depending on the type of hardware you are installing:

- If you are installing a pass-through driver on HP 9000 Series 800 containing an HP-PB interface card, go to the next section, [“To Install an HP-PB or NIO Pass-Through Driver” on page 417](#).

- If you are installing a pass-through driver on HP 9000 Series 700 or 800 containing an HSC or PCI interface card, go to the section, [“To Install a GSC, HSC, or PCI Pass-Through Driver”](#) on page 420.

▼ 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` (man 7 `scsi_pt`).

To install the HP-PB or NIO pass-through driver:

1. Run 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:

TABLE B-1 Output of `ioscan -f`

| Class | l | H/W Path | Driver | S/W State | H/W Type | Description |
|---------|----|------------|--------|-----------|-----------|---------------------------|
| ext_bus | 7 | 4/8 | c720 | CLAIMED | INTERFACE | GSC add-on 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 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 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, in this format:

```
driver path spt
```

In this command, *path* is the complete hardware path of the desired device (from the `ioscan` command).

If you were using 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> stape** might prevent the system from recognizing tape drives added to that SCSI bus. You might have to add lines for **driver <path> stape** for each tape drive that your system was previously able to recognize.

6. Save the old file system:

```
# 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:

```
# ioscan -f -C spt
```

Output from the `ioscan` should resemble the following:

TABLE B-2 Output of `ioscan -f -C spt`

| | 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 the section, [“To Create a Device File” on page 421](#).

▼ 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. Scroll to SCTL and select it. 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 421](#).

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 menu bar, select Actions>Add Drivers to Kernel and press [Return].
6. From the menu bar, select Actions>Create a New Kernel and press [Return].
7. Respond “Yes” to the prompt `Are you sure?` 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 421](#).

▼ To Shut Down the System

Before installing the autochanger device hardware, shut down your system.

To shut down the server:

1. Halt the server by entering the following command:

```
# sync  
  
# shutdown -h time
```

The time option included in this shutdown command 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 main power switch of the server to the “Off” position.

Installing the Autochanger Hardware

To install the autochanger hardware, follow the instructions in the appropriate hardware installation manuals, as provided by your vendor. Use the SCSI addresses that you selected in the earlier section, [“Select SCSI Addresses for the Autochanger” on page 416](#), 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 your vendor.

▼ To Start the System

1. Turn the main power switch of the autochanger to the “On” position.
2. Turn the main power switch of the server to the “On” position. This reboots the server, as required after installing new autochanger hardware.

▼ To Create a Device File

1. Verify that the SPT has been successfully installed with the following command:

```
# ioscan -kfn
```

2. Verify that the autochanger has been claimed by the driver. If the driver has been claimed, CLAIMED should appear under the S/W State entry. 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
```

In this command:

- `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 this format: **0xIIITL00**
 - **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` complete one of the following commands:

- If the robotics is attached to an HP-PB adapter, complete this command:

```
# lsdev -d spt
```

- If the robotics is attached to an HSC or PCI adapter, complete this command:

```
# lsdev -d sctl
```

The output should resemble the following. The assigned number may differ from those displayed in this example:

TABLE B-3 Output of `lsdev -d`

| | Character | Block | Driver | Class |
|------------|-----------|-------|--------|-------|
| HP-PB | 75 | -1 | spt | spt |
| HSC or PCI | 203 | -1 | sctl | ctl |

The value for `majornum` is the number under the “Character” field.

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 will contain HP C6280-7000 in the Description column).
- Pertaining to the adapter to which the controller is connected (which will be the second line above the line for the controller and contains “ext_bus” in the class field). For example, if the robotics is attached to an HP-PB adapter card, the `ioscan` should resemble the following:

TABLE B-4 Output of `ioscan` When Robotics is Attached to an HP-PB Card

| Class | I | H/W Path | Driver | S/W State | H/W Type | Description |
|---------|----|------------|--------|-----------|---------------|-------------------------|
| ext_bus | 12 | 10/4/4 | scsi3 | CLAIMED | INTERRAC E | HP 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 example above, the correct `mknod` command would be:

```
# /usr/sbin/mknod /dev/sjidev_filename c 75 0x0c6000
```

In this example, HPA5501 is the user-defined device file name.

In the case of an HSC or PCI adapter, the `ioscan` output varies from the above 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:

TABLE B-5 Output of ioscan With schgr Driver Configured on the System

| Class | I | H/W Path | Driver | S/W State | H/W Type | Description |
|-------|---|------------|--------|-----------|----------|---------------|
| spt | 0 | 10/4/4.6.0 | schgr | CLAIMED | DEVICE | HP c6280-7000 |

If the schgr driver is not configured on the system, no driver appears to be associated with the library. The ioscan output line resembles:

TABLE B-6 Output of ioscan Without schgr Driver Configured on the System

| 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 |

If the previously outlined procedures fail (this might occur on a v-class machine), refer to the Hewlett Packard-defined procedures outlined in *How to configure a DLT autoloader on V-class for tar and fbackup*, document ID A5248544.

The document can be accessed from the Electronic Support Center on the Hewlett Packard web site.

▼ To Test the Device Driver and Device File Installation

After you have installed the device driver and created the device file run the autochanger information program as follows, to confirm proper installation:

```
# jbinfo (/dev/sjid1u1)
```

If proper installation is confirmed, jbinfo prints a message containing information about the autochanger. If you are testing a second autochanger, the control port is /dev/sjid1u2.

▼ 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 their doors open. Refer to the `jbexercise` man page for complete instructions.

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). The following runs the `jbexercise` program on an Exabyte 10i autochanger:

```
# jbexercise -c (/dev/sjid1u1) -m EXABYTE_10I
```

“Hardware Address Not Found” Message Encountered

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 Solstice Backup `/nsr/logs/messages` file. If you encounter this message, perform the following steps:

1. Log in as root to edit the `/etc/syslog.conf` file and change every instance of `daemon.notice` to `local7.notice`.
2. Edit the Notifications resource configured for the Log Default notification and change the value shown in the Action attribute from `daemon.notice` to `local7.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#` obtained from the `/etc/syslog.pid` file to kill the designated `pid#` 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 your system, the Solstice Backup server can make use of them exclusively to avoid the hardware address problems encountered with `bootpd`.

Errors from Unsupported Media in HP Tape Drives

Certain Hewlett-Packard tape drives can only read 4 mm tapes of a specific length. Some, for example, read only 60 meters tapes and do not support the use of 90- or 120- meter tapes. To determine the type of tape supported by your HP drive, consult the hardware manual provided with the drive.

If you attempt to use unsupported media in an HP tape drive, you might encounter the following types of error messages in the respective 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 10.x storage node, all of the tape drives inside the autochanger are unloaded to their respective slots. If you only want to unload a single drive to its corresponding slot, use the `nsrjb -u -f` command instead.

AIX

This section provides information specific to Solstice Backup software storage node software running on the AIX platform.

▼ To Activate Full Diagnostic Core Dumps on AIX

Earlier releases of Solstice Backup software on an AIX system do not automatically provide full diagnostic core dumps.

You need to set the `NSR_FULLDUMP` environment variable to 1 to activate full diagnostic core dumps.

To set the 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 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 your 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 STK-9840 to an AIX server for Solstice Backup 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 Solstice Backup software storage node software running on the IRIX platform.

IRIX Displays “No Space Left in Tape” Message

If you use an IRIX Solstice Backup storage node, you might receive the following message:

```
BSF invalid argument no space left in tape
```

This is not a Solstice Backup error but an indication that the tape drive you are using might not be configured within the machine’s kernel. 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, follow these steps:

1. If you have not already enrolled as a SurfZone (free program) member, point your web browser to the SGI Technical Assistance Center at www.sgi.com/support/patch_intro.html to enroll and obtain a password.
2. Point your browser to support.sgi.com/surfzone/patches (you will need to provide your SurfZone password).
3. Download the recommended patch.
4. Follow the instructions provided to apply the patch set.

