

IRIX NetWorker™ Administrator's Guide

Document Number 007-1458-050

CONTRIBUTORS

Written by Celia Szente, Rebecca Mason, Vickie Brown, and Bill Tuthill
Production by Michael Dixon
Engineering contributions by Legato employees

© 1996-1997, Legato Systems Corp. and Silicon Graphics, Inc.— All Rights Reserved
The contents of this document may not be copied or duplicated in any form, in whole or in part, without the prior written permission of Silicon Graphics, Inc.

RESTRICTED RIGHTS LEGEND

Use, duplication, or disclosure of the technical data contained in this document by the Government is subject to restrictions as set forth in subdivision (c) (1) (ii) of the Rights in Technical Data and Computer Software clause at DFARS 52.227-7013 and/or in similar or successor clauses in the FAR, or in the DOD or NASA FAR Supplement. Unpublished rights reserved under the Copyright Laws of the United States. Contractor/manufacturer is Silicon Graphics, Inc., 2011 N. Shoreline Blvd., Mountain View, CA 94043-1389.

Silicon Graphics and the Silicon Graphics logo are registered trademarks, and IRIX, IRIX NetWorker, and IRIS InSight are trademarks of Silicon Graphics, Inc.

AIX is a trademark of International Business Machines Corp.
Exabyte is a trademark of Exabyte Corporation.
Informix is a registered trademark of Informix Software, Inc.
Legato NetWorker is a registered trademark of Legato Systems Corp.
Macintosh is a registered trademark of Apple Computer, Inc.
Motif is a trademark of the Open Software Foundation.
NetWare is a registered trademark of Novell, Inc.
NFS is a registered trademark of Sun Microsystems, Inc.
Oracle is a registered trademark of Oracle Corporation.
PostScript is a registered trademark of Adobe Systems, Inc.
SunOS and Solaris are trademarks of Sun Microsystems, Inc.
Sybase is a registered trademark of Sybase, Inc.
UNIX is a registered trademark of Silicon Graphics, Inc. of X/Open Company, Ltd.
Windows NT is a registered trademark of Microsoft Corporation.
X Window System is a trademark of the Massachusetts Institute of Technology.

IRIX NetWorker™ Administrator's Guide
Document Number 007-1458-050

Contents

List of Figures xv

List of Tables xxi

About This Guide xxiii

Intended Audience xxiii

Book Organization xxiii

Conventions xxiv

Documentation and Support xxv

Definition of Terms xxv

NetWorker Product Family xxvi

1. Using NetWorker Windows and Menus 1

Starting the NetWorker Programs 1

 Running the NetWorker Administrator Program 2

 Running the NetWorker Backup Program 5

 Running the NetWorker Recover Program 6

Using NetWorker Menus 6

 File Menu 7

 Printing Window Contents 7

 Saving Window Contents 8

 Exiting NetWorker 9

 Clients Menu 9

 Server Menu 10

 Media Menu 10

 Save Set Menu 11

 Customize Menu 11

- Options Menu 12
 - Changing the Polling Interval 12
 - Displaying the Speedbar 13
 - Setting the Tone Option 13
- Help Menu 13
- View Menu 14
 - Displaying a Tabular View 14
 - Displaying Details 17
 - Searching for Resources 19
- Generic Buttons 20
 - Create Multiple Button 20
 - Print and Save Buttons 21
- Summary 22
- 2. Configuring and Monitoring Clients 23**
 - Configuring NetWorker Clients 23
 - Navigating the Clients Window 23
 - Using the Pre- and Post- Processing Backup Command 27
 - Example—Backup Command 28
 - Configuring Clients 29
 - Adding a New Client 30
 - Scheduling Large Client Filesystems 32
 - Removing a Client 34
 - Securing the NetWorker Environment 35
 - Restricting Client Access 35
 - Migration Setup 36
 - Manually Managing the Online Indexes 36
 - Online Indexes 37
 - Save Sets 37
 - Index Actions 39
 - Generating File Index Reports 39
 - Navigating the Indexes Window 40

-
- Removing the Oldest Cycle 42
 - Reclaiming Index Space 44
 - Viewing Save Set Details 45
 - 3. Configuring and Monitoring the Server 47**
 - Configuring the NetWorker Server 47
 - Navigating the Server Window 48
 - Changing Parallelism 50
 - Using Multiple Devices Concurrently 51
 - Adding or Changing Administrators 52
 - Changing the NetWorker Server 53
 - Monitoring and Controlling Backups 54
 - Navigating the Group Control Windows 54
 - Group Control Details Window 56
 - Previewing a Group Backup 57
 - Starting a Scheduled Backup Immediately 57
 - Stopping a Backup Group 58
 - Restarting a Backup Group 59
 - Handling Open Files 59
 - Archive Request Control 60
 - Registering NetWorker 60
 - 4. Managing Media and Backup Devices 61**
 - Labeling and Mounting Backup Volumes 61
 - Navigating the Label Dialog Box 62
 - Labeling and Mounting in One Operation 63
 - Understanding How to Label Backup Volumes 64
 - Mounting and Unmounting Backup Volumes 65
 - Mounting Backup Volumes 66
 - Unmounting Backup Volumes 67
 - Finding a Backup Volume Name 67
 - Automatically Managing Backup Volumes 67

- Manually Managing Backup Volumes 68
 - Navigating the Volumes Window 68
 - Volumes Scrolling List 69
 - Save Sets Scrolling List 71
 - Manually Managing Backup Volumes 72
 - Changing the Status of a Save Set 73
 - Setting the Location for Backup Volumes 73
 - Changing the Mode of a Backup Volume 74
 - Recycling Volumes 76
 - Removing a Backup Volume 77
- Cloning Backup Volumes 78
 - Navigating the Clone Volume Windows 79
 - Example—Cloning Backup Volumes 82
 - Understanding the Volume Cloning Feature 83
 - Entering Index Entries for Cloned Data 84
 - Recovering from Cloned Volumes 84
- Using Volume Pools 85
 - Navigating the Pools Window 86
 - Examples 90
 - Example 1—Creating a New Pool 90
 - Example 2—Separate Departments 93
 - Example 3—Secure Device 93
 - Example 4—Targeted Backups 93
 - Example 5—Specific Filesystem 94
 - Shortcut 94
 - Deleting a Volume Pool 95
 - Understanding Volume Pools 95
 - Using Label Templates With Volume Pools 95
- Adding and Deleting Backup Devices 96
 - Navigating the Devices Window 96
 - Adding a New Backup Device 99
 - Deleting a Backup Device 99
 - Using the Auto Media Management Feature 100

5.	Recovering and Cloning Save Sets	101
	Recovering Save Sets	101
	Navigating the Save Set Recover Window	102
	Save Set Recover Details Window	103
	Save Set Recover Status Dialog Box	104
	Examples	108
	Example 1—Recovering Individual Files	108
	Example 2—Recovering a Disk Partition	109
	Shortcut	111
	Understanding the Save Set Recover Feature	111
	Preserving Access Control Lists	111
	Recovering Nonbrowsable Data	112
	Recovering Large Amounts of Data	112
	Cloning Save Sets	113
	Navigating the Save Set Clone Window	114
	Save Set Clone Window	114
	Save Set Clone Details Window	117
	Save Set Clone Status Dialog Box	118
	Additional Browsing Choices	119
	Examples	121
	Example 1—Manually Cloning Save Sets	121
	Example 2—Automatically Cloning One Client	122
	Example 3—Automatically Cloning Several Clients	123
	Shortcut	128
	Understanding the Cloning Save Sets Feature	128
	Automatic Cloning	129
	Using Clone Pools	130

- 6. Customizing NetWorker Backups 131**
 - Setting Up Backup Schedules 131
 - Navigating the Schedules Window 132
 - Examples 134
 - Example 1—A Weekly Schedule 134
 - Example 2—A Monthly Schedule 136
 - Shortcut 139
 - Understanding Backup Schedules 140
 - Selecting the Backup Levels 141
 - Creating Index Policies 142
 - Navigating the Policies Window 143
 - Creating a New Policy 145
 - Deleting a Policy 145
 - Changing a Policy 146
 - Understanding Index Policies 146
 - Using Directives 148
 - Navigating the Directives Window 148
 - Adding or Changing Directives 149
 - Understanding Directives 149
 - Using Label Templates 150
 - Choosing a Label Template 151
 - Navigating the Label Templates Window 152
 - Examples 153
 - Example 1—A Storage-Oriented Template 154
 - Example 2—A Sequential Label Template 156
 - Example 3—A Two-Sided Media Label Template 157
 - Configuring Backup Groups 159
 - Navigating the Groups Window 159
 - Creating a New Group 161
 - Retrying a Failed Client Backup 162
 - Changing a Backup Time 163
 - Deleting a Backup Group 163

Setting Up Event Notification	164
Navigating the Notifications Window	164
Creating a New Notification	165
Summary	166
7. Autochanger Software Module	167
Configuring Jukeboxes	167
Navigating the Jukeboxes Window	168
Adding More Jukebox Devices	171
Changing Available Slots	173
Cleaning Jukebox Devices	174
Automatically Managing Media	175
Operating Jukeboxes	176
Jukebox Mounting Dialog Box	176
Jukebox Labeling Dialog Box	177
Starting a Labeling Session	178
Labeling With Bar Codes	180
Understanding Labeling With Bar Codes	182
Jukebox Inventory Dialog Box	184
Inventorying Backup Volumes	184
Disabling the Element Status Feature	186
Understanding the Inventory Process	187
Checking Jukebox Notifications	187
Suggestions for Operating Jukeboxes	188
Recycling Versus Adding More Backup Volumes	188
Using the Jukebox Cartridge Access Port	189
Depositing a Cartridge	189
Withdrawing a Cartridge	190
Using Pools With a Jukebox	190
Calibrating the Devices	190

- 8. Archive Application 191**
 - Navigating the Windows 191
 - Clients Window 191
 - Archive Requests Window 193
 - Archive Request Control Window 195
 - Archive Request Details 196
 - Archive Example 197
 - Creating an Archive Client 197
 - Making an Archive Request 197
 - Checking the Archive Request 198
 - Rescheduling the Archive Request 199
 - Clone and Verify 200
 - Archiving Shortcut 201
 - Understanding the Archive Feature 201
 - Archive Functions 202
 - Data Archiving 202
 - Data Verification 202
 - Data Retrieval 203
 - Methods for Protecting Data 204
- A. Recovering From a Disk Crash 205**
 - Preparing for a Crash 205
 - File the Bootstrap Information 207
 - File the Disk Information 207
 - Using Save Set Recover or Recover 209
 - Recovering a Secondary Disk 212
 - Recovering a Primary Disk 214
 - Recovering /nsr on a NetWorker Server 215
 - Replacing the /nsr/res Directory 218
 - Finding the Bootstrap Save Set ID 219
 - Recovering to a New Server 219
 - Disaster Recovery With Jukeboxes 220
 - Summary 222

B.	Managing the NetWorker Environment	223
	Guidelines for Choosing a Configuration	223
	Tape Drives	223
	Clients	224
	Network	224
	Server	224
	Jukeboxes	225
	Guidelines	226
	Configuration Examples	227
	Example 1	227
	Example 2	227
	Measuring Performance	228
	Server Performance	228
	Tape Drive Speed	228
	Network Speed	229
	CPU Speed	230
	Memory	230
	Client Performance	230
	Filesystem Traversing	230
	Data Generation Rate	230
	Data on Multiple Disks	231
	CPU Speed	231
	NetWorker Backup Schedules	231
	Backup Levels	232
	Full Backups Versus Incremental Backups	233
	Using Level Backups	233
	Typical Monthly Backup Schedule	235
	Backups Take Time	236
	Using Compression During Backup	237
	Staggering the Backup Schedules	238
	Convenience Versus Security	239

- NetWorker Browse and Retention Policies 239
 - Browse Policy 239
 - How Browse Policies Work 240
 - Reclaiming Disk Space 241
 - Recovering Files Removed From the Index 241
 - Media Retention Policy 242
 - Setting Policies When Using a Jukebox 243
 - Jukebox Capacity 243
 - Choosing the Jukebox Capacity 244
 - Choosing a Jukebox 245
- Preconfigured Selections 246
 - Preconfigured Backup Schedules 246
 - Preconfigured Policies 249
 - Preconfigured Pools 251
 - Preconfigured Label Templates 253
 - Preconfigured Directives 255
 - Preconfigured Notifications 255
 - Registration 256
 - Log Default 256
 - Index Size 256
 - Savegroup Completion 257
 - Backup Media Request Notices 257
- Summary 258
- C. Troubleshooting 259**
 - Checking the NetWorker Daemons 259
 - Displaying NetWorker 259
 - Renaming a Client 260
 - Recover Access Issues 261
 - Previewing a Backup 261
 - Halting a Network Backup 262
 - Backup Media Capacity 263
 - Determining Jukebox Capacity 263

- Savegroup Completion Messages 264
 - Binding to Server Errors 264
 - Saving Remote Filesystems 265
 - File Changed During Backup 266
 - Cannot Print Bootstrap Information 266
 - Copy Violation 266
- Maximum Filename Length 267
- Savegroup Completion Warning Messages 267
- X11 Errors 268
- Moving Indexes 268
- Recovering Files From an Interrupted Backup 269
- Determining the NetWorker Server 269
- Using nsrexecd 269
- D. Command Summary 271**
 - NetWorker System Commands 271
 - NetWorker Maintenance Commands 272
- E. Frequently Asked Questions 275**
 - Licensing and Installation 275
 - Tapes and Devices 277
 - Autochangers (Jukeboxes) 279
 - Backup Schedules 284
 - Browse and Retention Policies 284
 - Backing Up Data 285
 - Recovering Data 287
 - Directives 290
 - Compression 291
 - Logging and Notification 292
- Glossary 293**
- Index 301**

List of Figures

Figure 1-1	Notice About Administrator List	1
Figure 1-2	NetWorker Administrator Icon	2
Figure 1-3	NetWorker Administrator Window	3
Figure 1-4	Speedbar Buttons	4
Figure 1-5	NetWorker Administrator Pending Icon	5
Figure 1-6	NetWorker Backup Icon	5
Figure 1-7	NetWorker Recover Icon	6
Figure 1-8	File Menu	7
Figure 1-9	File Menu (Save and Print)	7
Figure 1-10	Print Options Window	8
Figure 1-11	Save Options Window	8
Figure 1-12	Clients Menu	9
Figure 1-13	Server Menu	10
Figure 1-14	Media Menu	10
Figure 1-15	Save Set Menu	11
Figure 1-16	Customize Menu	11
Figure 1-17	Options Menu	12
Figure 1-18	Change Polling Dialog Box	12
Figure 1-19	Speedbar Buttons	13
Figure 1-20	Help Menu	13
Figure 1-21	View Menu	14
Figure 1-22	Clients Window	15
Figure 1-23	Clients Window, Tabular View	16
Figure 1-24	Groups Window	17
Figure 1-25	Groups Window, Details	18
Figure 1-26	Search for Resource Dialog Box	19
Figure 1-27	Notice for Resource Not Found	19

Figure 1-28	Create Multiple Button	20
Figure 1-29	Print Options Dialog Box for Hardcopy	21
Figure 1-30	Save Option Dialog Box for Report	22
Figure 2-1	Clients Window as Preconfigured	24
Figure 2-2	Clients Window, Scrolled Down	26
Figure 2-3	Save Set in Clients Window	31
Figure 2-4	Remote Access in Clients Window	31
Figure 2-5	Filesystems as Client Save Sets	34
Figure 2-6	Full Versus Incremental Backups	38
Figure 2-7	Indexes Button	40
Figure 2-8	Indexes Window	41
Figure 2-9	Removing Oldest Cycle From the Index	43
Figure 2-10	Reclaim Space Window	44
Figure 2-11	Reclaiming Space From the Index	45
Figure 2-12	Instances Window	46
Figure 3-1	Server Window	48
Figure 3-2	Server Parallelism, Active Devices, Sessions per Device	51
Figure 3-3	Server Administrator List	52
Figure 3-4	Change Server Button	53
Figure 3-5	Change Server Window	53
Figure 3-6	Group Control Button	54
Figure 3-7	Group Control Window	55
Figure 3-8	Group Control Details Window	56
Figure 4-1	Label Button	62
Figure 4-2	Label Dialog Box	62
Figure 4-3	Mount and Umount Buttons	66
Figure 4-4	Volumes Button	68
Figure 4-5	Volume Management Window	69
Figure 4-6	Volume Save Sets List	71
Figure 4-7	Save Set Change Status Menu	73
Figure 4-8	Set Location Dialog Box	74
Figure 4-9	Volume Change Mode Menu	75
Figure 4-10	Notice to Change Volume Mode	76

Figure 4-11	Volume Recycle Menu	76
Figure 4-12	Volume Remove Menu	77
Figure 4-13	Notice to Remove Volume	78
Figure 4-14	Media Clone Volume Menu	79
Figure 4-15	Clone Volume Window	80
Figure 4-16	Clone Volume Details Window	80
Figure 4-17	Clone Volume Status Dialog Box	81
Figure 4-18	Media Pools Menu	86
Figure 4-19	Pools Window	87
Figure 4-20	Notice About Pool Overlap	88
Figure 4-21	Label Templates Window (Test)	91
Figure 4-22	Pools Window (Test)	92
Figure 4-23	Devices Window	97
Figure 4-24	Notice About Auto Media Management	100
Figure 5-1	Save Set Recover Window	102
Figure 5-2	Save Set Recover Details Window	104
Figure 5-3	Save Set Recover Status Dialog Box	105
Figure 5-4	Save Set Recover Volumes Required Window	106
Figure 5-5	Save Set Recover Options Dialog Box	107
Figure 5-6	Save Set Clone Menu	114
Figure 5-7	Save Set Clone Window	115
Figure 5-8	Save Set Clone Details Window	117
Figure 5-9	Save Set Clone Status Dialog Box	118
Figure 5-10	Save Set Clone Window (More)	119
Figure 5-11	Sort By Field for Save Set Clone	120
Figure 5-12	Automatic Cloning Scheme	124
Figure 5-13	Acct Incr Schedule	125
Figure 5-14	Acct Full Clone Schedule	125
Figure 5-15	Acct Incr Groups Window	126
Figure 5-16	Acct Full Clone Groups Window	127
Figure 6-1	Schedules Window	132
Figure 6-2	Schedules Window, Backup Level Menu	135
Figure 6-3	Schedules Window, Level Override Menu	138

Figure 6-4	Backup Levels, Full and Incremental	141
Figure 6-5	Policies Window	144
Figure 6-6	Notice to Delete Policy	146
Figure 6-7	Automatic Index Management	147
Figure 6-8	Directives Window	148
Figure 6-9	Pools Window	151
Figure 6-10	Label Templates Window	152
Figure 6-11	Rack Storage of Backup Volumes	154
Figure 6-12	Label Templates, Storage-Oriented	156
Figure 6-13	Label Templates, Two-Sided	158
Figure 6-14	Groups Window	160
Figure 6-15	Client Retries Slider Bar	162
Figure 6-16	Notifications Window	165
Figure 7-1	Jukeboxes Window	168
Figure 7-2	Jukebox Devices Field	169
Figure 7-3	Devices Window	171
Figure 7-4	Pathname in Jukebox Devices Field	172
Figure 7-5	Jukeboxes With Multiple Devices	173
Figure 7-6	Jukeboxes Available Slots Field	173
Figure 7-7	Jukebox Mounting Dialog Box	176
Figure 7-8	Jukebox Labeling Dialog Box	177
Figure 7-9	Notice to OK Jukebox Labels	179
Figure 7-10	Jukebox Labeling More Tapes	179
Figure 7-11	Jukebox Labeling One Tape	180
Figure 7-12	Jukebox Labeling for Bar Codes	181
Figure 7-13	Volume Management Window	182
Figure 7-14	Jukebox Inventory Dialog Box	185
Figure 7-15	Jukebox Inventory Finished	185
Figure 7-16	Jukebox Options Field	186
Figure 8-1	Enable Archive in Clients Window	192
Figure 8-2	Archive Requests Window	193
Figure 8-3	Archive Request Control Window	195
Figure 8-4	Archive Request Details Window	196

Figure 8-5	Archive Request Schedule Window	199
Figure 8-6	Archive Process Flow Chart	203
Figure A-1	Backups Needed for Disaster Recovery	210
Figure A-2	Save Set Recover Window	211
Figure A-3	Secondary Disk Damaged	213
Figure A-4	Primary Disk Damaged	214
Figure B-1	Full Backup Followed by Levels 9 to 1	234
Figure B-2	Full Backup Followed by Levels 1 to 9	235
Figure B-3	Full Backup Plus Incrementals and Level 5	236
Figure B-4	Backups and the Browse Period	240
Figure B-5	Schedules Window	247
Figure B-6	Policies Window (Details)	250
Figure B-7	Registration Notification	256
Figure B-8	Log Default Notification	256
Figure B-9	Index Size Notification	257
Figure B-10	Savegroup Completion Notification	257
Figure B-11	Tape Mount Request Notification	258
Figure C-1	Group Control Stop Button	262
Figure C-2	Jukebox Mounting Window	263

List of Tables

Table 4-1	Changing Volume Modes Manually	75
Table 5-1	Sample Acct Incremental Schedule	123
Table 5-2	Sample Acct Full Clone Schedule	123
Table 6-1	Label Template Example	157
Table 7-1	Cleaning Cartridge Notifications	174
Table B-1	Network Speed Comparison	224
Table B-2	Server Throughput Comparison	225
Table B-3	Amount of Backup Data at Different Levels	225
Table B-4	Sample Monthly Data Backup	226
Table B-5	Capacity of Different Autochangers	226
Table B-6	Staggered Backup Schedules Reduce Load	228
Table B-7	Storage Requirements for One Backup Cycle	245
Table C-1	Where to Start nsrexecd	270
Table D-1	Maintenance Commands and Reference Pages	272

About This Guide

The *IRIX NetWorker Administrator's Guide* contains information on how to configure and manage the IRIX NetWorker™ backup and recover software. Information on configuring and managing optional NetWorker add-on products available for complete network storage management solutions is also included.

This manual provides enough information for the basic use and administration of IRIX NetWorker installed in a typical networked environment. For more technical information on IRIX NetWorker commands, refer to the online reference pages after you have installed the software.

Intended Audience

This manual is intended for administrators who install software and maintain the servers and clients on a network. Operators who monitor the daily backups of the machines and maintain the backup media and devices may also find this manual useful.

Book Organization

The information in this manual has been organized to help you become familiar with the graphical user interface, get you started using NetWorker, and assist you in configuring and managing your NetWorker backups.

Where possible, information is divided into three sections. These sections help you locate the kind of information you need to accomplish a NetWorker task or understand a feature. A description of each section follows:

- Navigating—explains the contents of a window and how to accomplish the associated tasks.
- Examples—provide examples of NetWorker tasks and configurations.
- Understanding—provides a thorough explanation of a NetWorker feature.

Conventions

The following conventions are used in this manual to make information easy to access and understand.

- Command names are displayed in **bold** typeface. For example:
To start the NetWorker Administrator program, use the **nwadmin** command.
- Examples of what *you* type are shown in bold fixed-width typeface. For example:

```
# nwadmin &
```
- Entries that you type and are substituted by a variable are displayed as a bold italic fixed-width typeface. For example:

```
% man nsr
```
- Examples, shell prompts, and information displayed on the screen are displayed in fixed-width typeface. For example:

```
media waiting: recover waiting for 8mm 5GB tape volume name
```
- Names of NetWorker buttons, displays, menus, scrolling lists, and windows are displayed in Helvetica typeface. For example:
Click the Cancel button to close the Help window.
- Directory *pathnames*, machine names, and new terms defined in the glossary are displayed in *italic* typeface. For example:
When you use IRIX NetWorker to back up your */usr* files from the *client* machine, you are saving them to *backup volumes* on the IRIX NetWorker server.
- The following paragraph types indicate various kinds of information to help you use NetWorker productively.

Caution: Important pieces of information and cautionary notes that prevent you from making a mistake are marked "Caution."

Note: Helpful information that you might want to know is marked "Note."

Tip: Tips or suggestions that you do not necessarily have to follow, but may give you hints as to how to set up IRIX NetWorker at your site, are marked "Tip."

Shortcut: Step-by-step procedures that help you save time because they provide the minimum information you need to complete a task are marked as "Shortcut."

Documentation and Support

Both the *IRIX NetWorker Administrator's Guide* and *IRIX NetWorker User's Guide* are available online, as IRIS InSight™ books. To view the manuals, install the subsystems `networker4.books.NetWorker_AG` and `networker4.books.NetWorker_UG`.

To learn how to use the NetWorker Backup and Recover windows for manual backups, see the *IRIX NetWorker User's Guide*. To get information about a specific release of IRIX NetWorker, see the *IRIX NetWorker Release Notes*. If you would like more technical information about the NetWorker commands, see the online reference pages after you have installed NetWorker.

Silicon Graphics® offers a comprehensive product support and maintenance program for IRIX products. For information about using support services for this product, refer to the release notes that accompany it.

Definition of Terms

A set of client files backed up by IRIX NetWorker is called a *save set*. Save sets are written to *backup volumes* or *media* (for example, tapes or optical disk) mounted in a backup device attached to the NetWorker server. A collection of backup volumes tracked by NetWorker is a *media set*.

NetWorker uses *volume pools* to sort specific data to preselected backup volumes to help manage the media set.

The client save sets and backup volumes are maintained in two online indexes on the server: the *file index* and the *media index*. IRIX NetWorker uses these two indexes to track the saved files and the volumes to which they are backed up.

When backing up files, IRIX NetWorker may request that a writable backup volume be mounted in the server backup device. When recovering files, IRIX NetWorker requests a specific backup volume by name. You need to mount the requested backup volume(s) in a device attached to the server. If you use an autochanger (a jukebox) for backups, the requested volume is automatically loaded if it is stored in it.

The terms *autochanger* and *jukebox* are used interchangeably throughout this manual. Both terms refer to a variety of backup devices, including autoloader, carousel, library, near-line storage, and datawheel.

NetWorker Product Family

In addition to NetWorker products for UNIX® based servers and TCP/IP networks, Legato Systems Corporation offers other software products that provide additional network backup options.

- NetWorker for Windows NT® provides support for Windows NT servers and TCP/IP networks, and can back up a wide variety of UNIX systems, NetWare® systems, Macintosh® desktops, and PC desktops.
- NetWorker for NetWare backs up and recovers data on NetWare networks and is fully interoperable with UNIX systems.
- NetWorker ClientPak™ software allows a wide variety of client platforms on heterogeneous networks to back up to and recover data from NetWorker servers.
- The NetWorker Autochanger Software Module allows you to fully realize the automatic backup capabilities of NetWorker.
- NetWorker database modules reliably back up the most widely used databases, including Oracle®, Informix®, SQL Server®, and Exchange®.
- The NetWorker Archive Application provides you with a sophisticated method for archiving valuable client data.

All NetWorker server products support the latest backup devices that provide the highest performance and most reliable solutions for your network backup and recovery requirements.

Using NetWorker Windows and Menus

NetWorker is a software product that reliably protects against loss of files across an entire network of systems. NetWorker saves valuable administrator time by speeding and simplifying routine daily backup operations. An easy-to-use graphical user interface (GUI) guides you through administering and configuring your network for scheduled backups. As NetWorker backs up your data, it creates a database of the saved data, making it easy to locate a file for recovery. As the network and number of files expand, NetWorker has the capacity and performance to handle the load.

Starting the NetWorker Programs

The NetWorker software package consists of three programs:

- NetWorker Administrator—for managing and configuring NetWorker.
- NetWorker Backup—for backing up data.
- NetWorker Recover—for recovering files.

Users displayed on the Administrators list in the Server window have permission to use and maintain NetWorker Administrator configurations. Users not included on the list can start and open the NetWorker Administrator window. However, any attempt to modify settings will result in the display of a message like the one in Figure 1-1.



Figure 1-1 Notice About Administrator List

Running the NetWorker Administrator Program

Use the NetWorker Administrator window to manage, configure, and monitor NetWorker. The window serves a dual purpose: as a starting point for NetWorker tasks and as a monitor for server activity.

If the `nwadmin` icon already appears on your screen, simply click to open it. The icon (shown in Figure 1-2) includes a graphic and the hostname of the currently selected NetWorker server.



Figure 1-2 NetWorker Administrator Icon

To easily identify unique servers in a multi-server network, NetWorker includes the current server name in the icon. To change the icon to display the `nwadmin` caption, use X11 resources to change the resource `iconServerName`, following these steps:

1. Enter the `xrdb -m` command at the system prompt.
2. Change the `iconServerName` resource setting from `True` to `False`.

Refer to the *X Window System User's Guide* for more detailed instructions.

If the `nwadmin` icon is not currently displayed on your screen, start the program from the system prompt by entering the `nwadmin` command:

```
# nwadmin &
```

If you have a multi-server network and do not specify a server name, NetWorker Administrator starts on the currently selected server. To specify a server, enter the following command:

```
# nwadmin -s server_name &
```

The NetWorker Administrator window appears, as shown in Figure 1-3.

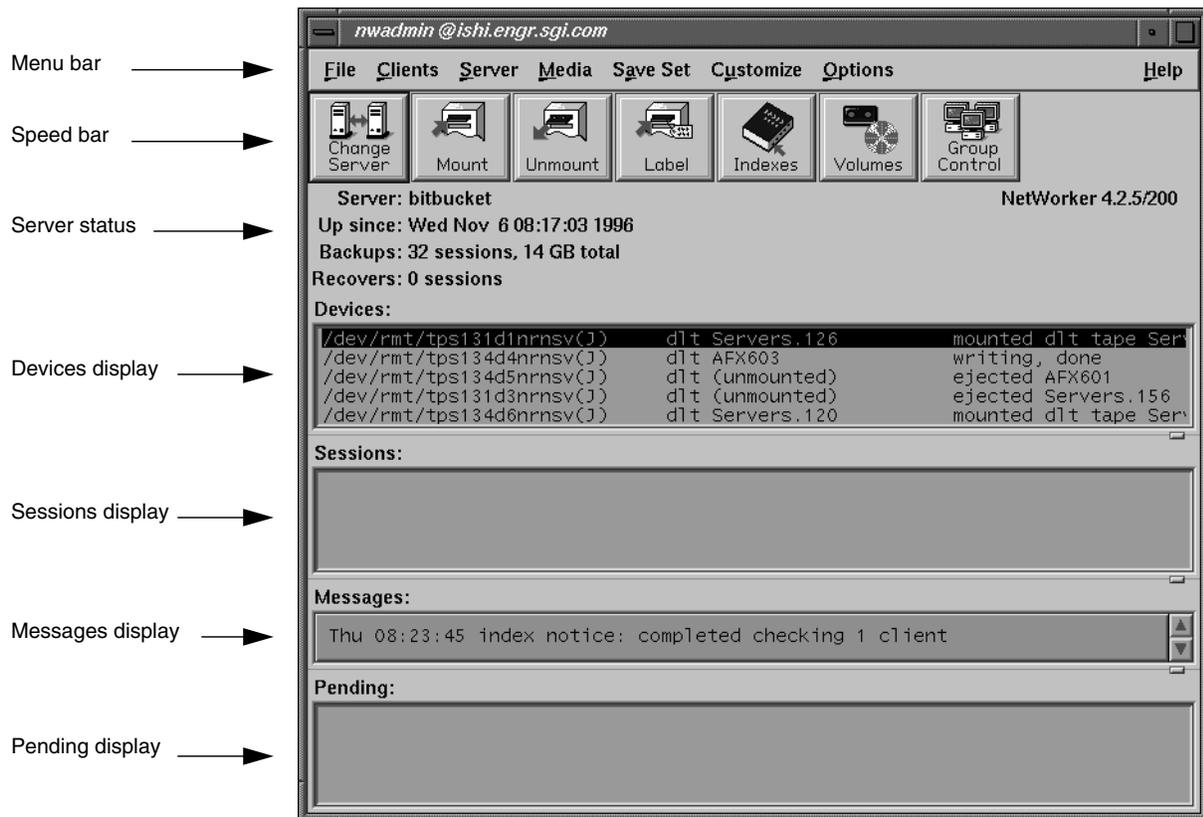


Figure 1-3 NetWorker Administrator Window

The NetWorker Administrator window has a menu bar, speedbar, and five displays.

- Menu bar—contains the File, Clients, Server, Media, Save Set, Customize, Options, and Help menus.
- Speedbar—provides convenient shortcut buttons for performing common NetWorker tasks. Throughout the manual, where appropriate, speedbar buttons are provided as an alternative to menu commands. Click a speedbar button to open its associated window (see Figure 1-4).

The speedbar buttons and their functions are described below:



Figure 1-4 Speedbar Buttons

- Change Server button—opens the Change Server dialog box; use this button instead of choosing Change from the Server menu.
 - Mount button—mounts a backup volume in a backup device; use this button instead of choosing Mount from the Media menu.
 - Unmount button—unmounts a backup volume in the backup device; use this button instead of choosing Unmount from the Media menu.
 - Label button—labels a volume in the backup device; use this button instead of choosing Label from the Media menu.
 - Indexes button—opens the Indexes window; use this button instead of choosing Indexes from the Clients menu.
 - Volumes button—opens the Volume Management window; use this button instead of choosing Volumes from the Media menu.
 - Group Control button—opens the Group Control window; use this button instead of choosing Group Control from the Server menu.
- Server status display—shows how long the current server has been online, the number of backups and restores the server completed since it has been online, and the amount of data the backups contain.
 - Devices display—shows the devices known to the server, backup volumes currently mounted, and the device status.
 - Sessions display—provides current information about which clients are backing up, browsing, or recovering files.
 - Messages display—provides a scrollable history of the server activity during the last 24 hours.
- Note:** The daily information shown in the Messages display is also added to the NetWorker `/nsr/logs/daemon.log` file. You can search the contents of this log file for messages no longer viewable in the Messages display.
- Pending display—displays messages requiring attention.

When an event requiring operator intervention occurs, the program icon changes to alert you that NetWorker Administrator is in a “pending” mode (see Figure 1-5).



Figure 1-5 NetWorker Administrator Pending Icon

Running the NetWorker Backup Program

Use the NetWorker Backup window to back up client files. This window does not provide the ability to administer or configure NetWorker. The backup program icon includes the “NetWorker knight” graphic and the caption `nwbackup`.

If the `nwbackup` icon (see Figure 1-6) is already displayed on your screen, click it to open the window.



Figure 1-6 NetWorker Backup Icon

If the `nwbackup` icon is not currently displayed on your screen, start the software program by entering the `nwbackup` command at the system prompt:

```
# nwbackup &
```

If you have a multi-server network and do not specify a server name, NetWorker Backup starts on the currently-selected server. To specify a server, enter the following command:

```
# nwbackup -s server_name &
```

Refer to the *IRIX NetWorker User's Guide* for detailed information about Backup.

Running the NetWorker Recover Program

Use the NetWorker Recover window to recover client files. This window does not provide the ability to administer or configure NetWorker. The recover program icon includes the “NetWorker knight” graphic and the caption `nwrecover`.

If the `nwrecover` icon (see Figure 1-7) is already displayed on your screen, click it to open the window.



Figure 1-7 NetWorker Recover Icon

If the `nwrecover` icon is not currently displayed on your screen, start the software program by entering the `nwrecover` command at the system prompt:

```
# nwrecover &
```

If you have a multi-server network and do not specify a server name, NetWorker Backup starts on the currently-selected server. To specify a server, enter the following command:

```
# nwrecover -s server_name &
```

Refer to the *IRIX NetWorker User's Guide* for detailed information about Recover.

Using NetWorker Menus

This section includes information about pulldown menus in the NetWorker Administrator window and how to use them. There are two choices for navigating the NetWorker menus: the point and select method with a mouse or keyboard shortcuts.

- To pull down a menu using a mouse, move the mouse so the arrow cursor points to your menu choice. Hold down the primary mouse button (usually on the left); the menu appears. With the mouse button still held down, move the mouse arrow to the desired command in the menu, and then release the button to make your choice.
- To pull down a menu using the keyboard shortcuts, follow these steps:

1. Hold down the [Alt] key.
2. Press the key that corresponds to the underlined letter of the desired menu.
3. Release both keys, then press the key that corresponds to the underlined letter in the desired command name.

File Menu

The File menu in the main NetWorker Administrator window contains the Exit command for exiting NetWorker (see Figure 1-8).



Figure 1-8 File Menu

Other windows opened through NetWorker menus have different versions of the File menu containing the Save, Print, and Exit commands (see Figure 1-9).

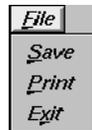


Figure 1-9 File Menu (Save and Print)

The Save and Print commands are available only when a NetWorker window is in the Tabular view. See “Displaying a Tabular View” on page 14 for more information.

Printing Window Contents

Use the Print command to print the contents of a NetWorker window displayed in the tabular view.

The ability to print the contents of a window makes it easy to maintain records of NetWorker activities and configurations.

Caution: Do not print the contents of a window if you choose Details from the View menu. The printer output will not be readable.

To print the contents of a NetWorker window, follow these steps:

1. Choose Tabular from the View menu. The window contents change to a tabular view.
2. Choose Print from the File menu. The Print Options dialog box appears (Figure 1-10).



Figure 1-10 Print Options Window

The Command field displays the “lp” default, or the NWPRINTER environment variable, if set. You may change the print command.

3. Enter the desired UNIX print command in the Command field.
4. Enter the number of copies in the Copies field.
5. Click Ok to print.

Saving Window Contents

Use the Save command when you are in the Tabular view of a NetWorker window. The Save command allows you to save the contents of a window for future reference.

To save the contents of a NetWorker window to a file, follow these steps:

1. Choose Tabular from the View menu. The window contents change to a tabular view.
2. Choose Save from the File menu. The Save Options dialog box appears (Figure 1-11).

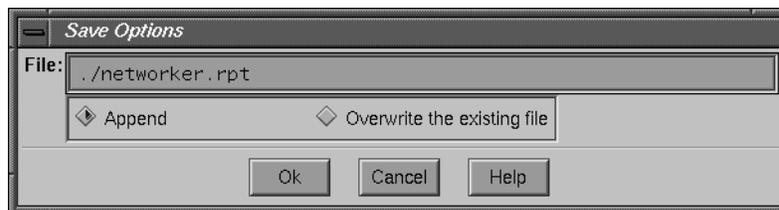


Figure 1-11 Save Options Window

3. Enter the name of the file in the File field.

The information is saved to a file specified by the environment variable NWREPORT, if set. Otherwise, the information is saved to *./networker.rpt*. You can save the data to a different data file by entering its name into the File field.

4. Specify an option for NetWorker to follow in case the file already exists. To update or append the existing file, click the Append button. To overwrite the existing file, click the Overwrite the existing file button.
5. Click Ok to save the contents of the window to the specified file.

Exiting NetWorker

There are two methods for exiting NetWorker to return to the system prompt or to exit a NetWorker command window to return the main NetWorker program window.

- Choose Exit from the File menu.
- Press both the [Alt] key and the [F] key at the same time, release them, and then press [X].

Clients Menu

The Clients menu (shown in Figure 1-12) contains commands for configuring your NetWorker clients and for monitoring and managing their indexes.

See Chapter 2, “Configuring and Monitoring Clients,” for a description of the commands available in the Clients menu.



Figure 1-12 Clients Menu

Server Menu

The Server menu (shown in Figure 1-13) contains commands for setting up, changing, and configuring the NetWorker server.

See Chapter 3, “Configuring and Monitoring the Server,” for a description of the commands available in the Server menu.

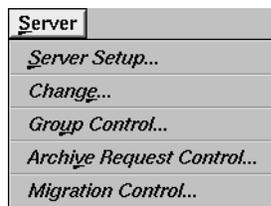


Figure 1-13 Server Menu

Media Menu

The Media menu (shown in Figure 1-14) contains the commands for using and managing the NetWorker server backup volumes. It also includes the commands for configuring and monitoring the devices used for backup and recovery, whether they are standalone backup devices or jukeboxes.

See Chapter 4, “Managing Media and Backup Devices,” for a description of the commands available in the Media menu.



Figure 1-14 Media Menu

Save Set Menu

The Save Set menu (shown in Figure 1-15) contains commands for recovering entire save sets and for *cloning*, or making identical copies of save sets.

See Chapter 6, “Recovering and Cloning Save Sets,” for a description of the commands available in the Save Set menu.



Figure 1-15 Save Set Menu

Customize Menu

The Customize menu (shown in Figure 1-16) contains selections for customizing and configuring your NetWorker backups. The menu contains commands for configuring schedules, policies, directives, label templates, groups, and notifications.

See Chapter 7, “Customizing NetWorker Backups,” for a description of the commands available in the Customize menu.



Figure 1-16 Customize Menu

Options Menu

Use the Options menu (shown in Figure 1-17) to make selections for customizing the NetWorker interface. You can set the polling interval, choose whether or not to display the speedbar, and sound a tone when a pending message arrives. The Options menu includes the Polling Interval, Speedbar, and Tone commands.

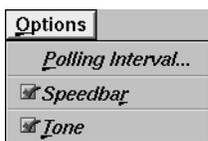


Figure 1-17 Options Menu

Changing the Polling Interval

The Polling Interval menu choice allows you to change how often, in seconds, the data in the NetWorker Administrator window is updated.

To change the polling interval, follow these steps:

1. Select Polling Interval from the Options menu. The Change Polling dialog box appears, as shown in Figure 1-18.



Figure 1-18 Change Polling Dialog Box

2. Use the mouse to slide the bar to the right or left to select the desired polling interval.
3. Click Ok to apply the change.

Displaying the Speedbar

NetWorker provides the convenience of a speedbar (shown in Figure 1-19) for the functions you perform most often. By default, the speedbar is enabled. Use the Speedbar toggle button in the Options menu to display or hide the speedbar.

This manual informs you when a speedbar button is available for accomplishing the same task as one of the commands from a pulldown menu.



Figure 1-19 Speedbar Buttons

Setting the Tone Option

Choose the Tone toggle button in the Options menu if you wish to have a tone sound when a pending NetWorker message arrives and when it clears. The tone sounds once when the message arrives and once after the task requested by the pending message has been completed.

Help Menu

All NetWorker windows contain a Help menu (shown in Figure 1-20) that offers you a variety of choices for help on using NetWorker.

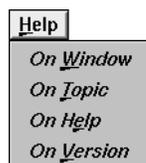


Figure 1-20 Help Menu

The Help menu contains four choices.

- On Window—displays a help screen with information about the currently selected window.

- On Topic—displays a list of topics from which you select a topic.
- On Help—gives you information on how to use the Help facility.
- On Version—displays the version of NetWorker software running on the server.

View Menu

The View menus located in many of the NetWorker Administrator windows contain the Tabular, Details, and Search choices (see Figure 1-21). Use these menu items to view the contents of a window in tabular format or to see more details in a current window.



Figure 1-21 View Menu

Displaying a Tabular View

NetWorker has the ability to display information in a tabular view for most of the NetWorker Administrator windows. This feature is useful for administering NetWorker because it enables you to view information or data in a compact format. The tabular view enables large sites to view a lot of data in a single display.

For example, the Clients window offers two ways to view client configuration information:

- one client at a time
- all of the clients at the same time, in the tabular view

To display the tabular view for the Clients window, follow these steps:

1. Choose Client Setup from the Clients menu. The Clients window appears, as shown in Figure 1-22.

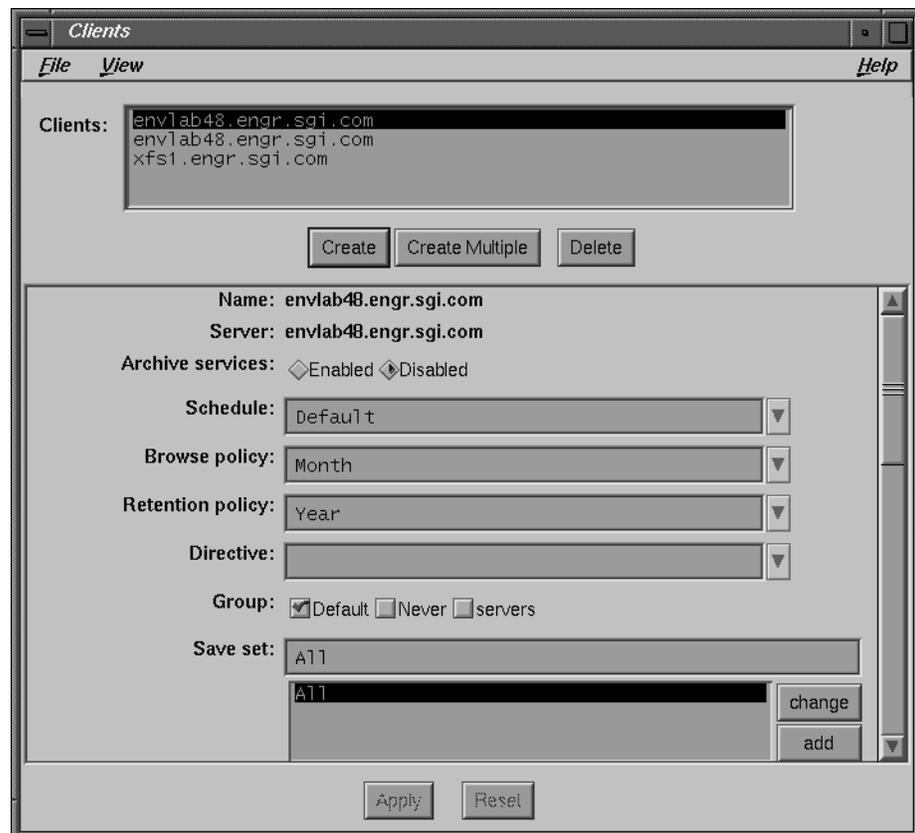


Figure 1-22 Clients Window

2. Choose the Tabular toggle button from the View menu.
The Clients window changes to a tabular view, as shown in Figure 1-23.

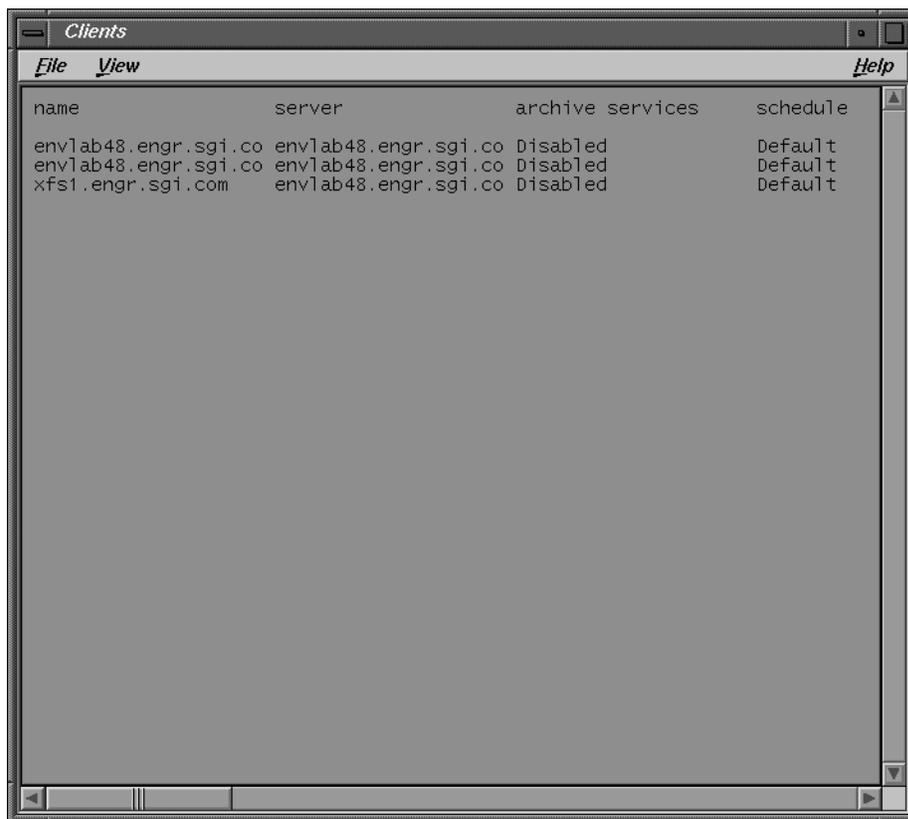


Figure 1-23 Clients Window, Tabular View

3. Use the scroll bars or resize the window to see all of the client configuration information.
4. To return to the normal Clients window, again choose the Tabular toggle button from the View menu.

Tabular view is available for most of the administrator windows. The Save Set Clone and Schedules windows are two exceptions.

Displaying Details

The Details toggle button in the View menu provides you with more detailed information about the current NetWorker window.

Caution: Do not print the contents of a window if you are using the Details command to view information. The printer output will not be readable.