SGI tests specific firmware revisions for each drive and robot with IRIX Solstice Backup software. Even if your site uses drives and robotics that SGI has qualified, you might need to apply patches to IRIX to make the combination work.

Visit the SGI Technical Assistance Center web site for information about and distribution of patches required for your particular configuration. You can also obtain information regarding the current list of supported drives, robotics, and firmware revisions supported by IRIX.

Glossary

This glossary contains terms and definitions found in this guide. Most of the terms are specific to Solstice Backup products.

| | |
|-----------------------------------|---|
| active group | A Solstice Backup backup group that has its autostart attribute enabled. |
| agent | The term used by Solaris to denote a cluster server. A cluster server is a shared resource or virtual machine containing the Solstice Backup server resources (indexes, /nsr directories, and tape devices). The Solstice Backup server is associated with (and known by) the cluster hostname rather than with the physical machine on which it is hosted. |
| annotation | A text string that the Solstice Backup administrator or user associates with a UNIX archive save set to help identify that data later. Annotations are stored in the media database for ease of searching during a retrieval operation and are limited to 1024 characters. |
| Application Specific Module (ASM) | 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, <code>compressasm</code> is Solstice Backup directive used to compress and decompress files. |
| archive | The process by which the Solstice Backup server 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 grooming . |
| archive pool | A volume pool that contains only archived save sets. A separate volume pool, the archive clone pool, contains only cloned archive save sets. Archived save sets are in a different format than regular backup save sets and must be maintained on separate media. |

| | |
|-----------------------|--|
| archive volume | A tape or other storage medium used to store Solstice Backup archive data, as opposed to a <i>backup volume</i> . |
| attributes | A feature of a resource. It is a service or information that the resource provides. |
| autochanger | 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 jukebox, carousel, library, near-line storage, datawheel, and autoloader. |
| auto media management | A feature that enables the storage device controlled by the Solstice Backup server to automatically label, mount, and overwrite a volume it considers unlabeled. Volumes that are eligible for reuse are also automatically recycled. |
| authorization code | A code that is unique to your network that you obtain by sending in the registration information you print after the enabler code is entered on the Solstice Backup server. The authorization code unlocks the software for permanent use. |
| backup, manual | A backup that a user requests from the client's <i>save</i> program. The user specifies participating files, filesystems, and directories. A manual backup does not generate a <i>bootstrap</i> save set. |
| backup cycle | The period of time from one level full backup to the next level full backup. |
| backup group | A group of Solstice Backup clients that begin their scheduled backups at the same time. |
| backup levels | <p>A measurement that determines how much data the Solstice Backup server saves during a scheduled or manual backup.</p> <p>A full backup backs up all files, regardless of whether they have changed.</p> <p>Levels one through nine (1-9) back up files that have changed since the last lower numbered backup level.</p> <p>An incremental (incr) backup backs up only files that have changed since the last backup.</p> |
| backup volume | A tape or other storage medium used to store Solstice Backup backup data, as opposed to an archive volume or migration store. |
| base enabler code | See <i>enabler codes</i> . |

| | |
|-------------------------|---|
| bootstrap | A save set that is essential for the Solstice Backup disaster recovery procedures. It is composed of two components that reside on the Solstice Backup server: the media database and the resource database. |
| browse policy | A policy that determines how long entries for your backup data remain in the client file index. |
| client file index | A database of information maintained by the Solstice Backup server that tracks every file or filesystem backed up. The Solstice Backup server maintains a single client index file for each client computer. |
| client-initiated backup | See backup, manual . |
| client failover | The process of a Solstice Backup client moving from one node to another within the cluster. This process only applies to the virtual cluster client. Failover can occur when the node on which the virtual client is located has a hardware or software failure or when the client is being moved to another node for administrative purposes. |
| clone | The Solstice Backup 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 because it leaves traceable information entries in both the client file index and the media data. |
| clone volume | A duplicated volume. Three types of clone volumes are tracked: backup clone, migration clone, and archive clone. Save sets of different types (for example, archive and migration) cannot be intermixed on the same clone volume. |
| cluster | A network of connected computers used as a single resource. |
| cluster aware | Recognizing a cluster environment. |
| cluster client | A Solstice Backup client that is either a physical cluster client or a virtual cluster client. |
| cluster host | A machine (node) in the cluster environment. |
| cluster mode | Describes a program element working as part of a cluster. Currently, the <code>NetWorker.cluster.svr</code> file specifies the cluster mode to the Solstice Backup server. |
| cluster node | A physical machine in the cluster environment. Also known as a node or physical host. |

| | |
|------------------------|--|
| cluster resource | The shared or managed resource that is registered with a cluster manager. Some platforms or products use the term service rather than resource. |
| cluster server | A shared resource or virtual machine containing the Solstice Backup server resources (indexes, /nsr directories, and tape devices). The Solstice Backup server is associated with (and known by) the cluster hostname rather than with the physical machine on which it is hosted. Also known as agent in Solaris clusters. |
| cluster unaware | Not recognizing a cluster environment. |
| command line interface | An interface with the Solstice Backup software, based on command text entered from the shell prompt. See also the Legato Command Reference Guide which can be found on your Installation . |
| daemon | A program that is not invoked explicitly, but lies dormant waiting for a specified condition to occur. |
| database | A collection of related data that can serve multiple purposes and support multiple users. |
| device | The unit connected to the <i>Solstice Backup server</i> or <i>storage node</i> – either as a stand-alone computer or in an <i>autochanger</i> – that stores data on media. |
| directed recover | A method of recovery where data originating on one client computer is re-created on another client computer. |
| directive | An instruction that directs the Solstice Backup software to take special actions on a given set of files for a specified client during a backup. |
| enabler codes | A special code provided by Sun that, when entered in the Registration resource for the Solstice Backup server, activates the software. The user must then register the software and enter the returned authorization code to permanently license the product. The enabler code that unlocks the base features for the version of the Solstice Backup software you purchased is referred to as a <i>base enabler</i> . Enabler codes for additional features or products (for example, autochanger support) are referred to as <i>add-on enablers</i> . |
| failover | The process of a program or service moving from one node to another within the cluster, either because of a hardware or software failure or for administrative purposes. |

| | |
|-------------------------|--|
| filesystem | <ol style="list-style-type: none"> 1. A file tree that is on a specific disk partition or other mount point. 2. The entire set of all files. 3. A method of storing files. |
| firewall | A system designed to prevent unauthorized access to or from a private network. 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 Solstice Backup software supports client backups from computers that are protected by <i>packet filtering</i> . |
| grooming | The process of removing the original files from a local disk after a successful archive operation. |
| group | A client or group of clients that starts backing up files at a designated time. |
| heterogeneous network | A network with systems of different platforms and operating systems that interact across the network. |
| highly available system | A highly available system allows the application services to continue despite a hardware or software failure. <i>Each cluster node has its own IP address.</i> Each cluster node also has <i>private</i> (local) resources or disks that are available only to that machine. |
| hostname | The name of either a physical or virtual host. |
| interoperability | The capability of software and hardware on multiple computers from multiple vendors to communicate. |
| LUS | An acronym for Legato User SCSI. The LUS driver is used by Legato software products as a proprietary device driver that sends arbitrary SCSI commands to an autochanger. |
| local cluster client | A Solstice Backup client that is not permanently bound to a physical machine, but is instead managed by a cluster manager. It can be bound to more than one physical machine in the cluster and can own its own data disks. It is referred to as a logical or virtual client. |
| local host | The node on which the client or server program is running. |
| media | The physical storage medium to which backup data is written. Solstice Backup software supports tape, magnetic or optical disk, and filesystems as backup media. |
| media database | A database that contains indexed entries about the storage volume location and the life cycle status of all data and volumes managed by the Solstice Backup server. |