For example, to display the details of the Groups window, follow these steps:

1. Choose Groups from the Customize menu. The Groups window appears, as shown in Figure 1-24.

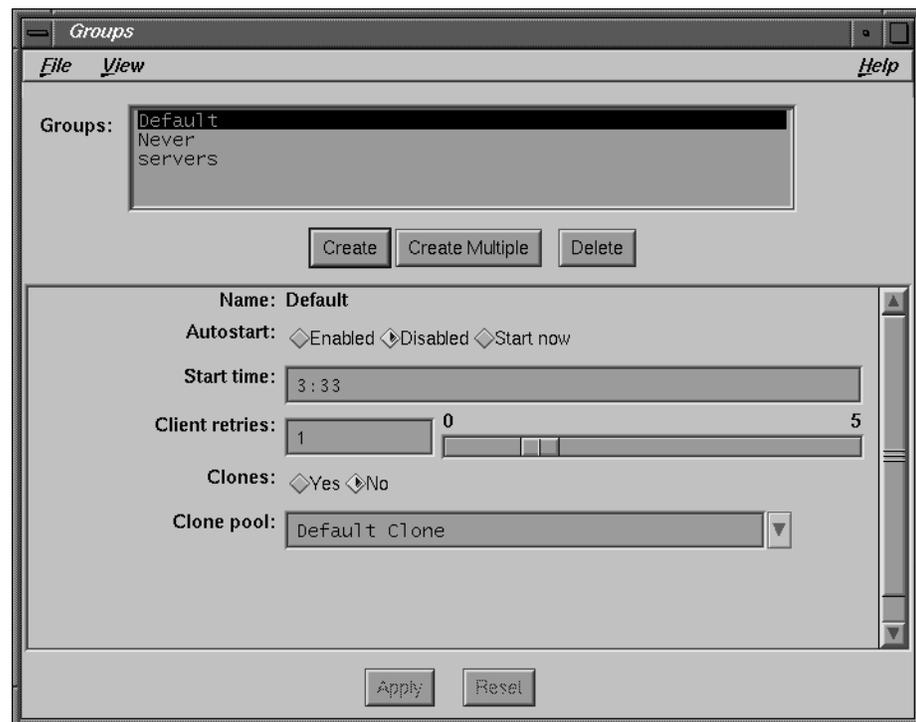


Figure 1-24 Groups Window

2. Choose the Details toggle button from the View menu.

The Groups window changes to show you more detailed information, as shown in Figure 1-25.

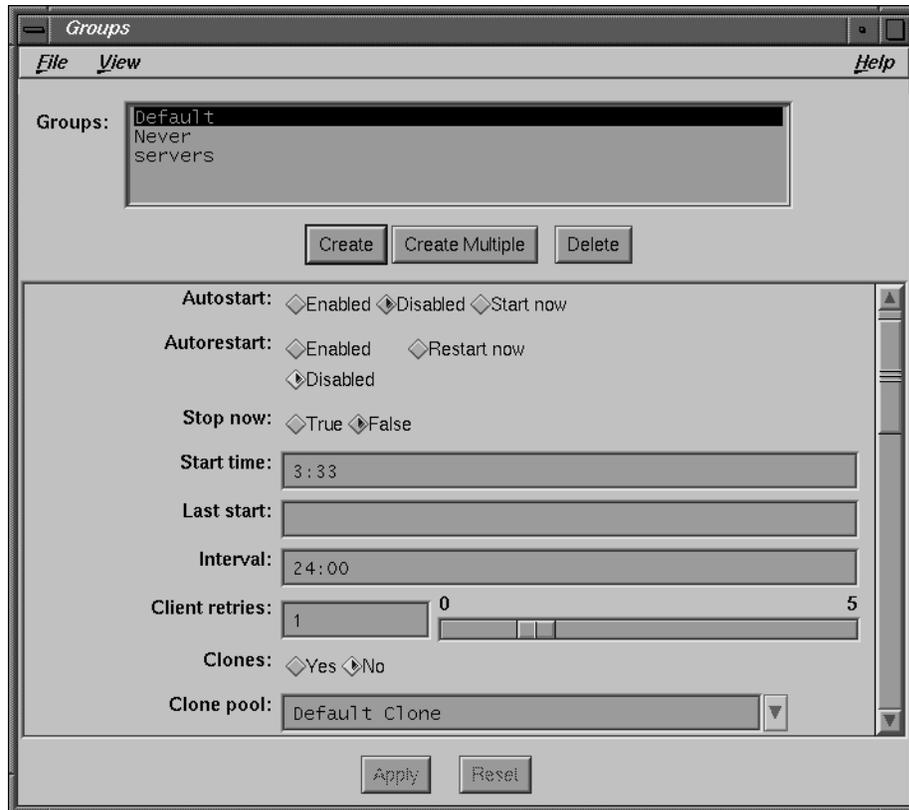


Figure 1-25 Groups Window, Details

The Groups window displays additional fields with more selections and information. To view the entire contents of the window, use your mouse to drag the scroll bar to the bottom of the window.

Searching for Resources

The Search command provides you with an alternate method for locating or finding a NetWorker resource. For example, if you back up a large number of clients, this feature saves you from having to search through the entire scrolling list of client names for a single client or alphabetized groups of clients.

The steps below show how to search for a NetWorker client after opening the Clients window:

1. Choose Search from the View menu. The Search dialog box appears, as shown in Figure 1-26.



Figure 1-26 Search for Resource Dialog Box

2. Enter all or part of the client name you want to locate in the Search for resource field.
Note: The search is not case-sensitive. You can enter the search string without having to recall exactly how a NetWorker resource is entered on your system.
3. Click Ok to begin the search.

If the resource is found, it is highlighted in the scrolling list. If the resource is not found, a notice window appears saying the resource could not be found, as in Figure 1-27.



Figure 1-27 Notice for Resource Not Found

Generic Buttons

NetWorker uses generic buttons in many of its windows. Most are self-explanatory. A brief description of each one follows:

- Cancel button—click to cancel changes or an operation.
- Create button—click to add or create a new resource.
- Ok button—click to close a window or to apply resource settings.
- delete button—click to delete a selected item.
- change button—click to change a selected item.
- add button—click to add a new item.
- Apply button—click to apply your selections to the highlighted resource.
- Reset button—click to cancel changes and display the current settings.

Create Multiple Button

Use the Create Multiple button to create multiple items or resources sharing the same characteristics. Using this button eliminates the need to make repeated, identical selections for NetWorker resources.

For example, to create multiple clients with the same configuration choices, follow these steps:

1. Open the Clients window by choosing Client Setup from the Clients menu.
2. Click the Create Multiple button, shown in Figure 1-28.



Figure 1-28 Create Multiple Button

3. Enter each client name in the Name field, one at a time. Click the add button after entering each name.
4. Select the common configurations that apply to all the new clients.
5. Click the Apply button to apply the configurations to all the new clients.

Print and Save Buttons

Several NetWorker windows contain Print and Save buttons for printing and saving the contents of the window. The following steps demonstrate how to print and save the contents of the Instances window.

1. Open the Indexes window by clicking the Indexes speedbar button or by choosing Indexes from the Clients menu.
2. Select a save set and click the Instances button in the Indexes window to open the Instances dialog box.
3. Click the Print button in the Instances dialog box. The Print Options dialog box appears, as shown in Figure 1-29.



Figure 1-29 Print Options Dialog Box for Hardcopy

4. If you want, change the print command (**lp**) in the Command field.
5. If you want, enter the number of copies in the Copies field.
6. Click Ok to print.

To save the contents of the Indexes window, follow these steps:

1. Open the Indexes window by choosing Indexes from the Clients menu.
2. Click the Instances button in the Indexes window to open the Instances dialog box.
3. Click the Save button in the Instances dialog box. The Save Options dialog box appears, as shown in Figure 1-30.



Figure 1-30 Save Option Dialog Box for Report

The information is saved to a file specified by the environment variable NWREPORT, if set; otherwise, the information is saved to *./networker.rpt*. You can save the data to a file of your choice by entering its name in the File field.

4. Specify an option for NetWorker to follow in case the file already exists. To update or append the existing file, click the Append button. To overwrite the existing file, click the Overwrite the existing file button.
5. Click Ok to save the information from the Instances dialog box to the specified file.

Summary

This chapter introduced you to basic NetWorker concepts, explained three different NetWorker programs, and showed how to navigate the NetWorker GUI.

The second chapter of the *IRIX NetWorker Installation Guide* describes how to test drive NetWorker. Detailed instructions lead you through the tasks required for a first-time backup of your files using NetWorker.

Configuring and Monitoring Clients

The Clients menu provides access to command windows for adding, configuring, and monitoring NetWorker clients and their associated save sets.

Use the Clients window for applying configurations directly to each client. These configuration settings include schedules, browse and retention policies, directives, and groups. The Clients window also provides the flexibility of scheduling large client filesystems individually and the ability to restrict access to client files.

The Indexes window provides information about backed-up files for each client, called save sets. Each time a client file is backed up, NetWorker creates an index entry for the file in the online file index. The Indexes window allows you to monitor and manage the size of the online file index for each client.

Caution: You must have administrator privileges to add a client to NetWorker's backups or change settings in the Clients window.

Configuring NetWorker Clients

NetWorker provides preconfigured settings for backing up clients. These settings also apply to the server, which automatically becomes a client of itself. You may use these preconfigured settings or create your own and apply them to clients in the Clients window. For more information on the preconfigured settings available, refer to Appendix A in the *Installation and Maintenance Guide*.

Navigating the Clients Window

This section describes the items contained in the Clients window and how to use them to configure clients.

Open the Clients window by choosing Client Setup from the Clients menu. The Clients window appears, with the preconfigured settings as shown.

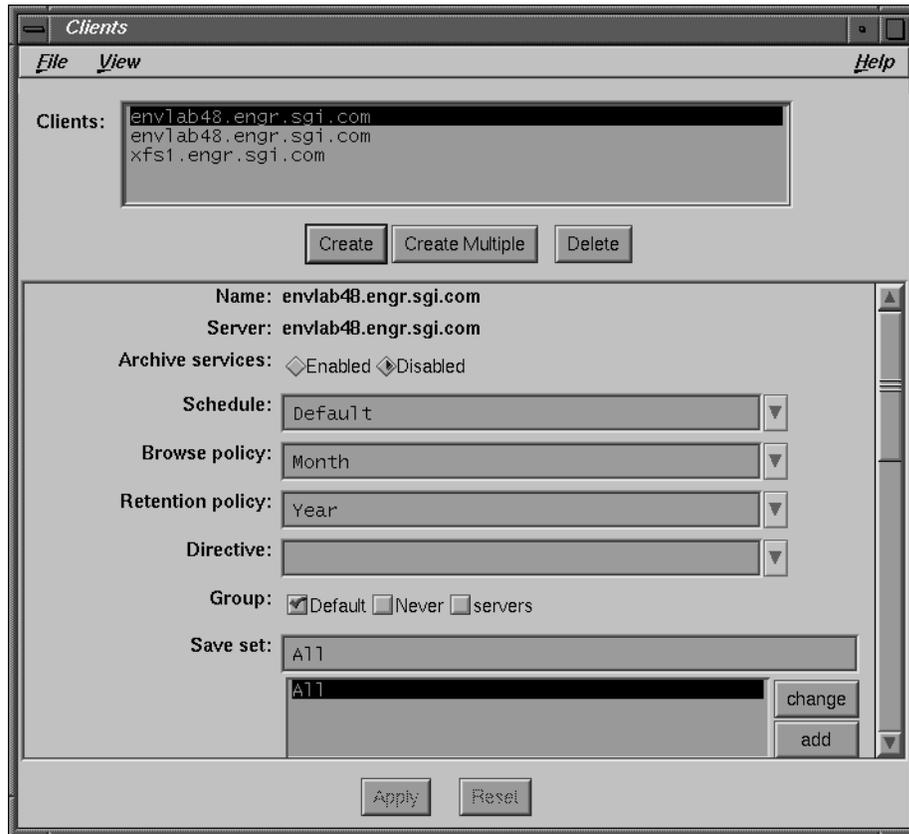


Figure 2-1 Clients Window as Preconfigured

Note: The server becomes a client of itself when you install NetWorker. NetWorker automatically adds the server *hostname* to the Clients scrolling list and provides the preconfigured settings shown above. You must enable the Autostart feature in the Groups window for NetWorker to back up the server automatically.

The following list describes items contained in the Clients window. The preconfigured settings described are automatically applied to the NetWorker server and newly created clients, unless you choose to make new settings.

- Clients scrolling list—displays an alphabetical list of *hostnames* for NetWorker clients of the selected server. Since the NetWorker server is a client of itself, its *hostname* appears in the Clients list when you open the Clients window for that server.

- **Name field**—displays the name of the selected client. To add a client: click the Create button, enter the name of the client in the Name field, select the appropriate settings, and then click Apply.
- **Server field**—displays the name of the NetWorker server currently in use.
- **Archive services choices**—allows enabling or disabling of archive services for clients. This option is available only with purchase of the optional NetWorker Archive Application. Backup clients are automatically enabled for archive services upon installation. Refer to Chapter 8, “Archive Application,” for more information.
- **Schedule field**—displays the assigned schedule for the selected client. Click the arrow button to access the Schedule scrolling list, which displays schedules included with NetWorker and any new schedules you create in the Schedules window.

One of the preconfigured Schedule settings included with NetWorker is Default, which performs a full backup every Sunday and an incremental backup every other day of the week.

- **Browse policy field**—displays the chosen browse policy for the selected client. Click the arrow button to access the Browse policy scrolling list containing the policies included with NetWorker and any new policies you create in the Policies window.

The default browse policy setting is Month. Entries for the client files backed up remain in the file index for one month. Backed-up files can be browsed and recovered in the NetWorker Recover window for one month.

- **Retention policy field**—displays the retention policy for the selected client. The retention policy determines how long the save sets on a volume will be recoverable before the volume is available for recycling. Click the arrow button to access the Retention policy scrolling list containing the policies included with NetWorker and any new policies you create in the Policies window.

The default retention policy setting is Year. The names of the backup volumes are retained for one year in the media index. Files can be recovered from the backup volumes for up to one year, unless the volumes are overwritten.

Note: NetWorker prevents you from setting a browse policy that exceeds a client’s retention policy. For example, a browse policy of one month with a retention policy of one week results in an error message. A volume is usually recycled once its retention has expired, when online file index entries are no longer valid.

- **Directive field**—displays the directives recognized by the NetWorker server. The Directive field is blank, by default. Click the arrow button to access the Directive scrolling list containing the directives included with NetWorker and any new directives you create in the Directives window.

- **Group choices**—displays all backup groups known to the current NetWorker server. To enable scheduled backups for the selected group, you must highlight the group name in the Groups window, then click the Enabled button for the Autostart feature.

NetWorker provides a preconfigured group named Default, which starts a backup daily at 3:33 a.m. Additional backup group settings are displayed in the Clients window once they have been set up within the Groups window. You can assign a client to as many groups as you want.

- **Save set field**—displays individual client save sets (typically comprising one or more client filesystems) for the selected NetWorker client. The Save set scrolling list displays all filesystems entered for a NetWorker client. To back up data from a single client filesystem, enter the filesystem name. The Save set setting included with NetWorker is All, which backs up all local filesystems for the selected client.

Use the scrollbar to access the lower portion of the Clients window, as shown in Figure 2-2.

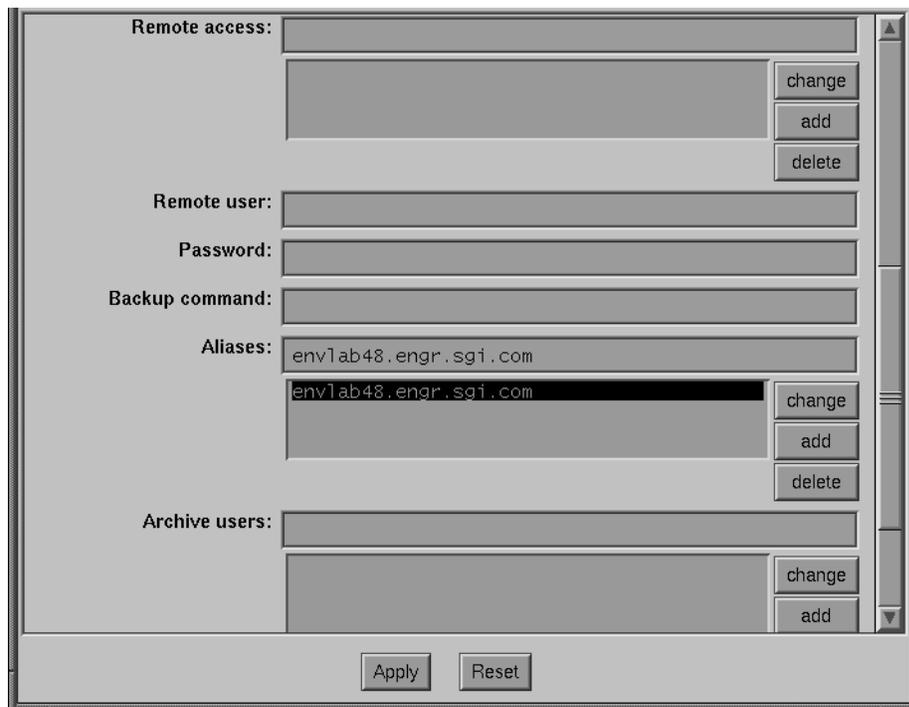


Figure 2-2 Clients Window, Scrolled Down

- **Remote access field**—displays the remote access permission assigned to the selected client. A blank field indicates that only the client itself can recover its backed-up files. If you want specific users or other clients to have access to data for another client, enter their *user@hostname* or *netgroup* name in this field.
- **Remote user field**—use this field to restrict client permissions by entering the client user name. NetWorker uses the *user name* in this field to execute **save** and **savefs** on a client during a scheduled backup. When this field is blank, the *user name*, by default, is *root*.
- **Password field**—use this field only if the client is a PC running NetWare, to include the PC client in a scheduled backup. Enter the password used by the NetWorker server to log into the client system in the Password field and press [Enter].
Note: If you enter a password in this field, you must also enter a user name in the Remote user field.
- **Backup command field**—enter the name of a program file you create that includes commands to execute before and after NetWorker backs up client data. Create the file in the directory where the **save** command is installed. See “Using the Pre- and Post- Processing Backup Command” in this chapter for details.
- **Aliases field**—use this field to enter an unlimited number of aliases for a selected NetWorker client. Client names may be in simple local name format (for example, *myclient*) or fully qualified DNS format (for example, *myclient.acme.com*). List all aliases for the client to help eliminate failed saves due to unrecognized or mismatched client names on the network.
- **Archive users field**—you can only access this field if the optional Archive Application is enabled on your NetWorker server. Refer to Chapter 8, “Archive Application,” for more information.

Using the Pre- and Post- Processing Backup Command

The Backup command field provides a method for customizing client backups. This field allows you to enter the name of a program you create that affects the way client data backs up. For example, you can create a program that shuts down either a mail server or database before a NetWorker backup and then restarts it after the backup completes. The program can be a binary executable or a UNIX shell script.

The backup command must follow these conventions:

- The program name must begin with the prefix *save* or *nsr* and cannot exceed 64 characters.

- The program file must reside in the same directory as the **save** command, typically */usr/bin*.
- The **save** command must be included in the program to ensure data is properly backed up.
- All commands within the program file must successfully execute, otherwise NetWorker cannot complete the remaining instructions.

Use your favorite text editor to create a program file in the directory where the **save** command resides, typically */usr/bin*. You can provide commands to execute either before or after the **save** command, if you do not need both. For example, your program may contain commands in the following order:

- command(s) to execute before a client backup
- command to back up the client data (**save** command)
- command(s) to execute after a client backup

Enter the program name in the Backup command field in the Clients window. When NetWorker begins a backup, it checks for an entry in the Backup command field for the client. If a command is present, NetWorker runs that command instead of the **save** command normally used to back up client data. Since NetWorker does not run the standard **save** command when a backup command is indicated, you must include the **save** command in your program. If you do not include the **save** command in the program, NetWorker will not back up the data.

Example—Backup Command

The following is an example of a valid Backup command program, named *savemsg*. This program prints a message before **save** begins, executes **save** on the client data, then prints a message after **save** completes.

```
#!/bin/sh
case $0 in
/* )   PATH=/bin:/usr/bin:~/bin/dirname $0`
      c=~bin/basename $0`
      ;;
* )   PATH=/bin:/usr/bin:/usr/sbin
      c=$0
      ;;
esac
export PATH
```

```
# print message at start of backup
echo "$c": backup started at `date`

# perform backup on client
save "$@" > /tmp/saveout$$ 2>&1

# print message at completion of backup
echo "$c": backup complete at `date`
# print save output last for savegrp
cat /tmp/saveout$$
rm -f /tmp/saveout$$
exit 0
```

Silicon Graphics strongly recommends that you immediately try backing up a client once you have created your program. This ensures that the backup command created does not prevent the client from backing up successfully during unattended backup sessions. Click the Start button in the Group Control window to start an immediate backup.

Configuring Clients

Before a client can back up to a NetWorker server, the client software must be installed. The software can be installed locally on the client disk or NFS[®] mounted over the network. After the software is successfully installed on the client, add the client to your NetWorker backups in the Clients window.

See the *IRIX NetWorker Installation Guide* for installation instructions.

NetWorker clients can manually back up and recover files from a NetWorker server and browse the online file index entries for files. Clients normally back up during a regularly scheduled network-wide backup.

To configure a client, make a selection for each of the following items:

- NetWorker server, if you have more than one
- backup schedule
- browse and retention policies
- directive
- backup group (one or many)
- individual client filesystems (rather than all filesystems for the client)

Adding a New Client

To add and configure a new NetWorker client, follow these steps:

1. Open the Clients window.
2. Click the Create button.
3. Enter the new client *hostname* in the Name field.
4. Add the client to one or more backup groups by selecting the group(s) you want from the Group choices.
5. Select a schedule from the Schedule choices.
6. Select both a browse and a retention policy for the client indexes. You may use the existing policies of Month for the browse policy and Year for the retention policy.
7. Select a set of directives for the client.
8. Add client aliases, if any, in the Aliases field.
9. Click Apply.

Note: For details on how to add groups of clients who share common settings, see the instructions in “Create Multiple Button” on page 20 of this guide.

If you try to add clients beyond the number of connections you purchased for your NetWorker server, you receive the error message “Too many clients—maximum is *number*.”

Next, choose the files you want NetWorker to back up for the client.

Tip: If you want all files for the client backed up, enter All in the Save set field and skip the instructions for backing up specific filesystems.

Follow these steps to back up a specific filesystem:

1. Select All in the Save set scrolling list.
2. Click the delete button.
3. Enter the *pathname* into the Save set field for the filesystem that you want backed up. For an example see Figure 2-3.

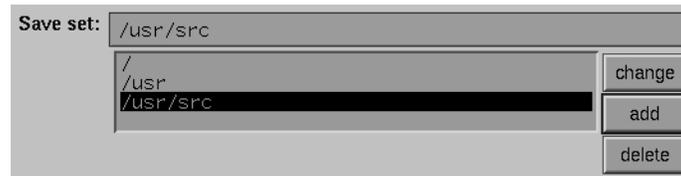


Figure 2-3 Save Set in Clients Window

4. Click the add button to add the client filesystem. The *pathname* you typed appears in the Save set scrolling list.

This feature allows you to back up scheduled filesystems at different times. For example, if a client has a large amount of data, you may want to schedule its filesystems separately for backups. For details, see “Scheduling Large Client Filesystems” on page 32.

By default, NetWorker clients can browse or recover only their own files. If your organization is concerned about security, Silicon Graphics recommends that you not change the remote access for any clients.

To give other machines access to the files of another client, follow these steps:

1. Enter the *hostname* of the client machine (or *netgroup* name, if you are using NIS) in the Remote access field, as shown in Figure 2-4.

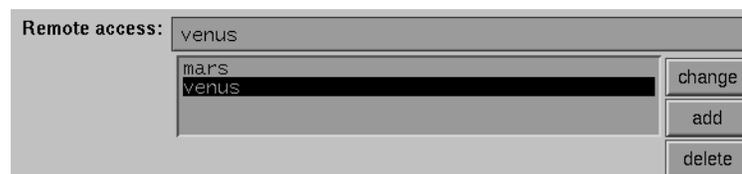


Figure 2-4 Remote Access in Clients Window

2. Click the add button next to the Remote access scrolling list.
The *hostnames* or *netgroup* names added to the Remote access field appear in the Remote access scrolling list for the client.
3. Decide whether to use the Remote user and Password fields. See “Navigating the Clients Window” on page 23 for details.
4. Enter all aliases for the client in the Aliases field, including the client’s fully qualified domain name (for example, *venus.company.com*).

Caution: If you do not enter all aliases for a client, you increase the risk of a failed backup because of unrecognized or mismatched client names on your network.

5. Click the Apply button to apply the settings and to add the new client to the NetWorker server. The new client appears in the Clients scrolling list. If you make a mistake, click the Reset button to display your previous set of choices.

Scheduling Large Client Filesystems

Using a moderate backup rate of 300 kilobytes per second, a full backup for a client with 5 gigabytes of data would take about 5 hours to complete. Consequently, it may not be convenient to complete a scheduled, unattended, full backup for this client because of the amount of time it will take.

You can, however, schedule the client filesystems to be backed up at different times. Splitting filesystem backups for a client enables you to save all files for that client without trying to do a time-consuming full backup of all of the local filesystems at one time.

The Save set scrolling list in the Clients window displays the filesystems that NetWorker backs up for a client. All of the local filesystems for a client are backed up at the same time, as long as the keyword All remains in the list.

To back up filesystems individually, add and configure the same client twice in the Clients window. Configure the first client to back up half its local filesystems with one backup schedule in one group. Configure the duplicate client to back up the other half of its local filesystems with a second backup schedule in another group.

The following example illustrates how to divide up the client filesystems so they back up at different times.

1. Find the names and sizes of the local client filesystems using the **df** command.

```
% df -k
Filesystem      Type  kbytes  use    avail  %use  Mounted on
/dev/root       efs   966304  619623 346681  65    /
/dev/dsk/dks0d2s7 efs  1966304 1573043 393261  80    /home
/dev/dsk/dks0d3s7 efs   966304  434836 531468  45    /nsr
/dev/dsk/dks0d4s7 efs  1210250  822970 387280  68    /export
/dev/dsk/dks0d5s7 efs  1210250  919790 290460  76    /var
/dev/dsk/dks0d6s7 efs  1210250 1077122 133128  89    /home/atlas
```

The largest filesystem in this example is */home*, with 1,966,304 KB of space.

2. Estimate how long it will take to back up the largest filesystem, using the following example:

```
1966304 kbytes ÷ 300 kbytes/second = 6554.34 seconds
6554.34 seconds ÷ 60 seconds/minute = 109.23 minutes
```

Repeat the same exercise for each filesystem to determine how to divide them into different backup schedules.

3. Choose Schedules from the Customize menu and create two new backup schedules:
 - "Monday Full," which performs a full backup on Mondays and incremental backups the rest of the week.
 - "Wednesday Full," which performs a full backup on Wednesdays and incremental backups the rest of the week.

Each full backup in our example should take less than three hours to complete.

Use the Clients window to configure the client:

1. Open the Clients window.
2. Select the client from the Clients scrolling list whose filesystems you plan to divide into different backup schedules. If the client has not already been added to the scrolling list, click Create to add the new client.
3. Configure the client by choosing a backup group, a directive, a policy, and the Monday Full schedule.
4. Delete All from the Save set list and add each filesystem you want to back up with the Monday Full schedule, one at a time.

Apply the Monday Full schedule to the following filesystems:

```
/dev/root          efs  966304  619623  346681  65  /
/dev/dsk/dks0d2s7  efs  1966304  1573043  393261  80  /home
/dev/dsk/dks0d3s7  efs  966304  434836  531468  45  /nsr
```

5. Click the Apply button.

Follow the same procedure for the remaining filesystems you want to back up with the Wednesday Full schedule.

1. Click the Create button to add the duplicate client.
2. Configure the client by choosing a backup group, directive, policy, and the Wednesday Full schedule.

3. Delete All from the Save set field, and enter each filesystem you want to back up on the Wednesday Full schedule. Apply the Wednesday Full schedule to the three other filesystems, as shown in Figure 2-5:

```
/dev/dsk/dks0d4s7  efs  1210250   822970   387280  68  /export
/dev/dsk/dks0d5s7  efs  1210250   919790   290460  76  /var
/dev/dsk/dks0d6s7  efs  1210250  1077122  133128  89  /home/atlas
```



Figure 2-5 Filesystems as Client Save Sets

4. Click the Apply button in the Clients window.

NetWorker backs up the filesystems on the client disk using two separate schedules, making it possible to run unattended full backups on this client.

Caution: When you schedule the filesystems, make sure you do not omit any of them from the Save set scrolling list. Any filesystem left off the list will not be backed up.

Removing a Client

To remove a client from the Clients window, follow these steps:

1. Select the name of the client in the Clients scrolling list.
2. Click the Delete button. NetWorker asks for confirmation to delete the client.

Deleting a client means it cannot back up or recover its files from the NetWorker server. The backup history for the client still remains in the file and media indexes until you specifically remove the entries. See “Removing a Backup Volume” on page 77 for instructions on how to remove the index entries for a NetWorker client.

Securing the NetWorker Environment

When NetWorker runs the **save**, **savefs**, and **recover** programs on a client, the programs attempt to establish a connection with the server. Before accepting the connection, the NetWorker server verifies that the user initiating the programs has the necessary permissions. The programs must meet the following conditions for the server to accept the connection:

- The request for the connection must be made from a secure port on the machine requesting the connection. A secure port can be opened only by *root*, so the **save**, **savefs**, and **recover** programs run **setuid** to *root*.
- The server verifies that the user executing the programs at the machine has permission to save or recover client files. One of the following criteria must be met for a user to have permission:
 - The machine name for the user is equivalent to the client name.
 - The user attempting to establish a connection must be a member of the Remote access list in the Clients window.

Once a connection has been established, the client programs **save** and **recover** set their effective user ID (UID) to the UID of the user who initiated the program. This ensures that all local filesystem and system call access is done as that user, preventing other users from recovering or backing up files to which they do not have access. The exception to this rule is that the user name *operator*, and users in the group *operator*, have filesystem access privileges of *root*. This allows the administrator to set up a login or group for the operators who initiate backups and recovers on behalf of other users, without giving the operators *root* access to client machines. For more information about security, refer to the nsr(1M) reference page.

Restricting Client Access

NetWorker's preconfigured settings allow clients to browse and recover only their own files. To give other clients recover access to files on another client, the administrator must explicitly add the access to the Remote access list in the Clients window. See "Configuring Clients" in this chapter for more information.

You can further tighten the access control for client programs by turning off the *set-uid* bit. This restricts the use of the **save**, **savefs**, and **recover** programs on client machines to *root*. To allow access by *root* and *operator*, but not by other users, change the group

ownership of these programs to *operator* and set the mode bits to allow execution by owner and group.

The **savegrp** command initiates the **savefs** command on each client machine in a backup group by sending a remote command request to the **nsrexecd** program.

The **nsrexecd** program runs on NetWorker client machines. This program provides a secure and restrictive way for NetWorker to start automatic backups on clients. The **nsrexecd** program allows you to restrict access to a select set of NetWorker servers. When you install NetWorker on a client, **chkconfig** automatically turns NetWorker on, so **nsrexecd** will be started each time the client reboots.

Migration Setup

The NetWorker Hierarchical Storage Management Application (HSM) provides file migration and recall services to a range of client machines. If you have purchased this optional extension for your Legato NetWorker server, the Migration Setup choice on the Client menu is enabled.

For complete instructions on the use of the NetWorker HSM Application, see the chapter on Hierarchical Storage Management in the Legato documentation.

Manually Managing the Online Indexes

This section describes NetWorker's powerful index management features for manually managing the online indexes. You can also automatically manage your indexes by using the different index policies in the Policies window.

Note: See "Creating a New Policy" on page 145 for more information on automatic index management.

Every time a backup completes, NetWorker creates entries for backed-up files for each client in the online file indexes. The indexes require disk space and must be monitored to ensure they do not become too large.

The remainder of this chapter provides a description of the following elements of online index management:

- online file and media indexes

- save sets
- four online index actions

The remainder of this chapter also includes instructions on manually managing the indexes and backup volumes:

- how to generate reports about the contents of the file index
- how to remove the oldest cycle and reclaim index space from the file index
- how to manage backup volumes by changing the mode of a volume or by removing a volume from the media index

Online Indexes

NetWorker maintains two types of indexes: a *file index* and a *media index*. The file index stores information about the files backed up by NetWorker. The media index stores information about NetWorker media and the save sets stored on the media. NetWorker uses these online indexes to locate files requested for recovery. NetWorker determines which volume to mount for recovering a file by mapping the saved files to their backup volumes.

Each entry in the file index typically includes the following information for a backed-up file: filename, number of blocks, access permissions, number of links, owner, group, size, last modified time, and backup time. The file index grows with each backup, as entries are added for the newly backed-up files. As long as an index entry for a file remains in the file index, you may recover the file by using the NetWorker Recover window.

The media index is usually much smaller than the file index because each volume contains many saved files. The size of an index is proportional to the number of save set entries it contains.

Save Sets

Save sets are groups of files, usually contained in a single filesystem, that have been backed up by NetWorker. Save sets are created each time a backup is started. Generating a save set creates one or more entries in both the file and media indexes.

To conserve both index and backup volume space, NetWorker provides a variety of backup levels. For example, a *full* backup contains all files in a given filesystem, while

an *incremental* backup contains only the files that have changed since the previous backup. Since directories often contain files that do not change along with files that do change, incremental backups conserve both index and volume space.

For information on different backup levels, see the section “Understanding Backup Schedules” on page 140.

You need files from both the incremental and full backups to recover a complete directory. Without the underlying full backup, you cannot completely recover the directory. Without the incremental backups, you can recover only the last full backup of the directory, which would not include any incremental changes since the full backup occurred. The incremental backup depends on the full backup. Since both the full and incremental backups are needed to recover the complete directory, NetWorker checks these dependencies when removing save sets from the indexes, either manually or automatically.

Figure 2-6 shows the relationship between incremental and full backups.

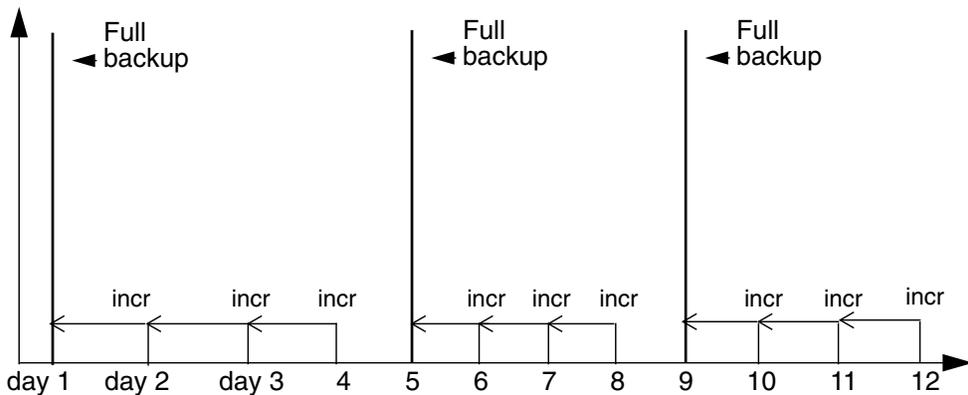


Figure 2-6 Full Versus Incremental Backups

Note: Incremental backups depend upon the previous full backup. NetWorker does not remove a save set until *all* of its dependent save sets have been removed.

Index Actions

There are four actions NetWorker performs on an index: inserting entries, browsing file and media entries, removing entries, and reclaiming space.

- *Inserting entries* in the file index occurs during a backup. If the index has no free space, NetWorker acquires more space from the filesystem to hold the new entries.
- *Browsing* neither increases nor decreases the size of an index. When you browse an index, you simply look through the index for information about your saved files or the contents of your backup volumes.

Browsing the media index occurs in the Volumes window where you view the save sets on the backup volume.

Browsing the file index occurs when you use the NetWorker Recover window to locate a file. You may also browse save sets that contain the files you see in the NetWorker Recover window by using the Indexes window.

- *Removing entries* frees up space in the index. NetWorker uses the free space to insert new entries. The Browse and Retention policies determine when entries get removed from the index automatically. You may also remove them manually by selecting Remove oldest cycle in the Indexes window or by choosing Remove from the Volume menu in the Volumes window.
- *Reclaiming space* returns empty space to the filesystem. Empty space is created when entries are deleted from the index, and removed when you click the Reclaim space button in the Indexes window.

Note: Remember to reclaim space after removing cycles.

Generating File Index Reports

NetWorker provides the ability to generate a report detailing the contents of the file index. The report shows current file index information for the backup *namespace* of all clients of the server. You can use command-line options to see details about file index contents for a specific client, backup time, or namespace. You can choose to view the report on your screen, print it out, or redirect it to a file.

See the `nsrinfo(1M)` and `mminfo(1M)` reference pages for more details on command line options used to customize the file index report.

The following example generates a report on the file index contents for */var* backed up during the past week for a specific client. The output of the report is redirected to a file.

1. Enter the **mminfo** command at the system prompt to list backup times, in GMT (Greenwich Mean Time), for the specific client save set during the last week:

```
% mminfo -r nsavetime -N save_set -t 'one week ago' -c client_name
```

The backup times returned by the *nsavetime* option are in GMT seconds format. Choose the time you want to use for the **nsrinfo** *time* entry.

2. Enter the **nsrinfo** command at the system prompt:

```
% nsrinfo -vV -t time client_name > report_file
```

The report appears similar to the example below, showing the following information for each file found: file type, path, file ID, file size, offset, file length, and namespace.

```
scanning client 'cl' for savetime 811756903 (Fri Sep 22 01:01:43 1995)
from the backup namespace
UNIX ASDF file `~/var/adm/wtmp', fid = 5383.2066, file size=105408,
off=6105488, len=105672, app=backup(1)
1 file found
```

Note: Queries for a specific time can take a while to complete because of the file index structure.

Navigating the Indexes Window

This section describes the Indexes window, its contents, and how to use the window to manage the online file index.

To open the Indexes window, click the Indexes speedbar button (see Figure 2-7) or choose Indexes from the Clients menu.



Figure 2-7 Indexes Button

The Indexes window appears, as shown in Figure 2-8.

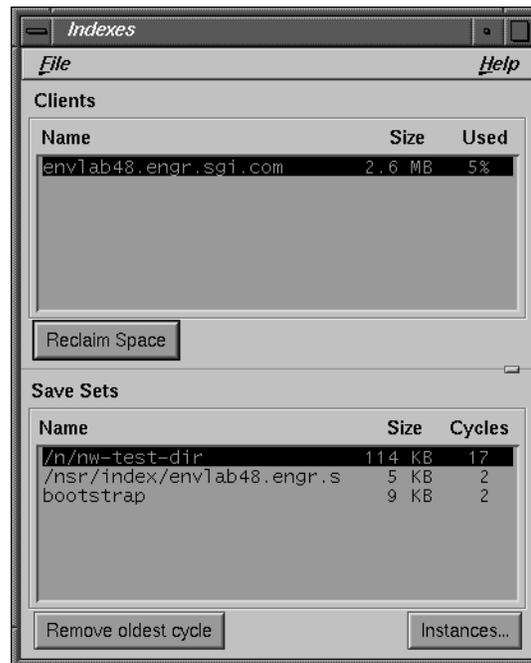


Figure 2-8 Indexes Window

The Indexes window contains the following items:

- Clients scrolling list—contains client names and client file index information:
 - Name—indicates the NetWorker client name.
 - Size—the allocated size of the client file index. The allocated disk space automatically grows as the index size increases.
 - Used—the percentage of the index file in use. If listed as 100%, the index has completely filled the allocated disk space and there is very little, if any, disk space to reclaim. The smaller the percentage, the more disk space there is to reclaim because the index is not using all of the currently allocated disk space.
 - If the percentage is less than 100%, you can reduce the size of the index by first selecting the client and then clicking the Reclaim space button. This removes any holes in the file index created by removing index entries. NetWorker automatically removes index entries based on the browse and retention policies you selected for managing online indexes.

- **Reclaim space button**—click to decrease the size of the client file index. NetWorker compresses the space in the index left by removed entries, thus freeing up disk space. A confirmation box appears when a reclaim space operation begins. Click OK to proceed or Cancel to discontinue.
- **The Save Sets scrolling list**—displays the save sets contained in the file index for the selected client. The save sets are grouped according to their unique save set names. For example, all backups of */usr* in a client file index are grouped under the save set named */usr*.
 - **Name**—shows client save set names.
 - **Size**—displays an estimate of the amount of index space used by the save set group in the client file index.
 - **Cycles**—displays the number of cycles for the save set group contained in the client file index. One cycle starts with a full backup and ends with the next full backup. A cycle includes the incremental and level 1-9 backups, if any exist. See “Selecting the Backup Levels” on page 141 for a description of NetWorker’s backup levels.
- **Remove oldest cycle button**—used to remove the oldest cycle for the save set, if the index requires further reduction. Select the save set, then click the **Remove oldest cycle** button. After removing a save set, click **Reclaim space** to remove the holes left by removing the oldest cycle. A confirmation box appears once the removal of an oldest cycle has begun.
- **Instances button**—used to display entries contained in the file index for the selected save set. Instances are displayed in order, starting with the oldest entry. The **Instances** dialog box displays the save set ID, number of files, size, date backed up, and backup level for each occurrence of the selected save set. Click the **Print** button to print a copy of the information displayed. Click the **Save** button to save a copy of the current save set information to a separate disk file.

Removing the Oldest Cycle

This section provides instructions for removing the oldest full-to-full cycle of a group of save sets from the client file index.

To remove the oldest cycle from the a client index:

1. Open the **Indexes** window by choosing **Indexes** from the **Clients** menu.
2. Select the client from the **Clients** scrolling list.

3. Select the save set group from the Save Sets scrolling list.
4. Click the Remove oldest cycle button to remove the oldest cycle for the save set group. A confirmation box appears, prompting you to confirm that you want to remove the oldest cycle of the selected save set.
5. Click Ok in the confirmation box to proceed or click Cancel to discontinue the operation.

Caution: If you do not select a save set in the Save sets scrolling list, NetWorker removes *all* the oldest cycles of *all* the save sets displayed in the scrolling list for the client selected in the Clients scrolling list.

Removing the oldest cycle opens up index space so that other entries can fit into the empty space without increasing the size of the index.

Figure 2-9 illustrates what happens when you remove the oldest cycle from the index.

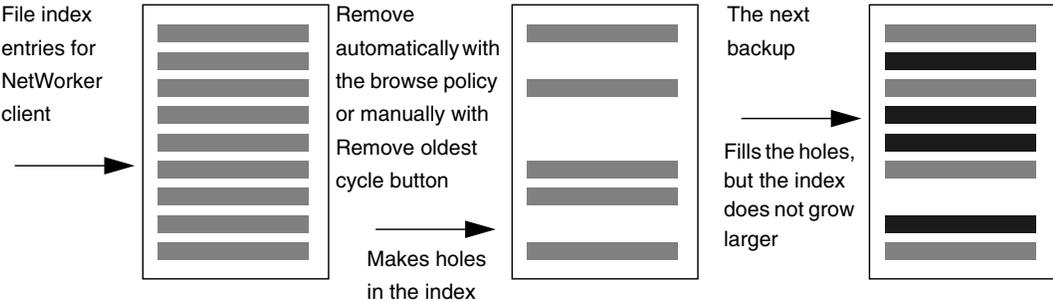


Figure 2-9 Removing Oldest Cycle From the Index

Note that removing index entries does not decrease the *size* of the file index, so it still takes up as much space as it did before. To decrease the amount of space the index uses, use the Reclaim space button.

Note: The last full cycle cannot be removed using the Remove oldest cycle button because it might be needed for recovery.

Reclaiming Index Space

Use the instructions in this section to reclaim empty space in the file index created when entries are automatically or manually removed.

1. Open the Indexes window by choosing Indexes from the Clients menu.
2. Click the Reclaim Space button. You do not need to select individual clients or save set groups from the scrolling lists.

The Reclaim Space dialog box appears, as shown in Figure 2-10.



Figure 2-10 Reclaim Space Window

3. Click the Ok button, to start the reclaim space operation.
The Reclaim Space dialog box changes. The Progress bar moves to indicate the progress of the index operation. When an index is busy, the Ok button is disabled.
4. Click the Cancel button to dismiss the dialog box. The reclaim space operation continues in the background. You can check the progress of the operation at any time by clicking the Reclaim space button to redisplay the confirmation box.

The index is rewritten without the deleted entries, shrinking its size and opening up space for new entries. The reclaimed space is returned to the filesystem.

Figure 2-11 illustrates what happens when you use the Reclaim Space button.

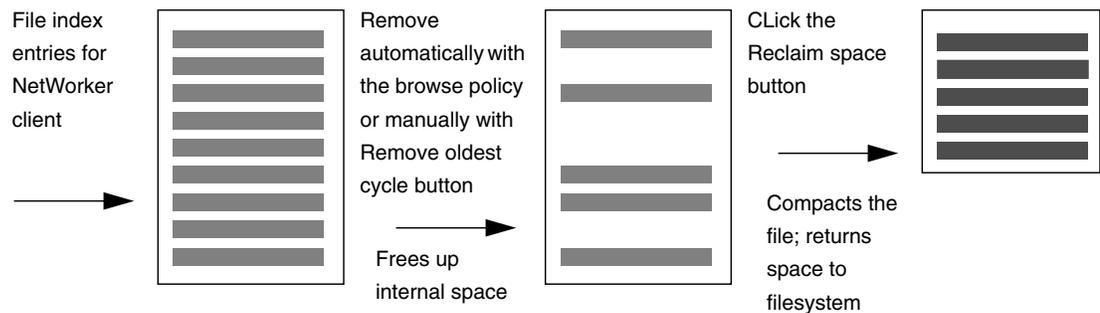


Figure 2-11 Reclaiming Space From the Index

You may not simultaneously reclaim space and remove an oldest cycle for one client. The index is busy during both operations.

After either the reclaim space or remove oldest cycle operations have finished, the statistics in the Indexes window are updated to reflect the current state of the file index.

Viewing Save Set Details

NetWorker provides you with more detailed information about save sets in the Instances dialog box.

The Instances dialog box contains information on the number of files in a save set, its size, the date it was backed up, and the level of the save set: full, incr (incremental), or level.

You can use the information in this dialog box to determine how your resources are being used. For example, you may need to see how large a save set is so you can plan the amount of disk space you need for the online indexes. Or you may need to generate a report that details your NetWorker client backups to fulfill administrative requirements. Follow these steps to view the Instances dialog box:

1. Select a client from the Clients scrolling list in the Indexes window.
2. Select a client save set from the Save Sets scrolling list.
3. Click the Instances button.

The Instances window appears, as shown in Figure 2-12.

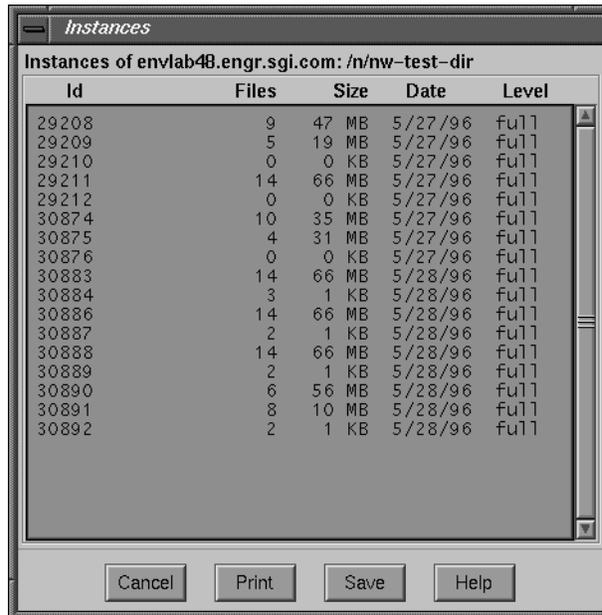


Figure 2-12 Instances Window

The title bar for the Instances window displays the name of the currently selected save set. The Instances scrolling list contains the following information:

- Id—internal NetWorker identification number for the save set.
- Files—number of files in the save set.
- Size—size of the save set.
- Date—date the save set was backed up.
- Level—level of backup: full, incr, or level.

Configuring and Monitoring the Server

Your NetWorker server is the source for creating clients, monitoring the progress of your backups, managing backed up data, and controlling your media and backup devices.

This chapter describes how to use the commands in the **Server** menu for configuring and controlling some of these NetWorker server operations.

In this chapter you will learn how to do the following tasks:

- Change server parallelism.
- Use devices concurrently.
- Add or change a NetWorker administrator.
- Change to another NetWorker server.
- Monitor backups in the Group Control window.
- Immediately start, stop, and restart a backup.
- Register NetWorker.