| | |
|------------------------------|--|
| media manager | The Solstice Backup component that tracks save sets to backup volumes. |
| multiplex | A Solstice Backup feature that permits data from more than one save set to be written to one storage volume on different storage devices. |
| Solstice Backup client | A computer that accesses the Solstice Backup server to back up or recover data. Clients can be workstations, PCs, or file servers. |
| Solstice Backup server | The computer on a network running the Solstice Backup server software, containing the client file indexes, and providing backup and recovery services to the clients and storage nodes on the same network and media database. |
| Solstice Backup storage node | See storage node . |
| NFS client | A computer that can access files on an NFS server. |
| NFS server | A computer that contains exported filesystems that NFS clients can access. |
| notification | A message generated to the Solstice Backup administrator about important Solstice Backup events. |
| \$NSR_BIN: | The location of the <code>NetWorker.clustersvr</code> file when the Solstice Backup software starts in cluster mode. The binary location varies from platform to platform. |
| operator | The person who monitors the server status, loads backup volumes into the server devices, and otherwise executes the day-to-day Solstice Backup tasks. |
| override | A Solstice Backup feature that allows you to configure a different backup level for a specific date listed in a Schedule resource. Refer to the online help in the <code>nwadmin</code> program for instructions on how to use this feature. |
| packet filtering | 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 firewall . |
| parallelism | A Solstice Backup feature that enables the Solstice Backup server to either back up save sets from several clients or many save sets from one client at the same time. Parallelism is also available during recovers. |

| | |
|-------------------------|---|
| pathname | A set of instructions to the operating system for accessing a file. An <i>absolute pathname</i> tells how to find a file beginning at the root directory and working down the directory tree. A <i>relative pathname</i> tells how to find the file starting where you are now. |
| physical cluster client | A Solstice Backup client that is bound to a physical machine in the cluster and can represent its own resources (private, or local). It can also be called the physical client. |
| physical host | Any one of the nodes (or machines) that forms the cluster. |
| pool | A feature that enables you to sort backup data to selected volumes. A volume pool contains a collection of backup volumes to which specific data has been backed up. |
| recover | A recovery method that re-creates an image of the client filesystems and database on the Solstice Backup server. |
| recyclable volume | A volume whose data has passed both its browse and retention policies and is now available for relabeling and use by a Solstice Backup server or storage node. |
| remote device | A storage device that is attached to a Solstice Backup storage node. |
| resources | Anything that you might need to manage or that a user might want to locate, such as a storage device, backup schedule, or event notification. In the Solstice Backup Administrator program, resources are represented as windows. Resources contain attributes. |
| resource owner | The cluster (logical, not physical) host that owns the resource. If a resource (for example, a shared disk) is not owned by any virtual host it is assumed to be owned by the local host (physical node) that hosts the resource. |
| retention policy | A policy that determines how long save set entries are retained in the Solstice Backup server's media database. |
| retrieve | The process of locating and copying back files and directories that a Solstice Backup server has archived. |
| root | <ol style="list-style-type: none"> 1. The UNIX superuser account (with username "root" and user ID). By extension, the privileged system-maintenance login on any operating system. 2. The top node of the system directory structure, the home directory of the root user. |
| save set | A group of files or a filesystem from a single client computer backed up onto storage media. |

| | |
|------------------------|--|
| save set consolidation | The process that merges an incremental backup with the last full backup of a save set to create a new backup. See also backup levels . |
| save set ID | An internal identification number that the Solstice Backup server assigns to a save set. |
| save set recover | The recovery of specified save sets to the Solstice Backup server. |
| save set status | The save set status indicates whether a given save set is restorable, recoverable, or recyclable. The save set status also indicates whether the save set has been successfully backed up. |
| savestream | The data and save set information being written to a storage volume during a backup. A savestream originates from a single save set. |
| server | The computer that runs the Solstice Backup software, contains the online indexes, and provides backup and recovery services to the clients on a network. |
| server-failover | The process that occurs when the Solstice Backup software moves from one node to another node in the cluster. In a cluster environment, failover occurs during hardware or software failure or when the Solstice Backup server is moved to another node for administrative purposes. |
| shared disk | The storage disk that is connected between multiple nodes in the cluster |
| shell prompt | A cue for input in a shell window where you enter a command. See also command line interface . |
| silo | A repository for holding hundreds or thousands of volumes. Silo volumes are identified by barcodes, not by slot numbers. They are controlled by silo management software on a server computer that might or might not be the Solstice Backup server computer. |
| SNMP | 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). The Solstice Backup server uses SNMP to send messages to the administrator about Solstice Backup events. |
| staging | The process of moving data from one storage medium to another, for example, to move backed-up, archived, or migrated save sets from a disk file to a tape. |

| | |
|------------------------|---|
| stand-alone device | A <i>storage device</i> that contains a single drive for backing up data. |
| stand-alone | In a cluster environment, if the <code>NetWorker.clustersvr</code> file is missing at the binary location, the Solstice Backup server will start in non-cluster mode, also called stand-alone mode. The stand-alone mode is sometimes referred to as the server backing-up itself. |
| storage device | The hardware that reads and writes data during backup, recovery, or other Solstice Backup operations. |
| storage node | A storage device physically attached to another computer whose backup operations are administered from the controlling Solstice Backup server. |
| target sessions | A feature that allows you to set the number of backup sessions accepted by an active device |
| trap | An unsolicited notification sent from a Simple Network Management Protocol (SNMP) agent to the SNMP event manager. |
| versions | The date-stamped collection of available backups for any single file. |
| virtual cluster client | A Solstice Backup client that is not permanently bound to a physical machine, but is instead managed by a cluster manager. It can be bound to more than one physical machine in the cluster and can own its own data disks. It is also referred to as logical cluster client or virtual client. |
| volume | A physical unit of media, such as magnetic tape, optical disk, or disk file. |

Index

A

administrators

- access, 48
- adding, 49
- privileges, 29, 48
 - directed recoveries, 223, 225
- removing, 49
- storage nodes, 67

ansrd daemon, 8

application information field

- EMC required values, 357

application specific modules

- always, 128
- compressasm, 128
- nullasm, 128
- pw1, 128
- pw2, 128
- skip, 128

Archive application. See archives

archive requests

- attributes, 254
 - annotation, 254
 - archive clone pool, 254
 - archive pool, 254
 - save set, 254
- control, 257
- creating, 255
- defined, 253
- deleting, 256
- details, 259
- disabling, 258
- editing, 256
- naming errors, 398

rescheduling, 258

- starting, 257
- stopping, 257

archives

- access, 246
- archive requests, 253
 - attributes, 254
 - creating, 255
 - deleting, 256
 - editing, 256
- archives, 355
- backups, comparison, 244
- benefits, 243
- clients
 - creating, 250
 - enabling, 251
 - privileges, 254
- clones, 265, 277
 - display errors, 397
- compression, 249
- grooming, 250
- indexed, 248
 - retrieval, 261
- licensing, 244
- management, 257
- manual, 251
- non-indexed, 247
 - retrieval, 261
- nsrchive program, 398
- nwretrieve program, 260
- options, 249
- permissions, 245
 - clients, 245