Configuring the NetWorker Server

The options available for configuring your NetWorker server depend upon which add-on modules you purchased for accomplishing your network-wide backups.

Setting up your server requires a minimal amount of configuring after you first install and run NetWorker. NetWorker provides preconfigured settings that enable you to immediately back up your server and clients to a backup device connected to the server.

You also have a high degree of control over the performance and security of your backups with the configuration choices available in the **Server** window.

Navigating the Server Window

Use the Server window to accomplish the following tasks:

- Change server parallelism.
- Use devices concurrently.
- Add or change NetWorker administrators.
- Provide important company and product information.

To open the Server window, choose Server Setup from the Server menu. The Server window appears, as shown in Figure 3-1.

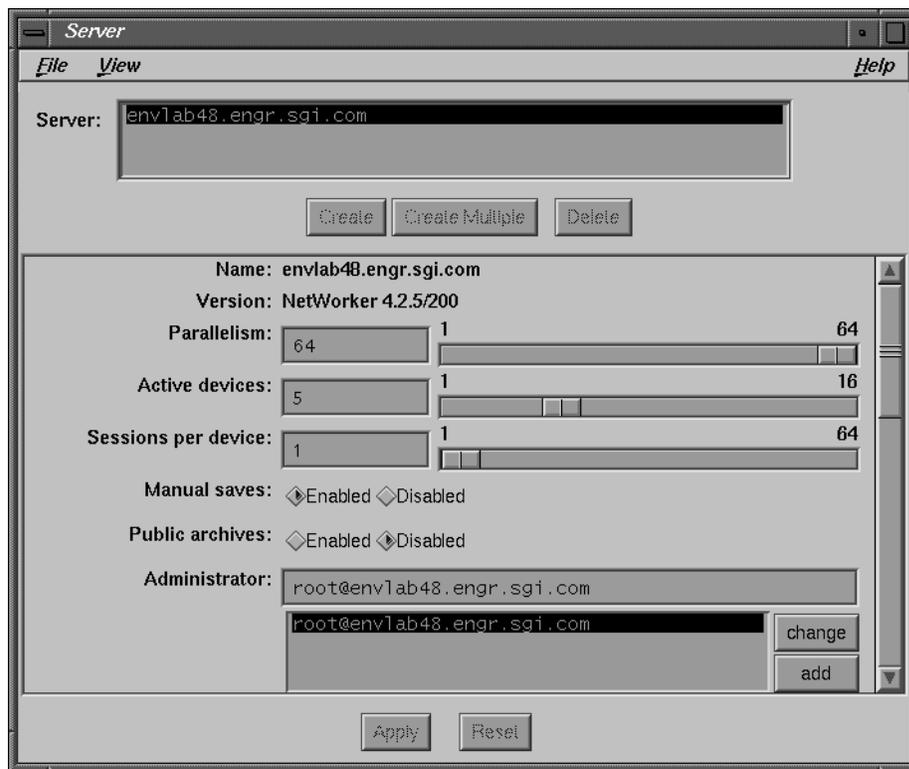


Figure 3-1 Server Window

Use the scroll bar on the right side of the window, if necessary, to view the entire contents of the Server window.

The Server window contains the following items.

- **Name field**—displays the name of the current NetWorker server.
If you have more than one NetWorker server on your network, you can use the Change Server speedbar button or command to change to a different server.
- **Version field**—displays the version of the NetWorker software currently installed on the NetWorker server.
- **Parallelism field**—displays the number of clients backing up in parallel. Use the mouse to slide the bar to the left or right to set the parallelism, or enter the number in the Parallelism field. The maximum value for parallelism varies, depending upon which NetWorker product you purchased.
- **Active devices field**—contains the maximum number of media devices NetWorker can use concurrently. An active device is one that the NetWorker server uses for backups and recovers. Use the slide bar to set the value, or enter the number in the Active devices field.

The Active devices and Sessions per device fields are not available with some versions of NetWorker. See the section “Using Multiple Devices Concurrently” on page 51 for more information.

- **Sessions per device field**—displays the target number of save sessions accepted by an active device. Use the slide bar to set the value, or enter the number in the Sessions per device field.

Tip: You can configure your backups so that one device is reserved for recovers only. For example, if Parallelism is 12 and Active devices is 3, you can set the Sessions per device to 6. This means that two of the devices handle all of the incoming save sets, while the third device remains available for performing a recovery.

- **Manual saves choices**—allow you to enable and disable manual saves performed by clients on the network. When the Disabled toggle button is selected, you prevent all NetWorker clients from performing manual backups. This provides more control over the backup devices and location of the backed-up data. When the Enabled toggle button is selected, any NetWorker client on the selected server can perform a manual backup.
- **Administrator scrolling list**—displays the names of users or user groups who have permission to change the configuration of the NetWorker server. To add a user name, enter the name in the Administrator field, then click the add button.

The remainder of the window contains fields for entering customer information required for registering NetWorker. Refer to Chapter 5 in the *Installation and Maintenance Guide* for more information about registering NetWorker.

Changing Parallelism

A NetWorker server can back up files from several clients in parallel and multiplex the files onto the same backup volume. This feature enables the backup devices to operate efficiently by keeping a steady stream of files supplied to the server. Use the Parallelism field to optimize the efficiency of your server by controlling the number of clients that back up their files at any one time. The maximum value for parallelism varies with the version of NetWorker.

You may want to adjust parallelism (the number of clients backing up simultaneously) under the following conditions:

- If the server takes too long to back up all of its clients, increase the parallelism.
- If the server is so busy backing up clients that it is unavailable for other tasks, decrease the parallelism.
- If the server runs out of swap space or memory, decrease parallelism.

To change the number of clients that can simultaneously back up their files to the NetWorker server, follow these steps:

1. Open the Server window by choosing Server Setup from the Server menu.
2. Select the Parallelism sliding bar and move it to the right or left to display the value you want, or enter the value in the field.
3. Click the Apply button.

The number of clients allowed to back up at one time cannot exceed the parallelism setting. For example, if you have eight NetWorker clients and set Parallelism to **4**, four clients back up at once. The fifth client will begin its backup after completion of the first client backup.

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

Using Multiple Devices Concurrently

In addition to setting the number of clients backing up in parallel to a NetWorker server, you can configure the server to back up to several devices concurrently. The devices operate simultaneously to back up data, thus speeding backups and recovers.

Enter the maximum number of active devices and the number of sessions per device in the Parallelism field (see Figure 3-2). A save session is generated when a client starts a backup. A single client can generate several save sessions simultaneously.

Parallelism:	<input type="text" value="4"/>	<input type="text" value="1"/>	<input type="text" value="16"/>
Active devices:	<input type="text" value="2"/>	<input type="text" value="1"/>	<input type="text" value="10"/>
Sessions per device:	<input type="text" value="4"/>	<input type="text" value="1"/>	<input type="text" value="16"/>

Figure 3-2 Server Parallelism, Active Devices, Sessions per Device

The Active devices value sets the maximum number of backup devices NetWorker uses for concurrent backups. An active device is used for backups or recovers.

The Sessions per device value is the target number of save sessions accepted by an active device. If a device is receiving the maximum number of save sessions, then the server sends the next session to the under-utilized backup device next in line. If all devices are receiving the target number of sessions, NetWorker overrides the Sessions per device value and forces the backup onto the device with the least activity.

Tip: The Parallelism, Active devices, and Sessions per device values control the backup traffic from the NetWorker server to the media devices.

To configure NetWorker to back up to multiple devices concurrently, follow these steps:

1. Open the Server window by choosing Server Setup from the Server menu.
2. Set the Parallelism value by moving the slider bar to display the value you want.
3. Set the Active devices by moving the slider bar to display the value you want.
4. Set the Sessions per device by moving the slider bar to display the value you want.
5. Click the Apply button.

Note: To distribute your client backups evenly across the available backup devices, simply divide the number entered in the Parallelism field by the number in the Active devices field. Then enter that number into the Sessions per device field. For example, if you have **8** entered in the Parallelism field and **2** in the Active devices field, enter **4** in the Sessions per device field.

Adding or Changing Administrators

NetWorker is shipped with preconfigured settings that give only *root@server_name* on the server permission to change system configurations. You can add, delete, or change permissions to give system administration privileges to other users or user groups. When using NetWorker for the first time, run it as *root* to add user group or NIS (network information service) *netgroup* names to the Administrator list.

To make additions or changes to the Administrator list, follow these steps:

1. Open the Server window by choosing Server Setup from the Server menu.
2. Highlight the current entry in the Administrator field, as shown in Figure 3-3.
3. Enter a new user group or *netgroup* name or make changes to the existing one. Netgroup names must be preceded by an ampersand (&) in the Administrator field.

You can restrict administrator access to a particular machine by listing it as shown in the following example:

root@server_name

4. Click the add button to add a new name, or click the change button to change the currently selected name.

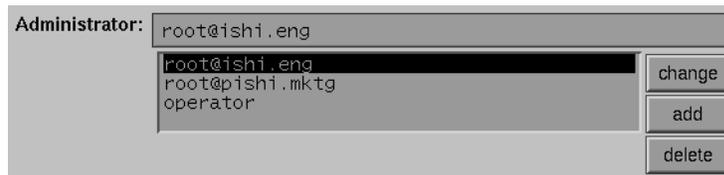


Figure 3-3 Server Administrator List

5. Click the Apply button to apply the changes or click the Reset button to cancel changes to the Administrator list.

To delete a name from the Administrator list, follow these steps:

1. Select the name in the scrolling list.
2. Click the delete button.
3. Click the Apply button (or click Reset to cancel the delete operation).

Every user or user group listed in the Administrator scrolling list has permission to modify the NetWorker server configuration.

Note: If you get the error message “user *user_name* needs to be on administrator list,” it means you do not have permission to make configuration changes.

Changing the NetWorker Server

This section describes how to change to another NetWorker server on a network.

To open the Change Server window, click the Change Server speedbar button (shown in Figure 3-4) or choose Change from the Server menu.



Figure 3-4 Change Server Button

The Change Server window appears, as shown in Figure 3-5.

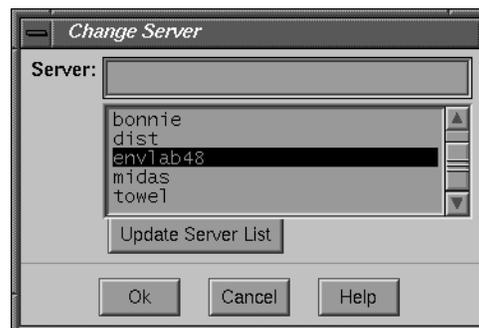


Figure 3-5 Change Server Window

Follow these steps to change to another NetWorker server:

1. Click to highlight the *hostname* of a server in the Server scrolling list, or enter a valid *hostname* in the Server field. If the hostname you want is not displayed in the Server list, click the Update Server List button to display all the NetWorker servers on your network.

Caution: When you click the Update Server List button, NetWorker searches for every server available to your machine. On large networks, this search could take awhile.

2. Click Ok to complete the operation and close the window.

Monitoring and Controlling Backups

NetWorker provides different windows for monitoring and controlling your backups: the NetWorker Administrator window, the Groups window, and the Group Control window.

In this section, you will perform the following tasks in the Group Control window:

- Monitor a scheduled backup.
- View the details of a scheduled backup.
- Preview a backup .
- Immediately start a scheduled backup.
- Stop a scheduled backup.
- Restart a scheduled backup.

Navigating the Group Control Windows

This section describes the Group Control windows and how to use them.

Open the Group Control window by clicking the Group Control speedbar button (shown in Figure 3-6) or by choosing Group Control from the Server menu.



Figure 3-6 Group Control Button

The Group Control window appears, as shown in Figure 3-7.

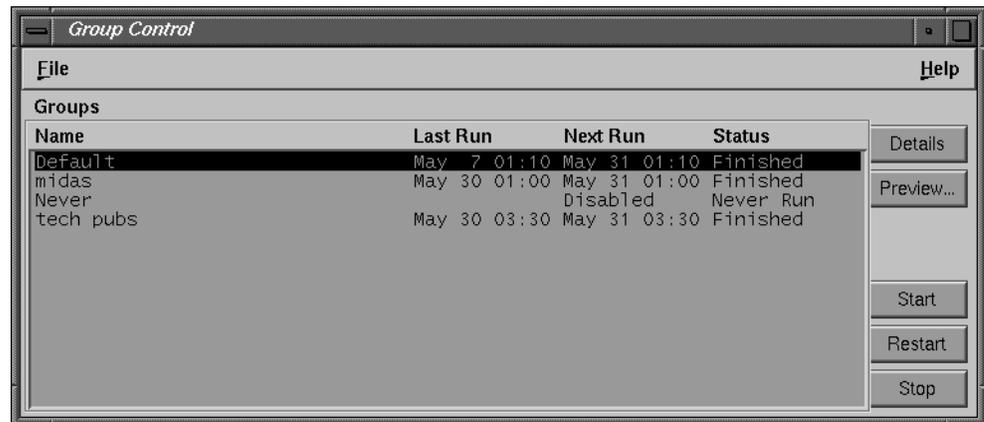


Figure 3-7 Group Control Window

The Groups scrolling list displays information about all of the backup groups known to the NetWorker server:

- Name—displays the name of the backup group.
- Last Run—lists the date and time the group was last backed up.
- Next Run—shows the date and time for the next enabled scheduled backup.
- Status—displays one of the following messages about the progress of your backup:
 - Running—the group is currently backing up.
 - Never Run—the group has never been backed up.
 - Finished—the backup has finished.
 - Not Finished—the backup is still running or has exited prematurely.
 - Preview Run—the preview of a backup is running.
- Start button—click to start a back up immediately.
- Restart button—click to restart a backup that has been stopped.
- Stop button—click to stop a scheduled backup.
- Details button—click to view details of completed backups.
- Preview button—click to preview an upcoming backup.

Use the Groups window to create new groups, enable group backup, and set backup time for a group. See “Configuring Backup Groups” on page 159 for more information.

Group Control Details Window

To view more detailed information about a completed group backup, click the Details button in the Group Control window. NetWorker displays the Group Control Details window, as shown in Figure 3-8.

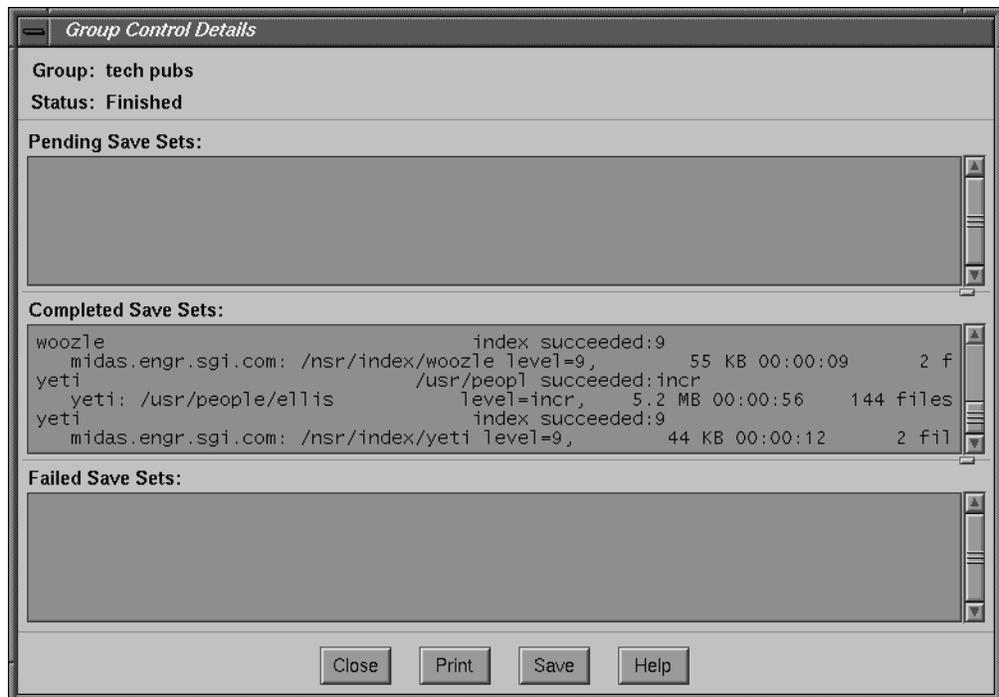


Figure 3-8 Group Control Details Window

There are three message fields in the Group Control Details dialog box:

- Pending Save Sets—displays the filesystems that have not yet been backed up.
- Completed Save Sets—displays the filesystems that NetWorker has successfully backed up.
- Failed Save Sets—displays the filesystems that NetWorker was unable to back up.

The Group Control Details dialog box displays save sets in the process of backing up and save sets that have completed their backups. Use this dialog box to determine which client save sets backed up successfully and which save sets have failed.

If NetWorker was unable to back up save sets for a client in the group, determine the cause for the failed backup.

Typical reasons for failed save sets include the following:

- The server crashed during backup.
- The client crashed during backup.
- The network connection failed during backup.

Tip: If you set Client retries to a number greater than zero in the Groups window and NetWorker fails to back up a client, you may temporarily see filesystems in the Failed Save Sets scrolling list while NetWorker retries backing up the client.

Previewing a Group Backup

You can preview a group backup by clicking the Preview button in the Group Control window. The Group Control Details dialog box appears, but instead of displaying past information about completed group backups, the dialog box displays information about how a group will perform during its next scheduled backup.

NetWorker actually simulates a backup for the group you select for preview in the Groups scrolling list in the Group Control window. The preview feature enables you to determine if there are potential problems with an upcoming group backup so you can take the necessary steps to avoid them.

Starting a Scheduled Backup Immediately

You can immediately start a scheduled backup by clicking the Start button in the Group Control window. This initiates an immediate backup of the group selected in the window.

To start a backup immediately from the Group Control window, follow these steps:

1. Select the group you want to back up immediately from the Groups scrolling list.
2. Click the Start button. A confirmation dialog box appears, as shown below.



3. Click OK to start the backup or Cancel to discontinue.

NetWorker immediately backs up the clients in the group, overriding the backup schedule, and displays "running" in the Status field of the Group Control window.

Tip: You can also start an immediate backup by selecting the Start now button in the Groups window (available from the Customize menu).

Stopping a Backup Group

If you need to stop a scheduled backup in progress, use the Stop button in the Group Control window.

To stop a scheduled backup in progress, follow these steps:

1. Select the group you want to stop from the Groups scrolling list.
2. Click the Stop button. A confirmation dialog box appears, as shown below.



3. Click OK to stop the backup or Cancel to continue.

As soon as the current save set completes backing up, NetWorker halts the scheduled backup and displays "Not Finished" in the Status field in the Group Control window.

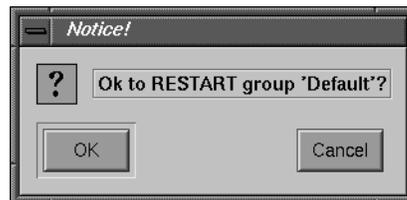
Restarting a Backup Group

If you want to restart a scheduled backup that you stopped, use the Restart button in the Group Control window.

Follow these steps to restart a backup:

1. Select the name of the halted backup group from the Groups scrolling list in the Group Control window.
2. Click the Restart button.

A confirmation dialog box appears, as shown below.



3. Click OK to restart the backup or Cancel to discontinue.

NetWorker resumes the scheduled backup for the group and displays “running” in the Status field.

Handling Open Files

If a client’s open files change during a scheduled backup, NetWorker backs up the file and “notices” that it is changing. A warning message similar to the following displays in the Group Control Details window:

```
warning: file name changed during save
```

You may restart the backup group, back up the client manually, or allow NetWorker to back up the client during the next scheduled backup.

Archive Request Control

The Legato NetWorker Archive Application provides file archiving and retrieval services to a range of client machines. If you purchased this optional extension for your NetWorker server, the Archive Request Control choice on the Server menu is enabled.

For complete instructions on using the NetWorker Archive Application, see Chapter 8, "Archive Application."

Registering NetWorker

To see licensed components of the NetWorker server, choose Registration from the Server menu. The Registration window appears, with an enabler code, host ID, and expiration date for each component. IRIX NetWorker does not use the Auth code field.

Most copies of IRIX NetWorker arrive with a temporary license. See the *IRIX NetWorker Installation Guide* for information on obtaining a permanent license.

If you fail to obtain a permanent license, NetWorker automatically disables itself 45 days after installation. Once the server disables itself, you will not be able to back up any more files. However, you *will* be able to recover previously backed-up files.

Managing Media and Backup Devices

This chapter discusses the choices in the **Media** menu used for managing media and standalone backup devices. Commands in the **Media** menu for operating an autochanger (also called a jukebox) are described in Chapter 7, “Autochanger Software Module.”

This chapter describes the following operations:

- labeling backup volumes
- mounting and unmounting backup volumes
- manually managing backup volumes
- cloning backup volumes
- using volume pools
- adding and deleting backup devices

Tip: If you back up to an autochanger (also called a jukebox), you need to purchase the Autochanger Software Module. A separate Autochanger Software Module license is required for each autochanger connected to the NetWorker server. Autochangers are licensed according to the number of storage slots they contain.

Labeling and Mounting Backup Volumes

This section applies only if you label and mount volumes in a standalone backup device. If you need to label multiple volumes for an autochanger, see Chapter 7, “Autochanger Software Module.”

You must label and mount a backup volume before NetWorker can use it to back up or recover files. NetWorker uses the volume labels to maintain a *record* of each backup volume in the online media index. This index is used by NetWorker to determine which volumes are needed for backing up or recovering data.

Your NetWorker server may have one or more backup devices connected to it. The current device selected in the Devices display of the NetWorker Administrator window also appears in the Jukebox Labeling dialog box.

Navigating the Label Dialog Box

Use the Label dialog box to label backup volumes in a standalone device.

To open the Label dialog box click the Label speedbar button (shown in Figure 4-1) or choose Label from the Media menu.



Figure 4-1 Label Button

The Label dialog box appears, as shown in Figure 4-2.

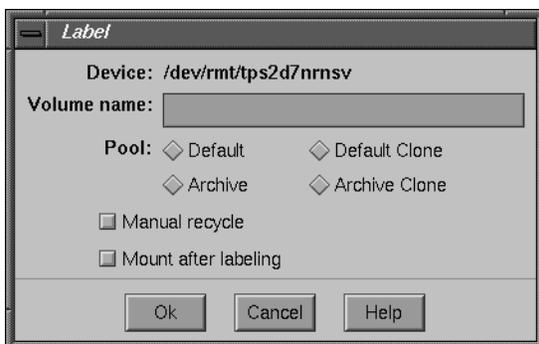


Figure 4-2 Label Dialog Box

The contents of the Label dialog box are described below:

- Device field—contains the name of the selected device. To select another device, click Cancel to close the Label dialog box, then highlight the new device in the NetWorker Administrator window, and reselect the Label command.

- **Volume name field**—contains the label that will be assigned to the volume and displays the name of the volume mounted in the backup device. NetWorker automatically assigns the next sequential label from the label template associated with the chosen pool. If you do not want NetWorker to use the preselected label template, delete the contents in the Volume name field and enter a new name for the backup volume.
- **Pool choices**—choose a volume pool from the Pool choices. Only pools enabled from the Pools window appear in the Label dialog box (most preconfigured pools are automatically enabled). NetWorker automatically uses the label template associated with the selected pool to label a backup volume. For more information on volume pools, see the section “Using Volume Pools” on page 85.
- **Manual recycle toggle button**—select if you want to manually recycle a backup volume. When a backup volume is marked “manual recycle,” NetWorker disregards the assigned retention policy. Manual recycle mode is explained in the section “Volumes Scrolling List” in this chapter.
- **Mount after labeling toggle button**—selected by default. Deselect if you only want to label volumes and not mount them. You may prelabel as many volumes as you like to label at one time without mounting them by deselecting the Mount after labeling toggle button.
- **Ok button**—click to label or mount a backup volume. If you make a mistake, click the Cancel button.

Any time you try to relabel a volume that already has a valid NetWorker label, NetWorker prompts you with a notice to make sure you really want to relabel the volume.

Labeling and Mounting in One Operation

To label and mount a backup volume in one operation, follow these steps:

1. Place a blank volume in the NetWorker server backup device.
2. Highlight the device in the Devices display in the NetWorker Administrator window.
3. Open the Label dialog box by clicking the Label speedbar button or by choosing Label from the Media menu.
4. Choose a volume pool from the Pool choices.
5. Delete the name from the Volume name field. If you want a label different than the one chosen by the label template, enter the new name.

6. Make your selection for Manual recycle and click the Mount after labeling toggle button.
7. Click the Ok button when you are ready to label and mount the volume. If you make a mistake, click the Cancel button to cancel the operation.

A dialog box appears, asking you to confirm your request. Click OK to confirm the labeling request or Cancel to quit the request.

When you complete labeling and mounting, the backup volume name appears in the Devices window of the NetWorker Administrator program beside the *pathname* of the device. (The name of the volume appears only if it has been mounted.)

Caution: Relabeling a backup volume destroys all records of its contents under the old label. That is why NetWorker asks for confirmation with the prompt “About to relabel.” For more information about relabeling backup volumes, see “Manually Managing the Online Indexes” on page 36.

Understanding How to Label Backup Volumes

You must label a backup volume before you can use it for backing up files. Labeling a backup volume provides NetWorker with a unique name for tracking and recognizing the media.

Note: Use an autochanger (jukebox) to automate volume labeling and mounting operations. See “Chapter 7, “Autochanger Software Module,” for details.

Every backup volume belongs to a volume pool. You may use the preconfigured pools or create your own. If you do not select a volume pool for your backup volumes, NetWorker automatically uses the preconfigured Default pool.

Each volume pool has a matching label template associated with it. NetWorker uses the label template associated with the pool you choose in the Label dialog box. Backup volumes are labeled according to the rules of these label templates. Label templates provide a method for consistently naming and labeling your backup volumes. Use the preconfigured label templates, or create custom label templates, using the Label Templates dialog box.

For more information about volume pools, see the section “Using Volume Pools” on page 85. For more information about label templates, see the section “Using Label Templates” on page 150.

NetWorker automatically labels backup volumes with the next sequential label from the associated template. Label names are recorded internally on the media, creating an internal volume label. Each backup volume should have an adhesive label attached to it that matches the internal volume label.

Each backup volume requires a unique label. If you choose not to use a preconfigured label template, you may create your own label templates with names appropriate for your network. No matter what naming scheme you choose, keep in mind that it must be easy for operators and administrators to understand.

Create individual label names not associated with a template by deleting the label template name from the Volume Name text box in the Label dialog box and entering a unique label name.

Tip: Each NetWorker server has its own collection of backup volumes. If you have more than one NetWorker server, you may find it helpful to label your backup volumes with the name of the server and a number. For example:

atlas.001 This backup volume is from the NetWorker server named *atlas*.

mars.010 This backup volume is from the NetWorker server named *mars*.

Use the Label dialog box to label a backup volume or to label and immediately mount a volume in the backup device. Once labeled and mounted, a volume is available for backups. You may label several volumes at once without mounting any of them by deselecting the Mount after labeling toggle button.

Mounting and Unmounting Backup Volumes

Before you can back up or recover files from a volume, you must mount the volume in the server backup device. If you use an autochanger, NetWorker automatically mounts backup volumes for backups and recovers.

All backup devices recognized by NetWorker appear in the Devices window of the NetWorker Administrator program. Select the device you need from the Devices window before you begin to mount or unmount a backup volume.

Use the speedbar buttons (shown in Figure 4-3) to mount and unmount backup volumes, or choose the Mount and Unmount commands from the Media menu.



Figure 4-3 Mount and Umount Buttons

Mounting Backup Volumes

When a backup volume is required for a standalone device, NetWorker uses the following hierarchy to select a volume from the appropriate backup volume pool:

- an already mounted, appendable volume
- an already mounted, recyclable volume not currently in use
- an already mounted, unlabeled volume not currently in use and in a device for which Auto media management is enabled in the Devices window
- an appendable volume not currently mounted in the device
- a recyclable volume not currently mounted in the device

To mount a backup volume in a specific device, follow these steps:

1. Select the device name in the Devices display of the NetWorker Administrator window.
2. Click the Mount speedbar button.

The Devices display changes to show the name of the mounted backup volume listed next to the device *pathname*.

Caution: To perform an unattended backup using a standalone device, you must premount backup volumes.

If you try to back up files when a backup volume is not mounted, or when a backup volume is full, NetWorker requests a new writable volume with the following message in the Pending display:

```
media waiting: backup to pool 'Default'  
waiting for 1 writable backup tape or disk
```

You receive a message suggesting that you mount a volume, relabel a volume, or label a new volume.

If you are recovering files, NetWorker requests the backup volume name you need to mount in the device:

```
media waiting: recover waiting for 8mm 5GB tape_volume_name
```

If you need more than one backup volume to recover the files, the Pending display gives you a list of all the backup volumes in the order they are needed. During the recovery process, NetWorker requests the backup volumes it needs, one at a time. If you back up to an autochanger, NetWorker automatically mounts backup volumes stored in the autochanger.

Unmounting Backup Volumes

To unmount a backup volume from a specific backup device, follow these steps:

1. Select the name of the device in the Devices display.
2. Click the Unmount speedbar button.

Finding a Backup Volume Name

If the physical label on the backup volume is missing or illegible, you can determine its name by using one of two methods:

- Click the Mount speedbar button to mount the volume in the backup device. The volume name appears next to the device name in the Devices display.
- Load the volume into a device and select the Label speedbar button. When the Label dialog box appears, the name of the magnetic label is displayed in the Volume name field. Click Cancel to avoid accidentally renaming the volume.

In either case, make sure to write the label name onto the physical tape label.

Automatically Managing Backup Volumes

NetWorker uses browse and retention policies to manage your backup volumes and their save sets automatically. However, you may manually override the automatic policies by using the different commands in the Volumes window.

For more information about browse and retention policies, see “Creating Index Policies” on page 142.

Manually Managing Backup Volumes

NetWorker provides flexibility in determining how to manage your backup volumes and their save sets. You can allow NetWorker to automatically manage your volumes, or you can use the Volumes window to manage your backup volumes manually.

Navigating the Volumes Window

The Volumes window contains information about the media index. It displays information about the collection of backup volumes known to the NetWorker server and the save sets they contain.

To view information about your backup volumes, open the Volumes window by clicking the Volumes speedbar button (shown in Figure 4-4), or by choosing Volumes from the Media menu.



Figure 4-4 Volumes Button

The Volume Management window appears, as shown in Figure 4-5.

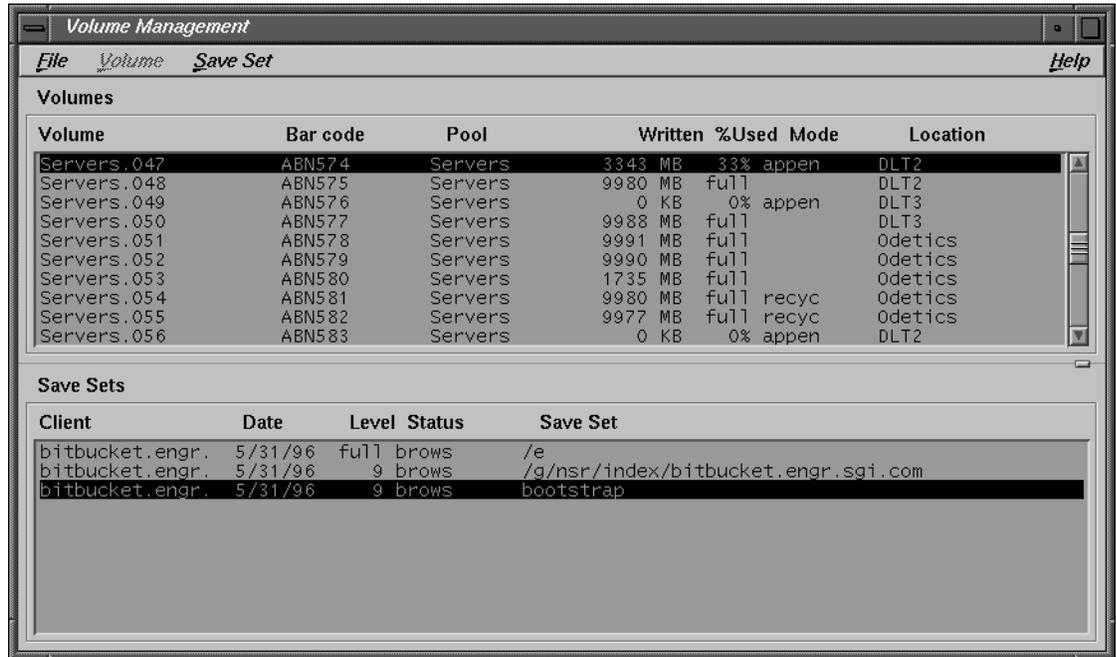


Figure 4-5 Volume Management Window

The Volume Management window displays the following information.

- Volumes scrolling list—contains the names of the backup volumes, sorted in alphabetical order for the current server.
- Save Sets scrolling list—contains the names of the save sets stored on the backup volume selected in the Volumes scrolling list.

Volumes Scrolling List

The Volumes scrolling list contains the names of the backup volumes for the current server. The list is sorted in alphabetical and numerical order and contains the following information:

- Volume—backup volume name. One of the following designations may appear at the end of the volume name:

- (A)—indicates an archive volume.
- (R)—indicates a read-only volume.
- Bar code—bar code label, if present.
- Pool—name of the pool to which the backup volume belongs.
- Written—amount of data written on the volume.
- %Used—percentage of the estimated total capacity of the backup volume used. 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 tape mark has not been reached.

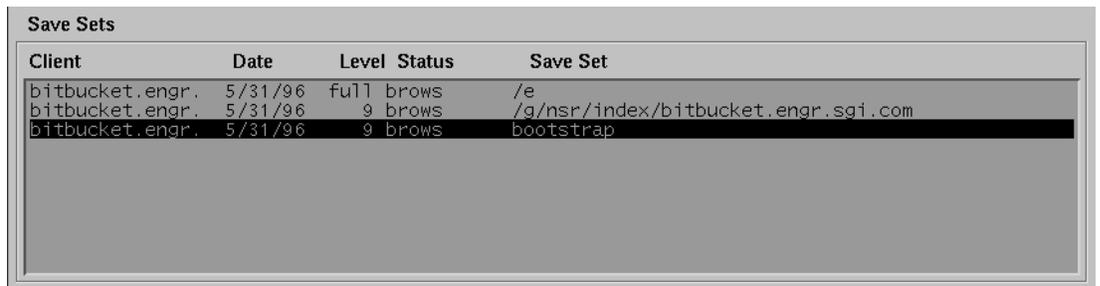
The %Used column displays full when there is no more space on the backup volume for additional data and the end of tape mark has been reached. If the save sets on the volume have not yet passed the time period specified by their browse and retention policies, the Mode column is blank. If the save sets have passed their browse and retention policies, recyc appears in the Mode column, indicating the volume is available for recycling.

- Mode—status or kind of backup volume.
 - appen—(appendable) there is room for more data on the backup volume.
 - recyc—(recyclable) all save sets on the volume have passed the browse and retention policies, and the volume is ready to be relabeled. Once the volume has been relabeled, it can be overwritten with new data.
 - man—(manual recycle) both the volume and its save sets are marked for manual recycling. If the appen or recyc modes are already assigned to the volume, the man mode is displayed as an additional mode.
 - The following volume mode indicator may appear next to the volume name in the Volume column:
 - (R)—volume has a mode of “read only” assigned manually with the Change Mode command. When the mode is manually changed to read only, the previously assigned mode is removed and the Mode column is blank for the volume.
 - See “Changing the Mode of a Backup Volume” on page 74 for information on manually changing the mode of a volume.
- Location—location of the backup volume, entered by the user as a reminder.

Note: When you back up to volumes marked `appen` (appendable), you are adding more data to the volume, not overwriting it. Therefore, you can leave volumes marked `appen` in the server backup device to receive more backed-up data. When a backup volume becomes full, it is marked full and NetWorker requests that the operator or autochanger mount another volume until the backup is complete.

Save Sets Scrolling List

In the Volume Management window, use the Save Sets scrolling list, shown in Figure 4-6, to view the save sets stored on a volume.



Client	Date	Level	Status	Save Set
bitbucket.engr.	5/31/96	full	brows	/e
bitbucket.engr.	5/31/96	9	brows	/q/nsr/index/bitbucket.engr.sgi.com
bitbucket.engr.	5/31/96	9	brows	bootstrap

Figure 4-6 Volume Save Sets List

The Save Sets scrolling list provides the following information about each save set stored on the currently-selected volume:

- **Client**—name of the machine that created the save set.
- **Date**—date the save set was created.
- **Level**—level of backup that generated the save set. Backup levels apply only to scheduled backups. If a Level is not displayed for a save set, the save set was generated by a manual backup.
- **Status**—type of save set:
 - **brows**—(browsable) indicates a file index entry exists for the save set. Files in a save set marked **brows** have not passed the time period specified by the client browse policy. These files can be browsed and marked for recovery in the NetWorker Recover window.
 - **recov**—(recoverable) indicates that the save set entry has been removed from the file index. The entry could have been removed automatically due to the browse policy for the client, or manually by clicking the Remove oldest cycle button in

the Indexes window. If you need to recover the data, use the Save Set Recover window. If you want to recover both the data and the file index entries, use the **scanner** command. See the scanner(1M) reference page for information on using this command.

- **recyc**—(recyclable) indicates that the save set has passed both the Browse and Retention policies time periods and is available for relabeling. The volume must be relabeled before you attempt to overwrite it with new backups.
- **scann**—(scanned-in) indicates that the save set was scanned in using the **scanner** command. Standard NetWorker index management policies do not apply to scanned data. You must remove these save sets manually from the index. See “Navigating the Indexes Window” on page 40 for instructions on manually removing save sets from the index.
- **inpro**—(in-progress) indicates the save set is in the process of being backed up.
- **abort**—(aborted) indicates that the backup was aborted manually by the administrator or that the system crashed.
- **susp**—(suspect) indicates that a previous recovery attempt failed. In this case, the susp status is displayed in addition to one of the previously described index status types.
- **Save Set**—path of the filesystem contained in the save set. The save set column also includes clone information. If the save set has a clone, it is marked has clones and the cloned save set is marked cloned *date time*.

Manually Managing Backup Volumes

Use the Volumes window to manually perform the following volume management tasks:

- change the status of a volume
- set a location for a volume
- change the mode of a volume
- mark entire backup volumes and their save sets as recyclable
- remove a volume from the media index

Changing the Status of a Save Set

NetWorker automatically marks the backup volume status as “suspect” and displays *susp* in the Volumes window if an error occurred while trying to read the media. If NetWorker does not encounter media read errors, the volume status remains “normal.”

If you know that data in a save set is really not suspect, you may want to manually change the status of the save set from Suspect to Normal. For example, if the drive heads were dirty when NetWorker tried to read the data, the data itself is still reliable.

To manually change the status of a save set, follow these steps:

1. Select the save set from the Save Sets scrolling list in the Volumes window.
2. Use the Change Status choice in the Save Set menu of the Volumes window to manually change the status of the save set, as shown in Figure 4-7.

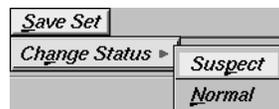


Figure 4-7 Save Set Change Status Menu

If you want to undo changes made manually to the status of a save set, you can use the Change Status command to return the save set status to its former setting.

Setting the Location for Backup Volumes

You can add a location to backup volumes as a reminder of where the volume is currently stored. This helps you find a volume when it is needed for recovering a file. Volume locations might include a vault, shelf, autochanger job pack, or any place you physically store backup media.

To set the location for a backup volume, follow these steps:

1. Highlight a backup volume in the Volumes scrolling list.
2. Choose Set Location from the Volume menu. The Set Location dialog box appears, as shown in Figure 4-8.



Figure 4-8 Set Location Dialog Box

3. Enter the location of the backup volume. You may want to specify the location of backup volumes stored on a shelf or volumes stored offsite.
4. Click the Ok button to designate the location for the selected backup volume. The Volumes window now reflects the location for that backup volume.

After you apply a location to a volume, that location is added to the scrolling list in the Set Location dialog box and can be applied to any volume listed in the Volumes window.

If you move backup volumes from one autochanger to another, inventory them so NetWorker can recognize them and update their location with the name of the new autochanger. If you move volumes offsite, manually assign them the location of the offsite facility before removing them from the autochanger.

Tip: When NetWorker requests that you mount a backup volume during recovery, you can find the volume by referring to its location in the Volumes window.

Changing the Mode of a Backup Volume

The mode of a backup volume determines whether NetWorker can write data to it during a backup. See “Volumes Scrolling List” for a description of volume modes. NetWorker automatically assigns the appen and recyc modes. This section provides a detailed description of how to change volume mode manually.

The Volume menu in the Volumes window contains the Change Mode choice, which allows you to change the mode of a currently selected volume to appendable, read-only, or recyclable, as shown in Figure 4-9.

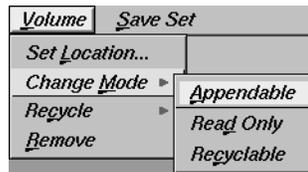


Figure 4-9 Volume Change Mode Menu

Table 4-1 provides a guide to existing volume modes that you can change manually.

Table 4-1 Changing Volume Modes Manually

Mode of Current Volume	Change to Appendable?	Change to Read Only?	Change to Recyclable?
Recyclable	yes ^a	yes	no
Appendable	no	yes	yes

a. You can mark these volumes appendable, but NetWorker may not be able to write data to them for a variety of reasons.

When the %Used is full or the mode of a volume is read only, NetWorker cannot write data to that volume. NetWorker automatically indicates a %Used of full when a volume is filled with data. An administrator manually assigns a status of read only to a volume so data cannot be written to it, whether it is full or not.

Note: If the status of all of the save sets is recyc, NetWorker automatically marks the volume as recyclable.

To change the mode of a volume, follow these steps:

1. Click the Volumes speedbar button or choose Volumes from the Media menu.
2. Highlight the volume in the Volumes scrolling list.
3. Choose Change Mode from the Volume menu in the Volumes window.
4. Drag and release on a mode from the Change Mode choices: Appendable, Read Only, or Recyclable. A confirmation box appears asking if you want to change the volume to the mode you have chosen, as shown in Figure 4-10.

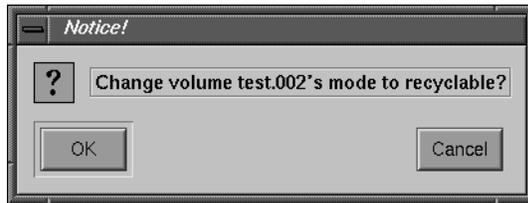


Figure 4-10 Notice to Change Volume Mode

5. Click the Ok button if you want to change the mode of the selected volume. Notice the mode of the volume changes in the Volumes scrolling list. Click Cancel to leave the mode unchanged.

The previously assigned mode is removed and the new volume mode is indicated next to the volume name in the Volumes display.

Caution: Do not manually change the mode of a volume to recyclable if you still need any of the save sets on the volume and do not have a cloned volume. This could lead to loss of data.

If a cloned volume does not exist when you manually change the mode of a volume to recyclable, all of the save sets on the volume change to a status of recyclable, regardless of the save set browse or retention policies. If a cloned volume exists, the Status of the save sets remains unchanged. Original backup volumes and cloned volumes share the same set of save set entries. As long as one of the volumes exist, NetWorker does not remove save set entries, since they are still needed for the remaining volume.

Recycling Volumes

You can override the volume retention policy by changing it to manual recycle. Choose Recycle from the Volume menu; the Manual and Auto choices appear to the right, as shown in Figure 4-11.

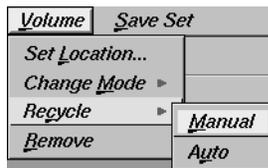


Figure 4-11 Volume Recycle Menu

You may want to change to manual recycle if you need to keep the save sets on a volume longer than its retention policy specifies. You can change a volume marked manual recycle back to automatic recycle, reverting to the original retention policy.

To change a backup volume to manual recycle, follow these steps:

1. Select a volume from the Volumes scrolling list.
2. Choose Recycle from the Volume menu in the Volumes window; the Manual and Auto subchoices appear.
3. Choose Manual from the Recycle menu.

NetWorker notes the volume by displaying man (manual) as the volume mode in the Volumes scrolling list of the Volume Management window.

If you want to remove manual recycle from the volume, simply highlight the volume again and choose Auto from the Recycle menu.

Removing a Backup Volume

To remove a volume from the media index, choose Remove from the Volume menu in the Volumes window, as shown in Figure 4-12.

Typically, you never remove a backup volume from the media index unless the volume has become physically damaged or unusable.



Figure 4-12 Volume Remove Menu

When you select a volume for removal, NetWorker checks to see if the volume has a clone. If so, NetWorker does not remove the entries from the online file index. If not, NetWorker removes the entries from the online file index.

Tip: If a backup volume is lost or destroyed, remove its entries from both the file and media indexes. Otherwise, NetWorker may request the nonexistent volume for data recovery.

To remove a volume from the media index, use the following steps:

1. Highlight the volume in the Volumes scrolling list. The volume save sets appear in the Save Sets scrolling list.
2. Choose Remove from the Volume menu. The notice “Remove volume *name* from the media index?” appears, as shown in Figure 4-13.

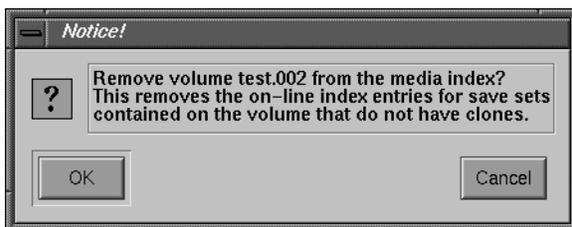


Figure 4-13 Notice to Remove Volume

3. Click the Ok button to remove the volume from the media index.
If a removed volume does not have a clone, the save sets are also removed.
4. Click the Cancel button if you do not want to remove the volume from the media index.

Caution: If you accidentally delete a backup volume from the media index and have not relabeled it yet, you can still recover the data by using the **scanner** program. See the scanner(1M) reference pages for more information.

Cloning Backup Volumes

NetWorker protects the data on your network by backing up files every day. Cloning provides additional protection by allowing you to make identical copies of the data on your backup volumes as needed.

The original backup volume is the *source* volume, and the volume that receives the cloned data is the *destination* volume. The destination volume must be different from the source volume. NetWorker tracks cloned volumes (destination volumes) separately and marks them as clones in the media database.

You must possess all of the following before you can clone save sets onto backup media using NetWorker:

- NetWorker with concurrent devices support (IRIX NetWorker has this)
- multiple backup devices
- in the Server window, the value of Active devices set to 2 or more
- source and destination backup volumes

NetWorker also has the ability to clone individual save sets manually and automatically. See “Cloning Save Sets” on page 113 for more information.

Navigating the Clone Volume Windows

Use the Clone Volume window for cloning all of the data on a backup volume. This window provides information about the backup volumes for the server currently in use and provides a starting point for beginning the cloning process.

To open the Clone Volume window, choose Clone Volume from the Media menu, shown in Figure 4-14.



Figure 4-14 Media Clone Volume Menu

The Clone Volume window appears, as shown in Figure 4-15.

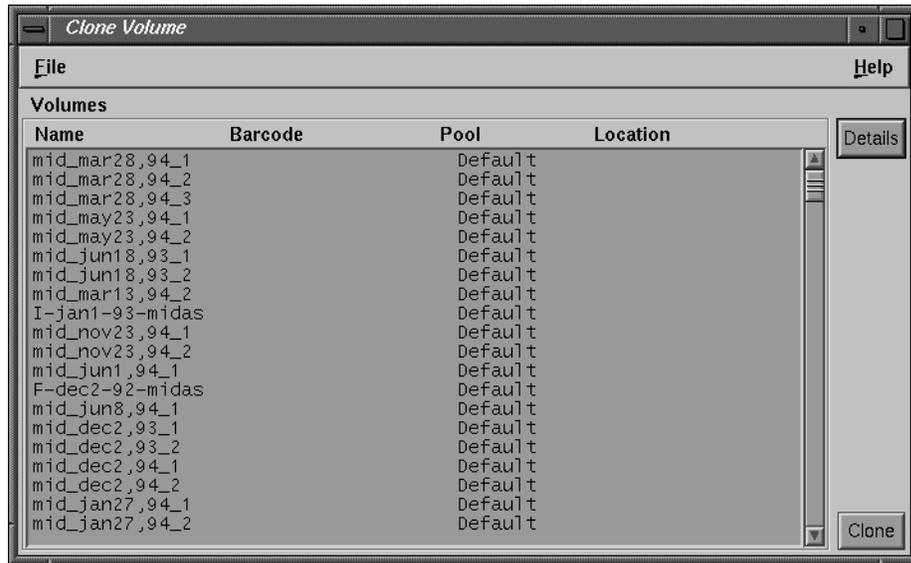


Figure 4-15 Clone Volume Window

The contents of the Clone Volume window are described below:

- Volumes scrolling list—contains a list of all backup volumes known to the server. The list includes the volume and pool name, bar code label, and the volume location. You can clone save sets from an unlimited number of backup volumes at the same time. Clone backup volumes after a scheduled backup has completed to avoid overloading your NetWorker server and backup devices.
- Details button—click to view backup volume details. The Clone Volume Details dialog box appears, as shown in Figure 4-16.