- server, 245
- policies
- pools, 147, 248
 - configuration, 144
 - errors, 398
- remote requests, failure, 397
- requirements, 244
- retrieval, 259, 260, 261
 - indexed, 261
 - non-indexed, 261
 - permissions, 260
- save sets
 - multiple, 397
 - retrieval, 261
- scheduled, 253
 - details, 259
 - disabling, 258
 - management, 257
 - rescheduling, 258
 - starting, 257
 - stopping, 257
- troubleshooting, 397
- usage, 246
- volumes, 250
 - tracking, 249
- asavegrp daemon, 8
- asm. See application specific modules
- attributes, hidden, 33
- audience, for documentation, xxxv
- authorization code. See licensing
- authorization. See licensing
- auto media management, 108
 - autochangers, 305
 - clones, 265
 - defined, 162
- auto media verification, 146
 - media position errors, 378
- autochangers, 156
 - adding, 282
 - AIX, 428
 - ATL 2640 libraries, 409
 - attributes, 288, 303
 - new, 400
 - auto media management, 305
 - autodetection, 405
 - bar codes
 - labeling, 310, 311
 - cartridges, cleaning, 302, 303
 - replacing, 304
 - Celestra devices, 346
 - cleaning, 302
 - cartridges, 304
 - non-default slot, 303
 - control ports access, 409
 - defined, 5, 281
 - deleting, 291
 - destination component, 407
 - devices
 - adding, 288
 - calibrating, 318
 - cleaning, 302, 303
 - drivers, 403
 - remote, 292
 - editing, 290
 - element status, disabling, 318
 - emulation, 282
 - HP-UX
 - device file, 421
 - error messages, 425
 - installation, 416, 421
 - major number, 422
 - media, unsupported, 426
 - minor number, 423
 - pass-through drivers, 416
 - SCSI addresses, 416
 - shutdown, 420
 - system start, 421
 - testing, 424, 425
 - unloading drives, 427
 - installation, HP-UX, 416
 - inventories
 - slots, 410
 - inventorying, 313, 406
 - I/O errors, 314
 - jbconfig program, 282
 - labeling, ranges, 309
 - licensing, 281, 293
 - maintenance commands, 402
 - NDMP support, 345, 348, 350
 - NetApp support, 347
 - notifications, 317
 - operations, 301
 - optical, 319
 - pathname order, 289
 - pools, 318
 - remote management, 291
 - requirements, 281

- resetting, 290, 319
- resource, 7
- sharing, 293, 296
 - backups, 299
 - block size, 299
 - cloning, 298
 - configuration, 295
 - defined, 294
 - devices, 295
 - licensing, 293
 - recoveries, 296, 299
 - Solstice Backup Network edition, 293
 - Solstice Backup Power edition, 293
 - Solstice Backup Workgroup edition, 293
 - volume management, 298
 - volumes, 298
 - volumes, unmounting, 299
- sleep function, setting, 292
- slots
 - available, 301
 - changing, 301
 - cleaning, 303
 - depositing volumes, 410
- STK-9840, 428
- storage nodes, 166, 291
 - defining, 292
- testing, 287
- tips, 316
- troubleshooting, 400
- volumes, 305
 - depositing, 315
 - exporting, 315
 - importing, 315
 - inventorying, 313
 - labeling, 308, 310
 - mounting, 307
 - recycling, 317
 - selecting, 306
 - unmounting, 308
 - withdrawing, 316
- X11 libraries, 406

B

- Backup Command
 - savenpc program, 120
- backup groups. See groups
- backup schedules. See schedules

- backups, 26, 101
 - abandoned, 185
 - archiving, 243
 - comparison, 244
 - autochangers, sharing, 299
 - balancing, 182, 196, 197
 - bootstrap, 69, 143, 191
 - storage nodes, 69
 - clearcase vob, 124
 - clients
 - adding, 28
 - licenses, 133
 - cluster, 334
 - completed, 191
 - consolidated, 196
 - customization, 118
 - criteria, 118
 - scripts, 119
 - cycle, using levels, 195
 - daemons, 12
 - details, viewing, 190
 - directives, 127
 - filesystems, 101
 - large, 198
 - forced incremental, 187
 - groups, 182
 - management, 188
 - preview, 191
 - restarting, 190
 - starting, 189
 - stopping, 189
 - immediate, 27
 - incremental, 142
 - forced, 187
 - levels, 194, 202
 - 1-9, 203
 - consolidated, 196, 203, 204
 - full, 203, 204, 383
 - incremental, 203
 - level, 204
 - options, 203
 - overriding, 202
 - save set consolidation, 211
 - skip, 203
 - usage, 203, 204, 205
 - management, 188
 - manual, 213
 - media database, 143
 - policies, 114

- pools, 143
 - monitoring, 27
 - NDMP, 354, 360, 362
 - details, viewing, 362
 - non-scheduled, 213
 - online indexes, 69, 204
 - open files, 193
 - order, 200
 - pools, 142
 - preview, 191
 - programs, 12
 - recoveries, 374
 - rejection, 377
 - restarting, 190
 - save sets, 101
 - completed, 190
 - failed, 190
 - load balancing, 101
 - omitted, 199
 - pending, 190
 - status, 190
 - scheduled, 26
 - clones, 267
 - server, forced, 385
 - starting, 189
 - stopping, 189
 - failure, 375
 - time intervals, setting, 186
 - troubleshooting, 374, 383
 - unattended, 142, 173
 - verification, 27
 - volumes
 - clones, 264
 - selection criteria, 168, 169
- bar codes
 - labels, 310
 - identical, 311
 - range, 312
 - tips, 311
 - types, 312
 - unmatching, 312
 - volumes, 312
 - troubleshooting, 310
 - volumes
 - labeling, 310
- block size configuration, 299
- bootstrap, 12
 - backups, storage node, 69

- clone, 69
 - file, 192
 - NDMP, 342
 - pools, 139, 144
 - printing, 191, 192
 - failure, 385
 - recovery, 86
 - storage nodes, 69
- browse period, 237
- browse policies, 30, 104, 105, 217
 - archiving, 243
 - clones, storage nodes, 279
 - creating, 113
 - defined, 104
 - editing, 114
 - granularity, 117
 - save sets, 118
 - recoveries, 233, 234
 - usage, 104

C

- cartridge access port, 315
- Celestra
 - devices, 345, 346
 - NDMP
 - limitations, 342
 - NetApp backups, 364
 - support, 346
- cleaning cartridges. See autochangers
- clearcase vob backups, 124
- client file index, 2, 4, 104, 236
 - backups, 204
 - browse policies, 75, 105
 - clones, 265
 - storage nodes, 279
 - conversion, 78, 79
 - multiple clients, 80
 - single client, 79
 - entries, 71
 - adding, 72
 - removing, 72, 75, 77, 78
 - returning space, 72
 - integrity checking, 377
 - location, designating, 72
 - management, 70
 - size, 75

- messages, missing, 383
- moving, 73
- NDMP, 362
- policies, 75
- pools, 139, 144
- recoveries, 236, 237, 238, 239
- restoration, 233, 240
- save sets
 - clones, 269
 - cycles, removing, 76
 - entries, 108
 - removing, 75
- scanner program, 240
- size, 71
 - management, 75
 - notification, 378
- structure, 71
- clients
 - access, remote, 94
 - adding, 28
 - administration, directed recoveries, 225
 - aliases, problems, 380
 - archives
 - creating, 250
 - enabling, 251
 - permissions, 245
 - privileges, 254
 - retrieval, 260, 261
 - backups, 101
 - customization, 118
 - filesystems, large, 198
 - renamed, 384
 - client ID, 93, 94
 - clones, storage nodes, 69
 - cluster, 333
 - communication, configuration, 44
 - configuration, 30, 89, 90
 - default, 90
 - creating, 28, 90
 - daemons, 9
 - defined, 89
 - deleting, 92
 - editing, 92
 - groups, 182
 - assigning, 182
 - indexes
 - location, 72
 - moving, 73
 - installation, 90
 - licensing, 95, 133, 376, 377
 - multiple, 100
 - name, changing, 93
 - NDMP, 341, 356
 - NetWare, 99
 - permissions, 90, 94
 - archiving, 245
 - ping command, testing, 391
 - platform, 95
 - policies, multiple, 112
 - priority, 132
 - programs, 9
 - recoveries
 - directed, 227
 - failure, 384
 - NetWare, 227
 - UNIX, 227
 - Windows 98/95, 227
 - Windows NT, 226, 227
 - resource, 6
 - restoration, 94
 - save sets, 100, 101
 - server
 - communication, 388
 - servers
 - converting, 86
 - Solaris
 - binary location, 415
 - testing, 391
 - UNIX, 96
 - unsynchronized client clock, 375
 - Windows 95, 98
 - Windows 98, 98
 - Windows NT, 96
- clones
 - archives, 277
 - display errors, 397
 - autochangers, shared, 298
 - NDMP, 354
 - online indexes, storage nodes, 279
 - pools
 - configuration, 143
 - recoveries, 277
 - reports, 264
 - save sets
 - archives, 277
 - automated, 273, 276
 - client file index, 269
 - clients, 276

- defined, 268
- details, viewing, 272
- devices, 269
- manual, 270
- media database, 269
- methods, 270
- multiple, 274
- online indexes, 269
- performance, 269
- performing, 271
- pools, 269
- requirements, 270
- retention policies, 269
- tasks, 275
- volumes, 268, 269
- storage nodes, 69
 - affinity, 278
 - online indexes, 279
- volumes, 266
 - archives, 265, 277
 - auto media management, 265
 - automatic, 264
 - backups, scheduled, 267
 - client file index, 265, 267
 - creating, 264, 267
 - defined, 263
 - deleting, 266
 - details, viewing, 268
 - devices, 263, 264
 - disaster recovery, 265
 - group, 267
 - manual, 264
 - media database, 265, 266, 267
 - online indexes, 264, 265, 267
 - performance, 266
 - pools, 266
 - purging, 266
 - recovery, 267
 - requirements, 265
 - save sets, 264
 - servers, 264
 - storage nodes, 264
 - usage, 263, 265
- clusters
 - backups, 334
 - client, 333
 - defined, 331
 - failover, 333
 - licensing, 336

- nodes, 331
- recovery, 336
- server, 333
- type, 333
- configuration
 - order for resources, 7
- crontab, 82, 83
- cross-platform management, 2
- customize
 - script, Backup Command use, 123

D

- daemons
 - log file, 81
 - troubleshooting, 374
- dasadmin program, 323
- data life cycle, 110
- devices
 - adding, 157
 - auto media management, 162
 - autochangers
 - adding, 288
 - calibrating, 318
 - cloning, 298
 - naming, 295
 - concurrent, 156
 - configuration, 30
 - creating, 157
 - defined, 155
 - deleting, 158
 - disk lable errors, 384
 - drivers, testing, 403
 - editing, 158
 - errors, limiting, 161
 - filesystem, 159
 - configuration, 159
 - defined, 159
 - restrictions, 160
 - staging, 160, 178
 - maintenance commands, 402
 - media
 - nonrewinding, 156
 - pools, 144
 - multiple, 156
 - nonrewinding, 156, 173, 409
 - pools, 144

- media, 144
- remote, 164, 166
 - autochangers, 166
 - naming, 292
 - priveleges, 165
 - See also storage nodes
 - storage, 66
 - timeout settings, 166
- resource, 6
- save sets, clones, 269
- sessions, 169
- silos, 320
 - cleaning, 325
 - sharing, 325
- storage nodes, 66, 69
 - autochangers, 68
 - configuration, 67, 164
 - silos, 68
- support, 156
 - Solstice Backup edition type, 155
 - Solstice Backup Network edition, 155
 - Solstice Backup Power edition, 155
 - Solstice Backup Workgroup edition, 155
- troubleshooting, 400
- volumes, clones, 263, 264

diagnostic tools, 413

- AIX core dumps, 427
- communications tests, 414
- reports, 414

directed recoveries, 221

- access
 - remote, 223, 225
 - server, 223
 - UNIX, 222
- administration, 222
- advantages, 236
- clients
 - NetWare, 227
 - types, 226
 - UNIX, 227, 228
 - Windows 98/95, 227
 - Windows NT, 226, 227
- platforms, 226
- recover program, 227
- usage, 222

directives, 30

- application specific modules, 127
 - always, 128

- compressasm, 128
- nullasm, 128
- pw1, 128
- pw2, 128
- skip, 128
- creating, 131
- customizing, 131
- defined, 127
- deleting, 132
- editing, 132
- preconfigured, 130
- resource, 6
- syntax, 128
- writing, 127

disaster recovery, 56, 221

- clones, 265

documentation

- CD-Based, 20
- choices, 20
- printed, 21
- web-accessible, 21

E

- element status, disabling, 318

EMC

- application information values required, 357

enabler codes. See licensing

environmental variables, 86

- ACI_MEDIA_TYPE, 324, 413
- CSI_HOSTNAME, 413
- DAS_CLIENT, 323, 413
- DAS_MEDIUM, 413
- DAS_SERVER, 323, 413
- NSR_DEV_BLOCK_SIZE_, 87
- NSR_DEV_LOAD_TIME_, 87
- NSR_DEV_TAPE_FILE_SIZE_, 87
- NSR_DEV-DEFAULT_CAPACITY_, 87
- silos, 412

F

- failover. See clusters
- file conversion, sparse to fully-allocated, 387
- filesystem devices, 159
 - configuration, 159

- defined, 159
- restrictions, 160
- staging, 160, 178
- filesystems
 - backups, large, 101
 - remote, saving, 396
- firewall support, 57
 - packet filtering, 57
- firmware, verification, 394

G

- grooming, 250
- groups, 10, 30
 - attributes, 185
 - backups
 - abandoned, 185
 - management, 188
 - preview, 191
 - bootstrap, 191
 - clients
 - policies, multiple, 112
 - clones, 267
 - completed, 191
 - creating, 185
 - default, 184
 - defined, 182
 - deleting, 188
 - details, viewing, 190
 - editing, 188
 - filesystems, large, 198
 - forced incremental, setting, 187
 - NDMP, 362
 - details, viewing, 362
 - override regular schedule, 202
 - overview, 181
 - pools, 182
 - preconfigured, 184
 - previewing, 191
 - resource, 6
 - restarting, 190
 - starting, 189
 - stopping, 189
 - time intervals, setting, 186
 - usage, 182

H

- hidden attributes, 33
- host device selection, 296
- host transfer affidavit, 17, 18
- hosts table
 - ping command, testing, 391
 - setting, 389

I

- indexes. See online indexes
- inquire command, 289, 350
- inventorying, autochangers. See autochangers
- ioscan program, 416
- IP
 - errors
 - hosts table, 389
 - troubleshooting, 388
 - name search order, 390

J

- jbconfig program, 164, 165, 166, 291, 295
 - hanging, 405
 - silos, 322
 - storage nodes, 67, 68
 - X11 libraries, 406
- jbexercise program, 287, 355, 403, 405, 425
 - failure, 405
- jbinfo program, 424
- jukebox. See autochangers

L

- label templates, 148
 - attributes, 151
 - components, 153
 - creating, 153
 - default, 148
 - deleting, 154
 - editing, 154
 - naming strategies, 153
 - number sequences, 152
 - preconfigured, 149

- resource, 6
- usage, 148, 167
- labeling
 - bar codes, 311
 - range, 312
 - volumes
 - bar codes, 310
 - silos, 328
 - tips, 175
- labels
 - autochangers, 306, 308
 - bar codes
 - identical, 311
 - matching names, 312
 - types, 312
 - name reuse, 312
 - volumes, 148, 167
- Legato Command Reference Guide, 24
- Legato Compatibility Guide, 5
- Legato Error Message Guide, 24
- Legato NetWorker Disaster Recovery Guide, 22
- License Manager, 41
- licensing, 30
 - archiving, 244
 - authorization, 41
 - authorization codes, 39
 - authorized mode, 17
 - autochangers, 281
 - sharing, 293
 - clients, 95, 133
 - backup rejection, 377
 - enabler codes, 133
 - errors, 376
 - clusters, 336
 - copy violation, 386
 - enabled mode, 17
 - enabler codes, 18
 - certificate, 17
 - clients, 133
 - entering, 40
 - Power edition, 18
 - removal, 40
 - transfer, 17
 - evaluation mode, 17
 - modes, 16
 - NDMP, 345
 - overview, 39
 - reregistration, 85

- Solstice Backup clients, 133
 - storage nodes, 67
- logical backup volumes, 160
- lusdebug program, 403
 - failure, 403