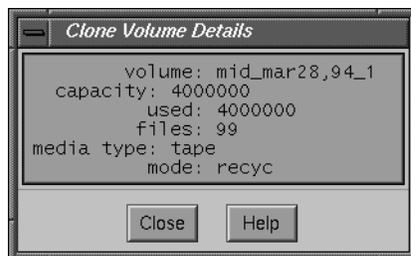


Figure 4-16 Clone Volume Details Window

The Clone Volume Details dialog box contains the following items:

- volume—name of the volume.
 - capacity—size of the media.
 - used—space used on the backup volume.
 - files—number of files on the backup volume.
 - media type—kind or type of backup volume.
 - mode—mode of a backup volume is either appen (appendable), full (full), or recyc (recyclable). See “Volumes Scrolling List” on page 69 for an explanation of the different volume modes.
- Clone button—click to open the Clone Volume Status dialog box and begin cloning selected volumes. The Clone Volume Status dialog box appears, as shown in Figure 4-17.

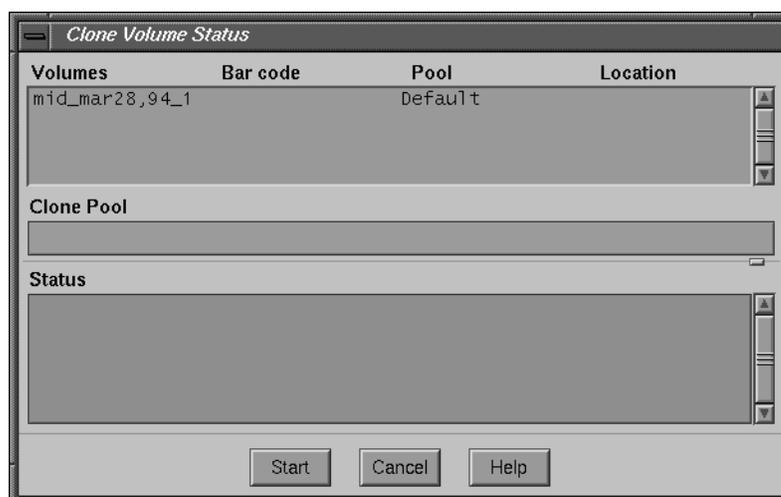


Figure 4-17 Clone Volume Status Dialog Box

The Clone Volume Status dialog box contains the following items:

- Volumes scrolling list—lists the volume(s) selected for cloning, the pool to which it belongs, and the volume location.

- Clone Pool field—enter the name of the clone pool. Cloned backup volumes must belong to a clone pool, either one you have created or one of the preconfigured clone pools. If you do not enter a pool name, NetWorker automatically selects the Default Clone pool.
- Status scrolling list—displays the progress of the cloning operation. The same messages appear in the NetWorker Administrator window.
- Start button—click to begin cloning.

Example—Cloning Backup Volumes

In this example, the network operator needs to clone the full backups from the engineering department for offsite storage. The operator wants only full backups from engineering cloned to the destination volumes. Because the operator does not want to mix the cloned data, a blank destination volume must be loaded.

The source volume in this example is already mounted in one of the backup devices because a full backup has just completed. If the source volume is not mounted when you begin cloning, you receive a message in the Pending display of the NetWorker Administrator window requesting that you load the volume.

1. Load, label, and mount a blank destination volume in the media drive.
Note: If you use a standalone device for NetWorker backups, a message appears in the Messages display of the NetWorker Administrator window requesting that you mount a volume. If you use an autochanger (jukebox) for NetWorker backups and do not load, label, and mount a volume, NetWorker selects an “appendable” volume from the clone pool entered in the Clone Status window.
2. Open the Clone Volume window by choosing Clone Volume from the Media menu.
3. Select the volumes containing full backups for the engineering department from the Volumes scrolling list.
4. Click the Clone button in the Clone Volume window to open the Clone Volume Status window.
5. Enter the name of the clone pool in the Clone Pool field.
6. Click the Start button in the Clone Volume Status window to begin the cloning process.

The Status scrolling list displays messages showing the progress of the clone operation.

Understanding the Volume Cloning Feature

Cloning backup volumes is useful for several purposes, including: sending data offsite, replacing old volumes, and testing the integrity of your data.

Most media has a limited shelf life. Consequently, you need to routinely replace your backup volumes. If you have old backup volumes that contain important data, clone them before they expire to continue to ensure the safety of the data.

Cloning also helps you determine the integrity of backed-up data. NetWorker displays an error message if it has difficulty reading data from the original backup volume during the cloning process.

When you clone backup volumes, you are actually copying all of the save sets that have been backed up to that volume; you are not creating an identical copy of the backup volume. Consider cloning your backup volumes after a scheduled backup has completed, so you do not overload your NetWorker resources and backup devices.

Creating clones of your media provides the ability to store copies of your data (save sets) offsite or in a vault for protection. Cloning backup volumes also allows you to share or transfer important data to other company facilities.

Note: Cloned media does not contain the file or media indexes used for recovering data. Since an offsite facility using NetWorker will not have these indexes, they must use the **scanner** command to recover the cloned data. For complete information about this command, see the scanner(1M) reference page.

Backup volumes can only be cloned manually. To clone selected data (save sets) automatically, use the Groups window to enable cloning for individual groups.

NetWorker may request more than one source volume to complete the cloning operation. This is because you are cloning the save sets, which may span multiple backup volumes. Furthermore, because media varies in size and capacity, you may need more destination volumes to hold the cloned data than the source volumes required.

NetWorker automatically selects the next available volume marked “appendable” from the clone pool for cloning if you have an autochanger. If you do not want to mix newer cloned data with existing cloned data, you must load, label, and mount a different volume in the backup device before starting the cloning process.

Entering Index Entries for Cloned Data

During a backup, NetWorker adds entries to both the *media* and *file* indexes after it sends data to the backup device. File index entries represent saved client files (save sets).

By contrast, during the cloning process, NetWorker copies data from one backup volume to another and creates entries for the cloned backup volume and its save sets in the media index only. During the cloning process, NetWorker does not back up client files, so no new client save sets are created or entered in the file index.

Both original and cloned volumes share the same file index entries since the data is identical on both volumes. When only one entry remains in the media index for a backup volume (cloned or original) you can manually remove file index entries for save sets that have passed their browse policy, or NetWorker can automatically remove the save sets according to their browse policy.

To see the save sets contained on a backup volume, select the volume in the Volumes scrolling list. The Save Sets scrolling list displays all the save sets contained on the currently-selected volume. If the save sets have been cloned, you see has clones next to the save set name. If the save set is a clone, you see cloned next to the save set name.

For more information about cloning data see “Understanding the Cloning Save Sets Feature” on page 128.

NetWorker indicates whether the save sets on backup volumes are cloned or if they have clones. However, NetWorker does not give any indication in the Volumes scrolling list whether or not a backup volume is an original or a clone. In other words, NetWorker does not distinguish whether the backup volumes are clones or originals, but it does distinguish if the save sets on the backup volumes are clones or have been cloned.

Tip: To distinguish between cloned and original backup volumes, assign the cloned volume to a clone pool and include the term “clone” in the label template. NetWorker provides the following preconfigured clone pools and associated label templates: Archive Clone, Default Clone, Migration Clone, and PC Archive Clone.

Recovering from Cloned Volumes

Cloned volumes provide added security by creating another source for recovering your data in case the original backup volumes become damaged or destroyed. By storing your cloned volumes offsite in a secure location, you maintain a reliable source for recovering lost or damaged data.

NetWorker makes no distinction between original (source) and cloned (destination) volumes. NetWorker chooses the first volume with reliable data for recovery. Consequently, if a cloned volume is already mounted in the drive, NetWorker uses it for recovering data, instead of searching for the original source volume. NetWorker finds the volume containing the data it needs and checks to see if the data has been marked suspect, regardless of whether or not it is a cloned or original volume.

NetWorker considers volumes for recovery in the following order:

1. A volume mounted in the drive.
2. A volume stored in an autochanger.
3. A volume stored someplace other than in an autochanger.

Using Volume Pools

This section describes the *volume pools* feature, preconfigured pools, and how to create new pools. It also provides several useful examples.

Volume pools allow you to sort your data to selected backup volumes. A volume pool contains a collection of backup volumes that receive specific data sorted during backup. All NetWorker volumes belong to a pool, either preconfigured or one you create.

Cloned backup volumes must belong to a clone pool, just as archived volumes belong to an archive pool. For information on archiving, see Chapter 9, "Archive Application."

Sort data to a pool by selecting the type of backup data you want included in that pool. NetWorker uses the choices you make in the Pools window to sort backed-up data to specific backup volumes.

You can create as many pools as you need to effectively organize and sort your backup volumes. There is no limit on the number of new pools you can create.

NetWorker provides the following preconfigured pools for your immediate use: Archive, Archive Clone, Default, Default Clone, Full, NonFull, Migration, Migration Clone, PC Archive, PC Archive Clone, and Offsite. These and any you create fall into the following "pool type" categories: Backup, Backup Clone, Archive, Archive Clone, Migration, and Migration Clone.

For information about the preconfigured pools provided by NetWorker, see the section "Preconfigured Pools" on page 251.

All NetWorker pools require a label template, described in the section “Using Label Templates With Volume Pools” on page 95. Before creating a new pool, create a label template with the same name. We strongly recommend that you keep the pool names and their corresponding label templates consistent. If you do not use the exact name, try to choose names that have a logical connection.

Tip: For each pool, you must select at least one of the following choices: a group, client, save set, or level. A pool cannot be created without at least one of these.

For example, you can sort data by the following categories:

- by backup group
- by NetWorker client
- by save sets (filesystems)
- by backup levels—fulls, levels 1-9, incremental, or manual

You can further manage backed-up data by one of these methods:

- specifying a special label template
- sending all the backup data to a particular backup device
- choosing whether to store the index entries in the file index

Navigating the Pools Window

To open the Pools window, choose Pools from the Media menu, shown in Figure 4-18.



Figure 4-18 Media Pools Menu

The Pools window appears, as shown in Figure 4-19.

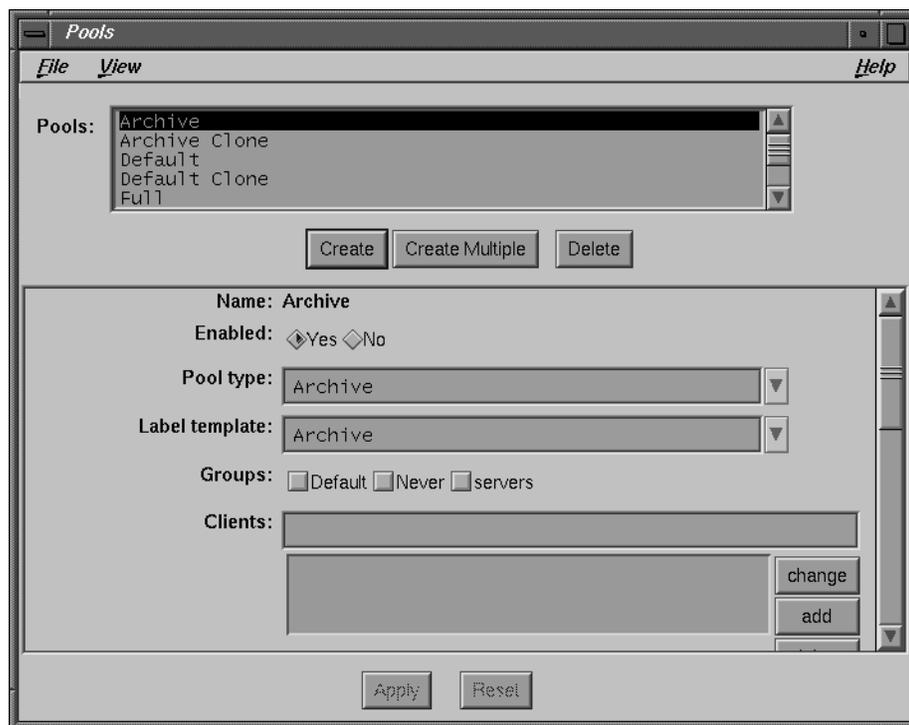


Figure 4-19 Pools Window

Use the scroll bar or resize the window to view the entire Pools window. The Pools window contains choices required for sorting client data to selected backup volumes.

You must make selections for at least one of the following choices: Groups, Clients, Save sets, or Levels. NetWorker does not sort data for Pools choices left unselected. For example, if you do not choose a level from the Levels choices, NetWorker includes all levels of backups for that pool. NetWorker sorts data based on the other selections you have made for that pool.

Tip: You cannot create pools that share identical settings for Pool type, Groups, Clients, or Save Sets. If any of these settings for a new pool match an existing pool, you receive a warning notice similar to that shown in Figure 4-20.



Figure 4-20 Notice About Pool Overlap

Below is an explanation of each of the elements in the Pools window and how to use them for modifying and creating custom pools:

- Pools scrolling list—displays all of the available pools (preconfigured pools and pools you have created).
- Create button—click to add a new pool.
- Delete button—click to eliminate a pool from the Pools scrolling list. You cannot delete a pool as long as it still contains backup volumes; you must first remove all of the backup volumes that belong to the pool by using the Volumes window.

Note: You cannot delete or modify any of NetWorker’s preconfigured pools.

- Name field—displays the name of the currently-selected pool and serves as the field for entering the name of a new pool.
- Enabled choices—enables and disables the pool. The preconfigured choice is Yes. If you want NetWorker to use a pool during the backup process, select Yes. If you want to exclude a pool during the backup process, select No.

Note: With the exception of the Full, NonFull, and Offsite pools, NetWorker’s preconfigured pools are already enabled.

- Pool type choices—select the appropriate type of pool for your backup volumes. You may select a preconfigured pool.
- Label template choices—select a label template, either a preconfigured label template or one created in the Label Templates window. Silicon Graphics highly recommends using the same name for both the label template and the pool.
- Groups choices—choose the groups you want included in the pool. The groups displayed represent all of the groups known to the NetWorker server. There is no limit on how many groups you can include in a pool. You can also select individual clients for a volume pool without selecting a group. If you do not make a selection from the Groups choices, NetWorker does not sort data based on a group.

- Clients field—entering a client name in this field is optional. You may want to enter the name of a NetWorker client for these reasons:
 - To add one or more clients to a pool without selecting a group.
 - To pick a specific client belonging to a group. If you pick a specific NetWorker client from a group, you exclude all the other clients from that group. If you select a group, the client you select must belong to that group.
- Save sets field—entering a save set in this field is optional. You may want to enter the name of a client save set for these reasons:
 - To send a specific filesystem, from a specific client, to a volume pool without selecting a group (see “Example 3—Secure Device” in this chapter).
 - To send a specific filesystem to a pool (see “Example 5—Specific Filesystem”).
- Levels choices—select the backup level(s) to associate with the media pool. For example, if you click full, the pool will include only full backups for the groups assigned to the pool. See “Selecting the Backup Levels” on page 141 for more information on backup levels.
- Devices choices—displays the devices recognized by NetWorker. This is where you indicate which device to use for the selected pool. For example, you might want one pool for optical disks and one for tapes, if you have both backup devices on the same server. If you do not choose a device, NetWorker uses any available device.
- Store file index entries choices—allows you to choose, with Yes and No choices, whether or not you want to include the index entries from the backup in the online index. If you choose No, the index entries will not be included in the file index. However, an entry for the save set is still made in the media database. If you are using the preconfigured pool Offsite, you may not want the index entries online.

You must select No for the Archive, Archive Clone, Default Clone and PC Clone pools. You must select Yes for the Migration and Migration Clone pools.

Tip: If you select No for Store file index entries, you will not have online file index entries available for browsing in the NetWorker Recover window. To recover data without file index entries, you must use either the Save Set Recover window or the **scanner** command. All parts of the save set must be intact in order for **scanner** to recover the data. See the scanner(1M) reference page for more information.
- Auto media verify choices—click Yes to enable automatic verification of your backup media. NetWorker compares a record near the end of the tape with the data stored on disk. If NetWorker detects a discrepancy, the backup terminates and the status of the offending volume is marked as full.

The savegroup completion message indicates whether or not media errors occurred during the backup.

- **Apply button**—click to apply all the selections you have chosen for a particular pool.
- **Reset button**—click to restore your original selections and override any new selections you may have chosen.

Use the Pools window to choose the preconfigured pools, create new pools, and change existing pools.

Examples

This section contains step-by-step instructions for creating a new pool and several suggestions for why you would want to create your own volume pools. If you are already familiar with the Pools feature, skip to the shortcut directions on page 94.

Example 1—Creating a New Pool

Suppose you want to create a pool named “Test” that selects all of the full backups from a group called “tech pubs.” The technical publications group backs up all of the data from the systems in the quality assurance lab. You want to sort the full backups so you can have a complete set of data on your backup volumes to store in a secured vault. You do not expect that you will need to recover this data, but plan to keep it for a year.

First, create a matching label template, following these steps:

1. Open the Label Templates window by choosing Label Templates from the Customize menu. See the section “Using Label Templates” on page 150 for complete instructions to create a label template.

After you create the Test label template, the Label Templates window looks similar to the image in Figure 4-21.

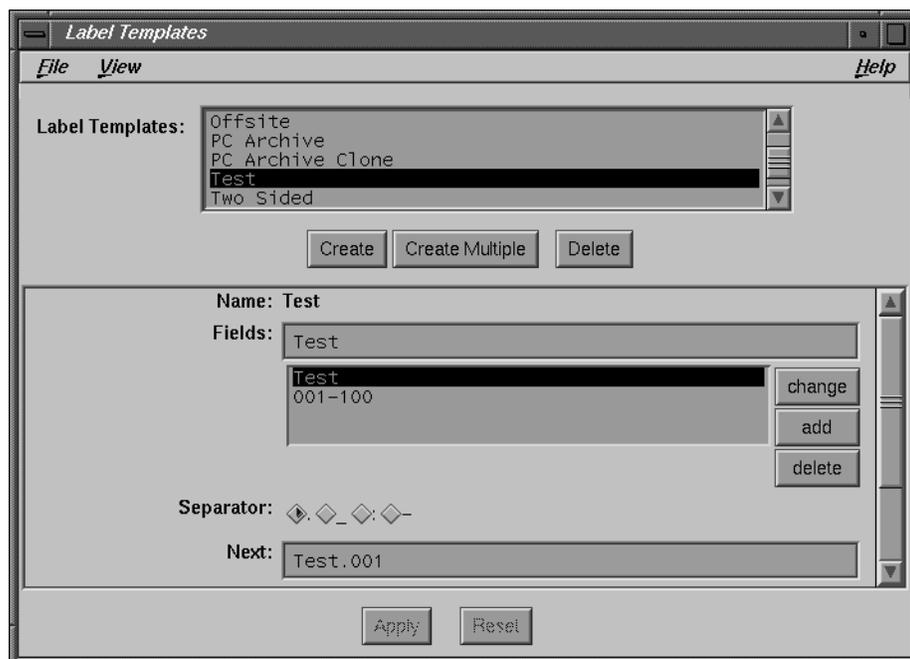


Figure 4-21 Label Templates Window (Test)

After creating the label template, open the Pools window and create a "Test" pool.

2. Click the Create button and enter "Test" in the Name field.

Note: You cannot create a new pool if NetWorker is busy doing a backup—this ensures that NetWorker sends the backup data to the correct pool.

3. Select Yes from the Enabled choices. Select Backup from the Pool type choices. Select the label template Test, which you previously created. Select the group tech pubs, since it contains clients you want to include in the Test pool.
4. Do not enter anything in the Clients or Save sets fields, since you do not need to selectively include any clients or save sets.
5. Select full from the Levels choices.
6. Select a backup device from the Devices choices.
7. Select Yes from the Store index entries choices. Select No from the Auto media verify choices.

- Click the Apply button to apply the selections. After you create the pool Test, the Pools window looks similar to the image in Figure 4-22.

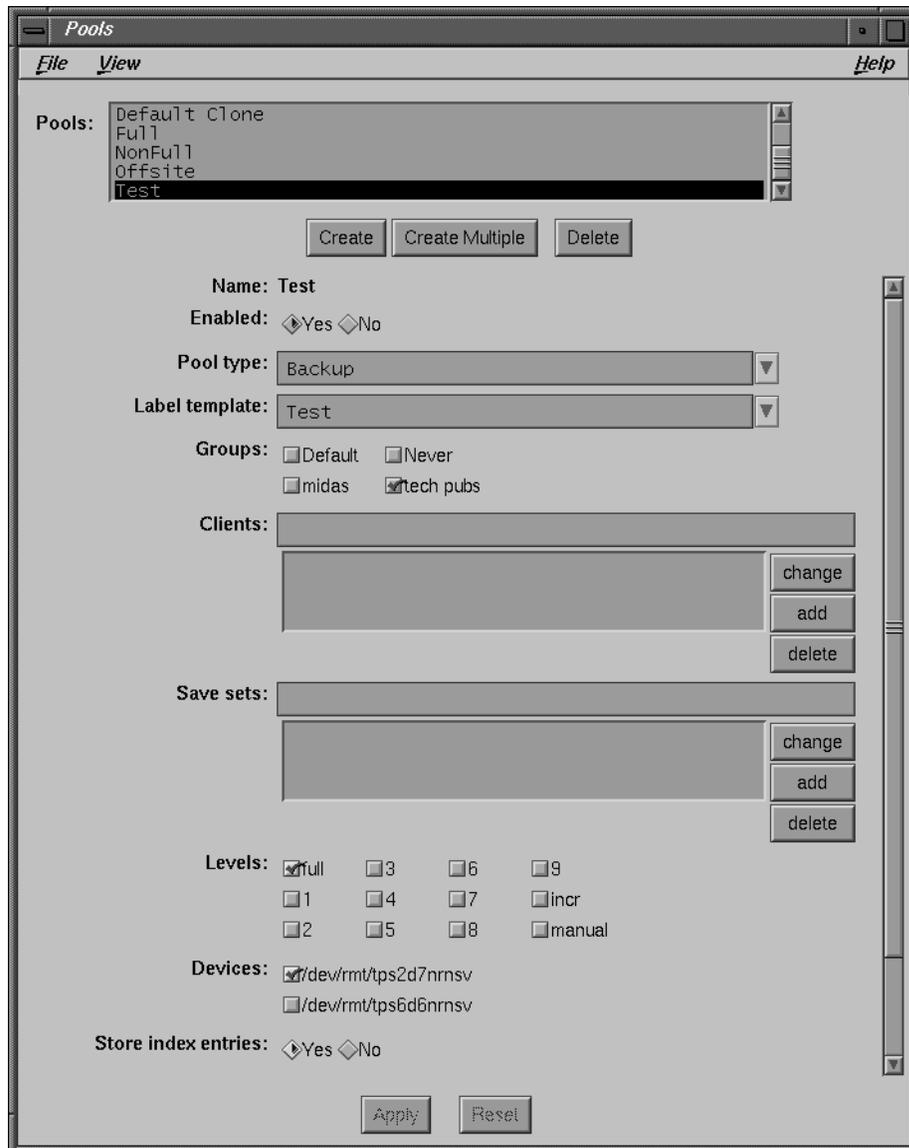


Figure 4-22 Pools Window (Test)

Once you create a new pool, you can easily modify it by choosing the pool from the Pools scrolling list and changing the current selections. When you are done making the modifications, click the Apply button. If you change your mind and do not want to implement the new selections, click the Reset button to restore the Pools window to its original settings.

Example 2—Separate Departments

Your company needs to track department resources for accounting requirements, including the related backup cost for each department. To sort data for each department, create a separate pool for each one. Sorting the data into different pools for each department makes it easy for the Accounting department to determine the backup resources required. These requirements include the number of backup volumes, number of backup devices, and the amount of time necessary for administering each department's backup needs.

To create a pool for each department, create label templates with the name of each department in the Label Templates window. Use the Pools window to create pools with the same names as the label templates, then select the groups for each department.