M

- man pages, 23
- media
 - auto media management, 162
 - handling errors, 161
 - nonrewinding, 156
 - silos, 325
 - supported, 156
 - viewing, 156
- media database, 4, 104
 - backups, 204
 - clones, 265, 266
 - storage nodes, 279
 - compression, 76
 - conversion, 78
 - entries, removing, 77, 78
 - management, 70
 - size, 76
 - restoration, 234, 240
 - retention policies, 104, 105
 - save sets
 - clones, 269
 - entries, 108
 - scanner program, 240
 - size
 - size, management, 76
 - structure, 71
- message logs
 - files, 57
 - management, 81, 82
 - reducing size, 81, 82
- mminfo program, 56, 115, 233
- mmlocate program, 56
- mmrecov program, 69, 86
- multiplexing, 3, 46
 - NDMP, 355

N

- name resolution, 390
 - name servers, disabling, 390
 - naming requirements, 394
 - NDMP
 - archives, 355
 - attributes, 356
 - autochangers, 345, 346, 347, 348
 - autochangers, configuration, 350
 - backups, 360, 361, 362
 - details, 362
 - parallel, 354
 - bootstrap, 342
 - browsing, 354
 - Celestra, 346
 - devices, 345
 - support, 342
 - client file index, 362
 - clients, 341, 356
 - configuration, 361
 - definitions, 354
 - clones, 354
 - configuration, 343
 - client, 361
 - local, 358
 - three party remote, 358
 - configuring, 355
 - connection software, 341
 - defined
 - failover support, 355
 - groups, 362
 - inquire command, 350
 - installation, 340, 341
 - licensing, 345
 - limitations, 342
 - local configuration, 358
 - media preparation, 361
 - multiplexing, 355
 - ndmpjbconf_NDMPhostname file, 348
 - NetApp
 - Celestra Inline Image backups, 364
 - devices, 346
 - remote tape drive backups, 364
 - support, 342, 347, 351
 - nsr resource attributes, 355
 - nsrmmmd program, 354
 - nsrndmp_recover program, 354
 - nsrndmp_save program, 354
 - operating systems supported, 341
 - prerequisites, 361
 - recoveries, 354
 - destructive, 363
 - test, 363
 - save set consolidation, 355
 - save set recovery, 355
 - savegroup probes, 354
 - scanner program, 355
 - Solstice Backup
 - configuring, 355
 - support, 339, 340
 - Solstice Backup support, 340
 - Solstice Backup, effects on, 354
 - storage nodes, 355
 - swap space, 343
 - terms, 339
 - three party remote configuration, 358
 - usage, 358
- ndmpjbconf command, 349
- network data management protocol. See NDMP
- newsyslog script, 82, 83
- notifications, 31
 - autochangers, 317
 - creating, 54
 - customizing, 51
 - defined, 50
 - deleting, 55
 - editing, 55
 - elements, 51
 - events, 52
 - preconfigured, 50
 - printing, 53
 - priorities, 53
 - programs, 52
 - resource, 6
- nsr_getdate program, 115, 116
- nsr_shutdown program, 81, 82, 83, 84, 224, 379, 384
- nsr_support script, 414
- nsradmin program, 31, 33
- nsrarchive program, 246
- nsrck daemon, 8, 10, 11, 86, 383
- nsrck program, 79, 238, 239
 - CPU usage, 80
- nsrclone program, 267
- nsrd daemon, 8, 10, 11, 14, 81
 - ports, 65

- tasks, 9
- nsrexec daemon
 - multiple processes, 374
 - ports, 65
 - troubleshooting, 374
- nsrexecd daemon, 9, 11, 81
- nsrexecd program, 246
 - functionality, 393
- nsrim daemon, 9, 12, 266
- nsrindexd daemon, 9, 15
 - ports, 65
- nsrinfo program, 56
 - defined, 56
- nsrjb program, 68, 291, 298, 299, 304, 410, 427
 - I/O errors, 315
 - o errors, 411
 - ports, 65
- nsrlic program, 134
- nsrmm command, 75
- nsrmm program, 115, 233, 379
 - defined, 56
- nsrmmmd command
 - NDMP
 - control timeout not supported, 355
 - polling interval not supported, 355
 - restart interval not supported, 355
- nsrmmmd daemon, 9, 11, 13, 14, 15, 68, 87, 267
 - ports, 65
- nsrmmdbd daemon, 9, 13, 15
 - ports, 65
- nsrndmp_recover command, 363
- nsrndmp_recover program, 363
- nsrndmp_save program, 360
- nsrstat program
 - ports, 65
- nsrtrap, 368
- nsrwatch program, 32, 56
- nwadmin program. See Solstice Backup Administrator program
- nwarchive program, 36, 249, 251
- nwbackup program, 34, 213, 215
- nwrecover program, 35, 104, 215
- nwretrieve program, 38, 260, 261

O

- online indexes
 - actions, 72
 - client file index, 4
 - clones, 264
 - conversion, 71, 78
 - disk space requirements, 79
 - entries
 - removing, 77
 - information, viewing, 74
 - management, 70
 - size, 75
 - media database, 4
 - recovered
 - relocating, 240
 - recovery
 - location, 380
 - relocating
 - recovered, 240
 - restoration, 233
 - save sets
 - clones, 269
 - size, 71
 - management, 75
 - structure
 - volumes
 - removing, 78, 177
 - volumes, removing, 77
- open files
 - backing, 193
- organizational structure, labeling for, 153
- overriding, 118

P

- packet filtering, 57
- parallelism, 3, 46
 - defined, 46
- performance factors, 45
 - multiplexing, 46
 - parallelism, 46
 - Solstice Backup, 3
 - Solstice Backup edition, 46
 - target sessions, 47
- permissions
 - client, 90
 - management of, 48

- retrieval, 260
 - storage nodes, 67
 - ping command
 - usage, 391
 - policies, 30
 - backups, manual, 114
 - browse, 104, 105, 217
 - data life cycle, 110
 - defined, 104
 - modification, 115
 - usage, 104
 - clients, 112
 - creating, 113
 - data life cycle, 110
 - deleting, 114
 - editing, 114
 - granularity, 117
 - manual backups
 - modifying, 116
 - multiple, 112
 - planning, 198
 - preconfigured, 112
 - recoveries, 217
 - resource, 6
 - retention, 104, 105, 107, 108
 - data life cycle, 110
 - defined, 104
 - modification, 115
 - usage, 104
 - volume relabeling, 105
 - save sets, overriding, 118
 - pools, 4
 - archive, 144, 147, 261
 - errors, 398
 - archiving, 248
 - auto media verification, 146
 - autochangers, 318
 - bootstrap, 139, 144
 - client file index, 139, 144
 - clones, 143, 266, 269
 - configuration, 142
 - archive, 144
 - clone, 143
 - criteria, 140, 141
 - devices, 144
 - incremental backups, 142
 - label templates, 148
 - manual backups, 143
 - conflicts, 141
 - creating, 145
 - archive, 147
 - default, 138, 143, 148
 - clone pool, 143
 - defined, 135
 - deleting, 147
 - devices, 144, 145
 - media, 144
 - editing, 147
 - expression matching, 138, 139
 - groups, 182
 - label templates, 148
 - PC archive, 261
 - preconfigured, 137
 - resource, 6
 - restrictions, 139
 - save set consolidation, 210
 - save sets, 135
 - types, 137
 - usage, 139, 167
 - volumes
 - labeling, 148
 - volumes, labeling, 148
 - portmappers
 - nsrexecd program, 393
 - registration, 65
 - verifying, 392
 - ports
 - nsrd daemon, 65
 - nsrexec daemon, 65
 - nsrindexd daemon, 65
 - nsrjb program, 65
 - nsrmmmd daemon, 65
 - nsrmmdbd daemon, 65
 - nsrstat program, 65
- ## R
- read-only volume mode, 171
 - recover program, 9, 13, 15, 85, 103, 108, 227, 336
 - recoveries, 215, 216
 - autochangers
 - sharing, 296, 299
 - backup times, 217
 - client file index, 238, 239
 - clients, renamed, 384

- clones, 267
 - volumes, 277
- cluster, 336
- components interaction, 13
- conflicts, 215, 218
 - resolution, 219
 - resolutions, 219
- daemons, 13
- data
 - relocation, 218
- directed, 221
 - access, 222
 - administration, 222, 223, 225
 - advantages, 236
 - clients, 226, 227
 - platforms, 226
 - recover program, 227
 - UNIX-based, 228
 - usage, 222
- disaster-related, 221
- failure, 382
- files
 - details, 220
 - searching, 220
 - versions, 220
- index-based, 236
 - advantages, 236
 - disadvantages, 237
- interruptions, 383
- methods, comparing, 235
- NDMP
 - destructive, 363
- NDMP, test, 363
- policies, 217
- program, 13
- save sets, 229, 230, 238
 - browse period, 230
 - client file index, 233
 - clones, 277
 - conflicts, 233
 - costs, 237
 - entire, 234
 - files, 235
 - media database, 234
 - NDMP, 355
 - omitting, 103
 - relocating, 232
 - requirements, 230
 - scanner program, 232
 - specifying, 231
 - usage, 229
 - volumes, 232
- type, determination, 238
- volumes
 - required, 221
- recovery
- recycling. See volumes
- registration resource, 6
- re-hosting, 18
- Release Supplement, 22
- remote access, 215
 - clients, 94
 - directed recoveries, 223, 225
 - recoveries, 396
- remote filesystems, saving, 396
- remote procedure call, 8
- reports, 55
 - clones, 264
 - daemon log file, 81
 - data movement operations, 55
 - diagnostic, 414
 - message logs
 - files, 57
 - management, 81, 82
 - newssyslog, 82
 - reducing size, 81, 82
 - save set policies, 117
 - savegrp completion, 12
 - server statistics and diagnostics, 57
 - storage management application, 56
 - system console log, 192
- resources, 30
 - autochanger, 7
 - client, 6
 - configuration order, 7
 - defined, 5
 - device, 6
 - directive, 6
 - groups, 6
 - label template, 6
 - notification, 6
 - policy, 6
 - pool, 6
 - registration, 6
 - schedule, 7
 - server, 6

- staging, 6
- retention policies, 30, 104, 105, 107, 108
 - archiving
 - clones, 269
 - storage nodes, 279
 - creating, 113
 - defined, 104
 - editing, 114
 - granularity, 117
 - read-only volume mode, 171
 - save sets
 - clones, 269
 - overriding, 118
 - usage, 104
 - volumes
 - relabeling, 105
- retrieval
 - annotations, empty, 398
 - clients, 260, 261
 - defined, 259
 - indexed, 261
 - non-indexed, 261
 - nwretrieve program, 260
 - permissions, 260
 - save sets, 261
- retrievals
 - troubleshooting, 397
- rpcbind command, 393
- rpcinfo command, 392

S

- save program, 10, 116, 119, 124, 335, 336
- save sets
 - archived
 - retrieval, 261
 - backups, 101
 - filesystem, 101
 - status, 190
 - client combination, 100
 - client file index
 - entries, 108
 - client priority, 132
 - clones, 268, 270
 - volumes, 264, 268
 - completed, 190
 - consolidation, 207
 - files, large, 208
 - limitations, 209
 - NDMP, 355
 - pools, 210
 - requirements, 208
 - starting, 210
 - usage, 207, 208
 - defined, 4
 - defining, 101
 - failed, 190
 - information, viewing, 74, 172
 - load balancing, 101
 - media database, entries, 108
 - omitted, 199
 - pending, 190
 - policies
 - modification, 115
 - overriding, 118
 - reports, 117
 - pools, 135
 - recoverable, 108
 - recoveries, 229, 230, 232, 238
 - browse period, 230
 - client file index, 233
 - clones, 277
 - conflicts, 233
 - costs, 237
 - entire, 234
 - files, 235
 - media database, 234
 - online indexes, 233
 - relocating, 232
 - requirements, 230
 - specifying, 231
 - usage, 229
 - volumes, 232
 - recoveries, omitting, 103
 - recovery
 - NDMP, 355
 - recyclable, 108
 - retrieval, 261
 - See also clones, 270
 - staging. See staging
 - status, 102, 107
 - aborted, 102
 - browsable, 102
 - changing, 103
 - in progress, 102
 - recoverable, 102

- recyclable, 102
 - scanned-in, 103
 - suspect, 103
- suspect, 278
- savefs program, 10
- savegrp program, 9, 10, 11, 12, 55, 185, 203
- savepnc program, 91
- savepnc
 - used in Backup Command, 120
- savepnc program, 118, 119
- scalability, Solstice Backup, 4
- scanner program, 78, 93, 103, 108, 110, 165, 232, 234, 235, 237, 240, 267
 - NDMP, 355
 - online indexes, 240
 - record size, 382
 - recovering clone volumes, 267
 - storage nodes, 67
 - usage, 240
 - volumes
 - read-only, 379
- schedules, 30, 182, 193
 - attributes, 199
 - backup cycle, 195
 - balancing, 196, 197
 - creating, 200
 - default, 194, 196
 - deleting, 201
 - editing, 201
 - load balancing, 101
 - order, 200
 - overriding, 202
 - overview, 181
 - planning, 198
 - preconfigured, 194
 - resource, 7
 - save sets, omitted, 199
 - staggered, 197
 - usage, 193
- schedules, large filesystems, 198
- SCSI
 - address selection for HP-UX, 416
- See also archive requests
- servers, 4
 - administration
 - directed recoveries, 223, 225
 - administrators
 - access, 48
 - adding, 49
 - privileges, 48
 - removing, 49
 - archiving, permissions, 245
 - changing, 43
 - clients
 - communication with, 388
 - converting, 86
 - license enforcement, 133
 - cluster, 333
 - configuration, 29, 44
 - connectivity, 43
 - termination, 44
 - daemons, 8
 - devices
 - volume management, 298
 - directed recoveries, administration, 223
 - environmental variables, 86
 - errors, binding to, 395
 - hostname creation, 84
 - index
 - backup, failure, 385
 - management, 71
 - index entry, defining, 84
 - licensing
 - reregistration, 85
 - maintenance, 81
 - moving, 83, 84, 85
 - multiplexing, 46
 - parallelism, 46
 - performance factors, 45
 - permissions
 - archiving, 245
 - ping command, testing, 391
 - programs, 8
 - resource, 6
 - setup, 44
 - storage nodes, converting, 165
 - target sessions, 47
 - testing, 391
 - volumes
 - clones, 264
- silos, 166
 - defined, 320
 - devices, 320
 - cleaning, 325
 - defining, 324
 - releasing, 324

- sharing, 325
- EMASS/Grau, 323
- environmental variables, 412
- IBM 3494, 323
- installation, 321
- jbconfig program, 322
- media management, 325
- NDMP support, 342
- slot numbers, 326
- storage nodes, 166
- StorageTek, 322, 330
- tape library interface (STLI), 320
- usage, 320
- volumes, 320
 - allocating, 328
 - deallocating, 329
 - exporting, 326
 - importing, 326
 - labeling, 328
 - mounting, 327
 - unlabeled, 328
 - unmounting, 327
- simple network management protocol. See SNMP
- SNMP, 367
 - Configuring, 368
 - defined, 367
 - notification configuration
 - notifications
 - configuration, 368
 - nsrtrap, 368
- Solstice Backup
 - administration interfaces, 31
 - archives. See archives
 - backups. See backups
 - client license
 - enforcement, 133
 - client/server communication configuration, 44
 - clients. See clients
 - command line, 33
 - configuration
 - sequence, 29
 - tasks, 29
 - configuring, 29
 - cross-platform management, 2
 - customizing, 29
 - daemons, 8
 - directives. See directives
 - documentation, 20
 - ease of use, 3
 - editions, 46
 - device support, 155
 - enabling. See licensing
 - firewall support, 57
 - functionality, 8
 - groups. See groups
 - how works, 4
 - installation, 29
 - interfaces, 34
 - hidden attributes, 33
 - nwarchive program, 36
 - nwbackup program, 34
 - nwrecover program, 35
 - nwretrieve program, 38
 - introduction to, 1
 - label templates. See label templates
 - licensing. See licensing
 - man pages, 23
 - modules, 3
 - NDMP
 - configuring, 355
 - effects, 354
 - Network edition, 19, 46
 - autochanger sharing, 293
 - device support, 155
 - notifications. See notifications
 - online help, 23
 - performance, 3
 - policies. See policies
 - pools. See pools
 - Power edition, 18, 19, 46
 - autochanger sharing, 293
 - device support, 155
 - products, 19
 - client, 19
 - server, 19
 - programs, 8
 - recovery. See recovery
 - related products, 20
 - remote access, 215
 - directed recoveries, 223, 225
 - reporting, 2
 - system console log, 192
 - resources, 30
 - scalability, 4
 - schedules. See schedules
 - SNMP Module, 367
 - Configuring, 368

- nsrtrap, 368
- starting, 26
- storage nodes, 2, 4
- usage, first time, 25
- volumes, 2
- Workgroup edition, 19, 46
 - autochanger sharing, 293
 - device support, 155
- Solstice Backup 6.0 Administrator's Guide, 21
- Solstice Backup 6.0 Installation Guide, 22
- Solstice Backup Administrator program, 31
 - starting, 26
 - xview, errors, 386
- staging
 - described, 177
 - filesystem devices, 160, 178
 - policies
 - creating, 178
 - deleting, 179
 - editing, 179
 - resource, 6
- status values
 - scanned-in, 241
- STLI. See silos
- storage nodes, 4
 - administrators, 67
 - affinity, 69
 - clones, 278
 - problems, 399
 - autochangers, 68, 166, 291
 - bootstrap backup, 69
 - clones, 69, 278
 - configuration, 67, 164
 - daemons, 14, 15
 - defined, 66, 164
 - devices, 68, 69
 - defining, 291
 - volume management, 298
 - installation, 68, 69, 165
 - licensing, 67
 - NDMP, 355
 - operations, 14, 15
 - prerequisites, 66
 - privileges, 165
 - programs, 14, 15
 - servers
 - conversion from, 68
 - converting from, 165

- silos, 68, 166
- timeout settings, 166
- timeouts, 400
- troubleshooting, 399, 400
- volumes
 - clones, 264
- sysconfig command, 347

T

- tapeexercise program, 405
- target sessions, 47
- technical support
 - troubleshooting information, 373
- tips. See troubleshooting
- Traps, 367
- troubleshooting
 - AIX, 427
 - core dump, 427
 - STK-9840, 428
 - archive pools, 398
 - archive requests
 - naming, 398
 - archives, 397
 - display errors, 397
 - multiple save sets, 397
 - nsrchive program, 398
 - remote request failure, 397
 - auto media verify, 378
 - autochangers, 400, 403
 - AIX, 428
 - ATL 2640 libraries, 409
 - attributes, 400
 - autodetected scsi errors, 405
 - control port access, 409
 - destination component, 407
 - drivers, 404
 - HP-UX, 416, 420, 421, 422, 423, 424, 425, 426
 - installation, 416
 - inventorying, 406
 - maintenance, 402
 - X11 libraries, 406
 - backups, 374
 - levels, 383
 - stopping, 375
 - bar codes, 310
 - bootstrap printing, failure, 385

- client file index
 - messages, missing, 383
 - size growth, 378
- clients
 - alias, 380
 - licensing, 376
 - Solaris, location, 415
 - unsynchronized errors, 375
- clones, display errors, 397
- daemons, 374
- devices, 400
 - drivers, testing, 403
 - maintenance, 402
 - nonrewinding, 409
- disk label errors, 384
- display errors, 397
- file conversion, 387
- firmware, 394
- hosts table, 389
 - ping command, testing, 391
- HP-UX, 416
 - autochanger installation, 421
 - autochanger testing, 425
 - device file, 421
 - error messages, 425
 - major number, 422
 - minor number, 423
 - pass-through drivers, 416
 - system shutdown, 420
 - system start, 421
 - testing, 424
 - unloading drives, 427
 - unsupported media, 426
- illegal characters, 381
- index failure, 377
- IP
 - errors, 388, 389
 - name search order, 390
- IRIX, 428
 - error messages, 428
- jbexercise program, failure, 405
- licensing, copy violation, 386
- lusdebug program, 403
- name resolution, 390
- name servers, disabling, 390
- naming requirements, 394
- no ecb counter, 379
- nsrexec processes, 374
- nsrexecd program, 393

- nsrjb program, 411
- online indexes
 - recovery, 380
- packet receive buffer, 379
- ping command, 391
- portmappers, verifying, 392
- ports, 393
- recoveries, 374
 - interrupted, 383
 - new installation, 382
 - online indexes, 380
 - remote access, 396
- remote access, recoveries, 396
- remote filesystems, saving, 396
- renamed client backups, 384
- retrievals, 397
 - annotations, empty, 398
- routers, 394
- rpcinfo command, 392
- scanner program, 379
 - record size, 382
- SCSI
 - bus length limitations, 404
 - IDs, 404
- server errors, binding to, 395
- server index, 385
- Solaris, 415
 - client location, 415
- storage nodes, 399, 400
- tapeexercise program, 405
- technical support information, 373
- volumes
 - capacity issues, 407
- X11 libraries, 406
- xview errors, 386

U

- uasasm program, 235
- uninventoried volumes, mounting on
 - autochanger, 307
- user ID, sending bootstrap to multiple, 192

V

- verification
 - auto media verification, 146

- virtual servers. See clusters
- volume pools. See pools
- volumes, 4, 5
 - archives, 250
 - clones, 277
 - tracking, 249
 - auto media management, 162
 - auto media verification, 146
 - autochangers, 305
 - bar codes, 311
 - depositing, 315
 - exporting, 315
 - importing, 315
 - inventorying, 313
 - labeling, 308, 309
 - mounting, 307
 - selecting, 306
 - shared, 298
 - slots, changing, 301
 - unmounting, 299, 308
 - withdrawing, 316
 - bad, removal, 177
 - bar codes, labeling, 310, 311
 - capacity issues, 407
 - client file index
 - removing, 77, 78
 - clones, 263, 266
 - archives, 277
 - creating, 267
 - deleting, 266
 - performance, 266
 - recovery, 267
 - recycling, 269
 - information
 - viewing, 171
 - labeling, 26, 148, 167, 173, 175
 - bar codes, 310, 311
 - time, 169
 - tips, 175
 - labels
 - bar codes, 312
 - unmatching, 312
 - logical backup, 160
 - management
 - server devices, 298
 - storage nodes, 298
 - media database
 - removing, 78
 - removing entries, 77
 - modes, 169, 170, 266
 - appendable, 170
 - changing, 176
 - manual recycle, 170
 - read-only, 170, 171
 - recyclable, 170
 - mounting, 26, 173, 307
 - names
 - reuse, 312
 - names, finding, 176
 - nonrewinding, 409
 - operations, 171
 - recoveries
 - clones, 277
 - required, 221
 - recyclable, 266
 - recycling, 76, 266, 307
 - comparison, 317
 - relabeling
 - selection criteria, 169
 - removing, 76
 - save sets
 - clones, 269
 - information, 172
 - recoveries, 232
 - scanner program
 - usage, 240
 - selection criteria, 168
 - silos, 320
 - allocating, 328
 - deallocating, 329
 - exporting, 326
 - importing, 326
 - labeling, 328
 - mounting, 327
 - unlabeled, 328
 - unmounting, 327
 - slots, depositing, 410
 - status, 169
 - StorageTek, ejecting, 330
 - unlabeled, 162
 - unmounting, 175

X

- xview, errors, 386