Example 3—Secure Device

You have confidential company information that needs to be backed up to a secure device in an area that only authorized personnel can access. You do not want to include an entire group, so just a select set of files in a pool named "Confidential."

To create the "Confidential" pool, create a label template and a pool with the name "Confidential." Enter the name of each client in the Clients field, then enter the filesystems for each of the clients in the Save sets field. Finally, select the device that will back up the data from the Devices choices.

Example 4—Targeted Backups

Your company creates large multi-gigabyte graphic files that require multiple volumes to complete a full backup. You want to back up these files to the autochanger on your network, so you do not need to manually load new backup volumes as they fill with data.

First, create a label template named "Graphics Full," create a new pool by the same name, choose the group(s), select full from the Levels choices, and choose the autochanger device from the Devices list. The graphic files will automatically back up to the autochanger.

Example 5—Specific Filesystem

Your company requires that you save electronic mail for all employees. To keep the e-mail files separate from the rest of the backups, create a pool named “E-mail.”

First, create a label template named “E-mail,” then create a pool with the same name. You do not need to select a specific group or client because you will back up electronic mail for all clients. Enter the following in the Save sets field:

```
/usr/spool/mail
```

NetWorker will back up all the electronic mail from all clients to the volumes in the “E-mail” pool.

Shortcut

Use this section if you have reviewed the examples or if you are an experienced NetWorker user.

To create a volume pool, follow these steps:

1. Create a label template using the Label Templates window.
2. Open the Pools window by choosing Pools from the Media menu.
3. Click the Create button and enter the name of the new pool into the Name field. (We recommend matching the name used for the label template.)
4. Select Yes from the Enabled choices.
5. Select the appropriate pool type from the Pool type choices.
6. Choose the matching label template.
7. Make a selection for at least one of the following choices: a group, client, save set, or level. You must make one selection from among these choices or NetWorker will be unable to sort data for the pool.
8. Make the remainder of your choices for Levels, Devices, and Store index entries.
9. Click the Apply button to apply the selections.

Caution: NetWorker does not allow you to create two different pools with duplicate choices for Groups, Clients, Save sets, and Levels. If you try to create two pools with the same choices for all of these items, NetWorker displays an error message.

Deleting a Volume Pool

To delete a pool, follow these steps:

1. Choose Pools from the Media menu to open the Pools dialog box.
2. Select the pool you want to delete in the Pools list box.
3. Click the Delete button. The pool disappears from the Pools list box in the Pools dialog box.

Tip: Give careful consideration to the pools you create. Once you create a pool, you cannot rename it or delete it, as long as it contains backup volumes.

Understanding Volume Pools

NetWorker is designed to back up clients in a group according to a schedule. Data from several clients are mixed together on backup volumes for optimal media usage. NetWorker tracks backed-up data for each client and all the volumes to which their data is backed up.

By contrast, volume pools establish a logical and systematic method for tracking, organizing, and sorting your data and backup volumes. For example, you may want to create a separate pool for each department in your company or a pool that contains only files with confidential information. Or you may want to assign all of your full backups to a particular pool so you can easily store them offsite. Volume pools allow you to select and sort specific files or data to preselected backup volumes.

You cannot delete a volume pool if any backup volumes still belong to that pool. If you want to delete a pool, you must first remove all of the backup volumes from the media database.

Using Label Templates With Volume Pools

Each pool has an accompanying label template, which you create in the Label Templates window before creating the pool. Use the same name for both your label template and volume pool to help track the backups and the pools or tapes to which they belong. If you do not select a label template in the Pools window while creating a pool, NetWorker automatically uses the pool name and creates a template for you. You have more control over the label templates design and content if you create them yourself.

Adding and Deleting Backup Devices

To perform backups, you must enter at least one device *pathname* for a backup device. IRIX NetWorker supports multiple devices, either concurrently as separate standalone drives, or concurrently in an autochanger (with optional jukebox support). Low-end versions of NetWorker support two drives in a series: NetWorker backs up to a volume in the first drive until full, and then backs up to a volume in the other drive.

Use concurrent devices to spread your backups evenly across many devices (up to 16 drives are supported). For unattended backups using a jukebox, purchase the optional Autochanger Software Module.

Note: Depending upon which NetWorker product you purchased, your server may have a limit on the maximum number of devices it supports.

NetWorker supports the following types of backup devices and their corresponding backup media:

- DLT
- half-inch magnetic tape drives
- quarter-inch cartridge tape drives
- 4 mm (DAT) drives
- 8 mm tape drives and 8 mm 5 GB tape drives
- 3480 tape drives
- VHS tape drives
- optical disk drives

Navigating the Devices Window

When you first install the NetWorker software, NetWorker automatically configures the SCSI devices for the backup device choices you make. If, after installing NetWorker, you decide you need to add, delete, or modify your NetWorker backup devices, begin by using the Devices dialog box.

Use the Devices window to view current devices or to add a new device. To open this window, choose Devices from the Media menu. The Devices window appears, as shown in Figure 4-23.

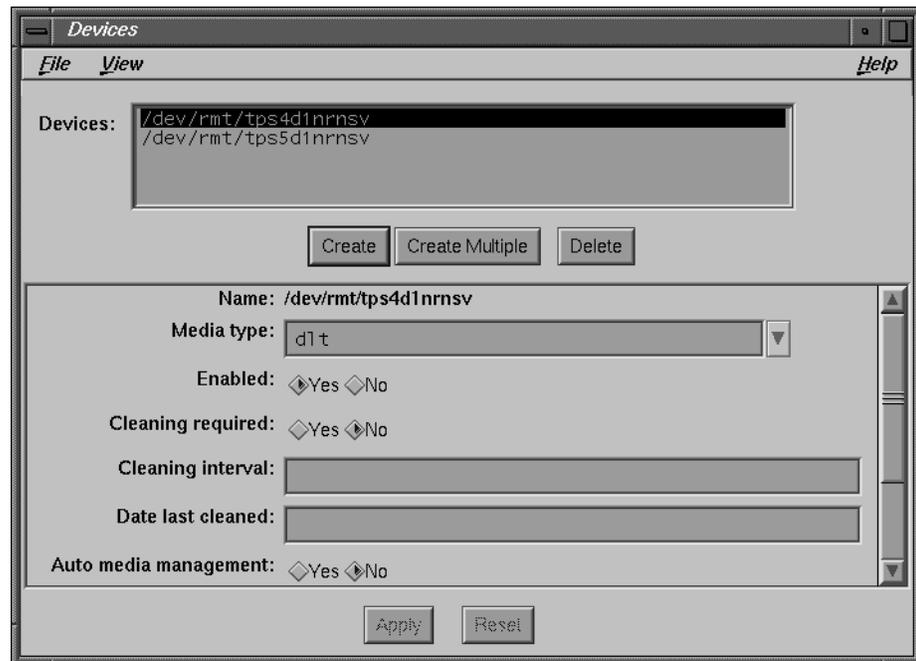


Figure 4-23 Devices Window

The Devices window contains the following items:

- Devices scrolling list—displays device *pathnames* of the backup devices connected to the NetWorker server: either standalone devices or devices located inside of an autochanger. Use these devices to back up and recover client files.
- Name field—displays the *pathname* of the currently selected device and accepts a *pathname* for a new device.
- Media type scrolling list—displays the media type of the currently selected device. You also use the Media type choices to select the type of media for new devices.

Note: You cannot change an existing media type. You must delete and then re-create the selected media type if you want to make changes to it.

Your choices are: half-inch magnetic tape drives (*hmt*), quarter-inch cartridge tape drives (*qic*), four millimeter tape drives (*4mm*), eight millimeter tape drives (*8mm*), eight millimeter five gigabyte tape drives (*8mm 5GB*), 3480 tape drives (*3480*), digital linear tape drives (*dlt*), VHS tape drives (*vhs*), and optical disk drives (*optical*).

The media types displayed in the Devices window vary, depending upon the type of backup devices NetWorker currently supports.

- **Enabled choices**—select **Yes** to enable a backup device and **No** to disable it. NetWorker automatically enables newly added devices. If your backup device is not functioning, select **No** so that NetWorker does not attempt to use the device for backups. When you disable a backup device, it is listed as disabled in the Devices display of the NetWorker Administrator window.

Note: You cannot disable or delete a backup device while a volume is mounted in it.

- **Cleaning required choices**—automatically set to **Yes** when NetWorker determines that an autochanger device needs to be cleaned and set to **No** after the device has been cleaned.

A user schedules future cleaning of a device by setting this attribute to **Yes**. When the value of **Cleaning required** changes from **Yes** to **No**, the value of **Date last cleaned** is automatically set to the current date and time. If required, you can manually enter a value for **Date last cleaned**.

- **Cleaning interval field**—specified number of days, weeks, or months between the **Date last cleaned** and the next scheduled cleaning of the device.

This field accepts any combination of numbers and time periods, with or without a space between: 2day, 2 days; 3week, 3 weeks; 6month, 6 months; and so on.

NetWorker keeps track of the time elapsed since the last cleaning and determines when cleaning is required. A blank value in this field disables the automatic cleaning function.

- **Date last cleaned field**—records the date a device was last cleaned and is updated whenever a device is cleaned.
- **Auto media management field**—gives NetWorker exclusive control over media loaded in the backup device. If you select **Yes**, NetWorker automatically labels, mounts, and overwrites a volume it considers unlabeled. If you select **No**, NetWorker ignores the media and does not consider it for backup.

Note: If the device is in a jukebox, you can enable only the Auto media management feature in the Jukeboxes window. See “Using the Auto Media Management Feature” in this chapter for more information.

Adding a New Backup Device

You must provide the following information to add a new backup device:

- the *pathname* of the device

The *pathname* for storage devices varies depending upon the system vendor. For example, on IRIX servers the *pathname* for a no-rewind tape device is something like */dev/rmt/tps0d3nr*, where **0** indicates the controller and **3** the SCSI address.

- the media type of the device

Caution: A no-rewind backup device is required (for tape drives) because NetWorker writes a file mark on the backup 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.

To add a new backup device to the server, follow these steps:

1. Choose Devices from the Media menu.
2. Click the Create button.
3. Enter a new device *pathname* into the Name field.
4. Select the appropriate media type for the device from the Media type choices.
5. Click the Apply button.

Any additions you make appear in the Devices window.

Deleting a Backup Device

To delete a backup device from the server, follow these steps:

1. Select the device you want to delete in the Devices scrolling list.
2. Click the Delete button. An “ok to delete?” confirmation box appears.
3. Click Ok to delete the device.

Using the Auto Media Management Feature

Auto media management further automates the backup procedure by freeing the administrator from the task of premounting and labeling a volume before backup. When Auto media management is enabled in the Devices dialog box, NetWorker assumes that the backup volume loaded in the device is handled exclusively as NetWorker media.

Backup volumes that appear unlabeled are considered blank and are automatically labeled, mounted, and overwritten with new data.

NetWorker does not recognize the following volume labels and considers these volumes unlabeled:

- volumes without a label
- volumes labeled with something other than a NetWorker label
- volumes with NetWorker labels that were written in a density different than the device in which it is currently loaded. For example, if you performed NetWorker backups on an older 8mm tape drive, a newer 8mm tape drive may not be able to read the volumes from the older device because a different density was used to write the data.

Caution: If your backup devices have different density capabilities, be cautious when sharing NetWorker volumes between devices when Auto media management is enabled. You could potentially relabel and overwrite valuable data.

The Auto media management feature is located in both the Jukeboxes and Devices windows. For autochangers (jukeboxes), you only need to enable Auto media management in the Jukeboxes window. If you attempt to enable Auto media management for an autochanger device in the Devices window, NetWorker displays a warning similar to the following.



Figure 4-24 Notice About Auto Media Management

Note: For more information on managing media automatically with autochangers see Chapter 7, “Autochanger Software Module.”

Recovering and Cloning Save Sets

NetWorker's Save Set Recover feature provides the ability to retrieve data whose file index entries have expired and to quickly restore large amounts of data.

The Save Set Clone feature provides an ideal solution for making identical copies of selected data for storing offsite, sending to another location, or verifying backed-up data.

You can also clone entire NetWorker backup volumes for offsite storage. See "Cloning Backup Volumes" on page 78 for details.

Recovering Save Sets

When an index entry for a file passes its browse time period, the index entry is removed from the online file index to conserve disk space. This means that you cannot recover the data through the NetWorker Recover window. However, as long as the backup media the data was stored on has not been relabeled, you may still be able to recover the data.

To recover data from the backup media when the browse period has expired and the index entry has been removed, use the Save Set Recover window. When you recover by save set, you recover only the data; you do not recover file index entries for the data.

Recovering by save set makes the process of restoring large amounts of data, or data lost due to a disk crash, efficient and straightforward. System administrators who are familiar with the setup of their system disks can easily determine which save sets they need to recover a disk partition or an entire disk.

Tip: Recover data whose entries have been removed from the online file index with the save set recover feature in NetWorker Administrator (**nwadmin** command). Recover data whose index entries are still in the online file index with NetWorker Recover (**nwrecover** command).

Navigating the Save Set Recover Window

This section describes the contents and use of the Save Set Recover window and its associated windows.

To open the Save Set Recover window, choose Recover from the Save Set menu. The Save Set Recover window appears, as shown in Figure 5-1.

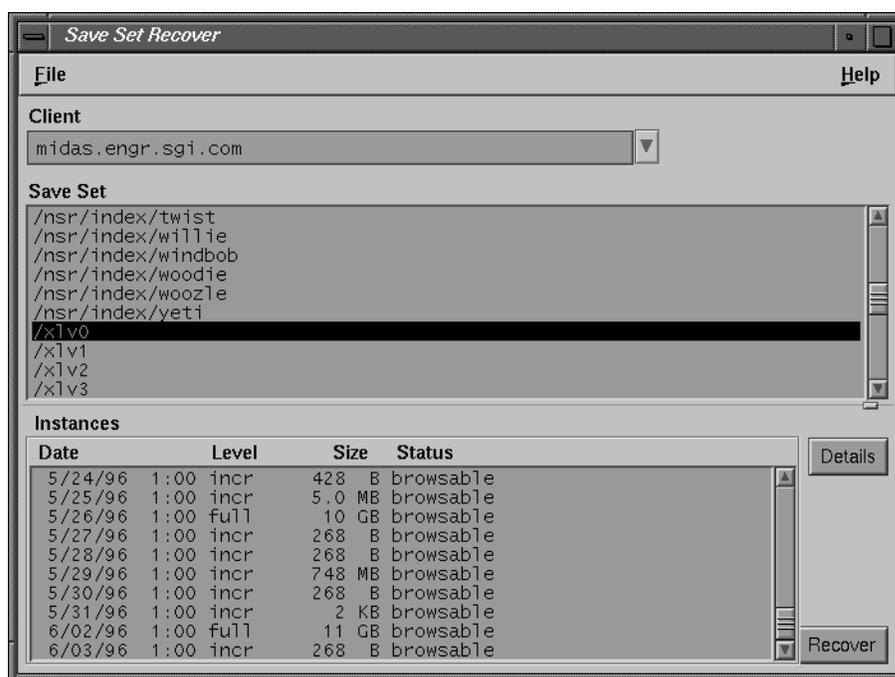


Figure 5-1 Save Set Recover Window

The Save Set Recover window contains the following items:

- Client field—enter the name of the client system for which you want to recover a save set, or use the arrow button to select the client from the Client scrolling list. The list contains the clients known to the current NetWorker server. If you do not enter a client name, NetWorker searches for all client save sets.
- Save Set scrolling list—contains the save sets for the selected client. Click one save set at a time to see its versions in the Instances scrolling list.

- **Instances scrolling list**—contains versions of the selected save set. The list includes the save set **Date**, **Level**, **Size**, and **Status**. The save set status also includes clone information. If the save set has a clone, it is marked **has clones** and the actual cloned copy is marked **cloned** *date time*.

To select a save set, click once. To deselect a save set, click again. The following describes the save set status choices:

- **browsable**—the save set still has an entry in the online file index.
- **recyclable**—all the save sets have passed both the **Browse** and **Retention** policy time periods, and the volume is available for relabeling and overwriting. The data is still recoverable as long as the volume has not been relabeled.
- **scanned-in**—the save set was scanned in using the **scanner** command.
- **recoverable**—the entry for the save set has been removed from the client file index, but the save set is still recoverable from the media.
- **in-progress**—the save set is in the process of being backed up.
- **aborted**—the save set was aborted manually by the administrator during backup or by a system crash.
- **suspect**—a previous recovery attempt might have failed.

You may select an unlimited number of versions of the same save set for recovery at the same time. If you do not choose a save set, NetWorker automatically selects the last save set version from the bottom of the Instances scrolling list.

- **Details button**—click to open the **Save Set Recover Details** dialog box for more detailed information on a specific save set version. Details for the last save set selected will be displayed.
- **Recover button**—click to open the **Save Set Recover Status** dialog box to begin the save set recover process.

Save Set Recover Details Window

The **Save Set Recover Details** window provides detailed information, including the size of the data you are recovering and the name and location of the backup volume.

To open a **Save Set Recover Details** window, click the **Details** button in the **Save Set Recover** window. The **Save Set Recover Details** window appears, as shown in Figure 5-2.

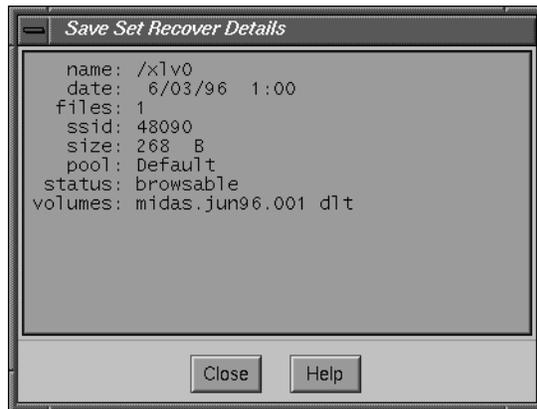


Figure 5-2 Save Set Recover Details Window

The Save Set Recover Details window contains the following items:

- name—the save set name.
- date—day and time the save set was backed up.
- files—number of files in the save set. If the save set is an incremental backup and the files were not changed since the last backup, the number of files is set to zero.
- ssid—internal number assigned to the save set by NetWorker for purposes of tracking the data.
- size—the size of the save set.
- pool—name of the pool to which the backup volume belongs.
- status—the status of the save set.
- volume—backup volume name and its location, if it has one.

Save Set Recover Status Dialog Box

The Save Set Recover Status dialog box allows you to customize save set recovers. It offers you a chance to select individual files rather than entire save sets and allows you to determine where and how to locate recovered data. To open the Save Set Recover Status dialog box, click the Recover button in the Save Set Recover window.

The Save Set Recover Status dialog box appears, as shown in Figure 5-3.

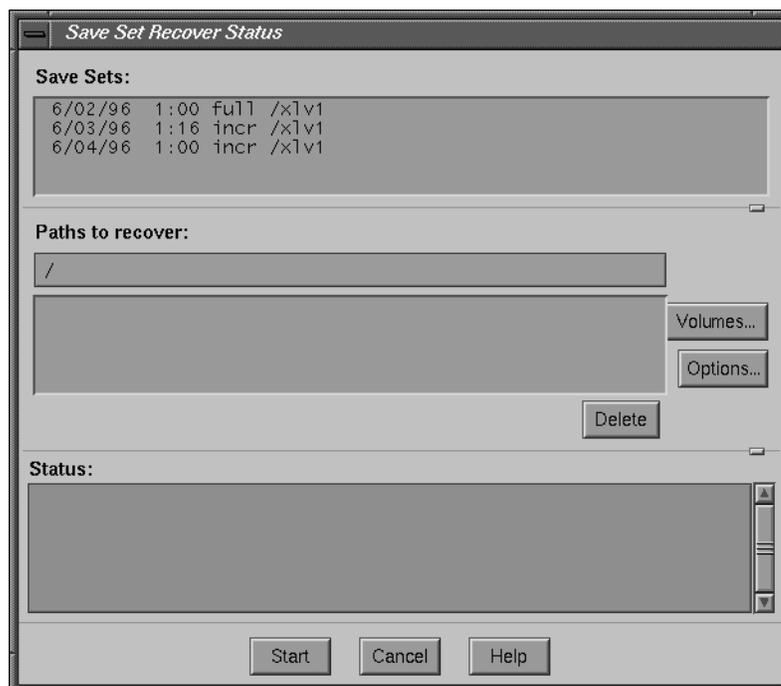


Figure 5-3 Save Set Recover Status Dialog Box

The Save Set Recover Status dialog box contains the following items:

- **Save Sets scrolling list**—contains the save sets selected for recovery. After you open the Save Set Recover Status dialog box, you may continue adding or deleting save sets to the Save Sets scrolling list by selecting and deselecting them from the Instances scrolling list in the Save Set Recover window.
- **Paths to recover field**—recover individual files or directories by entering their complete *pathname*s. Press [Enter] to add them to the Paths to recover scrolling list. Entering an individual file or directory automatically excludes all other data in the save set from being recovered.

If there are no entries in the Paths to recover list when you start the recover, the entire save set is recovered.

Example 1: The save set *pathname* in the Save Sets scrolling list is */nsr/index/cutey*. You need to recover only a single file named *help* from the save set.

To recover the file *help*, enter the following in the Paths to recover scrolling list:

```
/nsr/index/cutey/help
```

Press [Enter] to add *help* to the list. NetWorker recovers only the file *help*.

Example 2: The save set *pathname* in the Save Sets field is */usr/src/pubs*. You want to recover only the directory *Templates* from the save set. To recover the directory *Templates*, enter the following in the Paths to recover field:

```
/usr/src/pubs/Templates
```

Press [Enter] to add *Templates* to the list. NetWorker recovers the *Templates* directory and all of its data, including any directories and files. It does not recover any data above the *Templates* directory level.

- Delete button—click to delete a file or directory selected from the Paths to recover scrolling list.
- Status scrolling list—displays messages that report on the status of the recover. The Messages display of the NetWorker Administrator window displays the same information.
- Start button—click to start a recover.
- Volumes button—click to open the Save Set Recover Volumes Required window, shown in Figure 5-4. This helps determine which backup volumes are needed to recover data. NetWorker displays all the volumes required to recover the save sets listed in the Save Sets scrolling list. This information may be useful, especially if you discover that a backup volume is located offsite and you must make special arrangements to retrieve it.

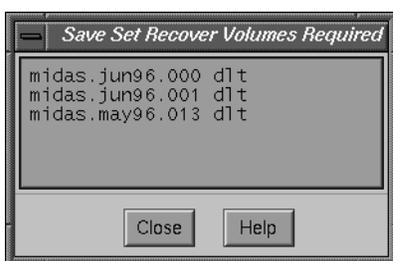


Figure 5-4 Save Set Recover Volumes Required Window

- Options button—click to open the Save Set Recover Options dialog box, shown in Figure 5-5. Use this dialog box to control where and how you will recover data.

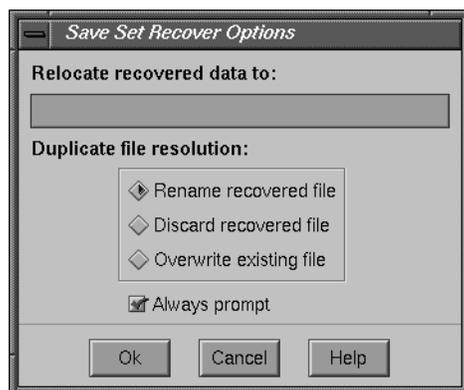


Figure 5-5 Save Set Recover Options Dialog Box

The Save Set Recover Options dialog box contains the following items:

- Relocate recovered data to field—enter the full pathname of the directory where you want to locate the recovered data. If you do not enter a pathname, NetWorker automatically recovers the data to its original location.
- Duplicate file resolution choices—select one of these items to determine what NetWorker should do if a file with the same name already exists. If you do not select one of the choices, NetWorker automatically selects the Rename recovered file choice.
 - Rename recovered file—renames the recovered file with a *.R* extension when a filename conflict occurs. If the *.R* file exists, NetWorker adds another *.R* extension. For example, a file named *document.R* will be renamed to *document.R.R*.
 - Discard recovered file—discards the recovered file. The existing file remains current.
 - Overwrite existing file—replaces the existing file with the recovered file. The existing file is lost and the recovered file becomes current.
- Always prompt toggle button—select for prompts when there is a filename conflict, and turn off if you do not want prompts when there is a filename conflict. If the toggle button is not selected, the Duplicate file resolution choice is applied without confirmation.

Examples

This section provides two examples of how and why you would recover a save set.

Example 1—Recovering Individual Files

In this example, a user needs to recover several files and has not been able to locate them in the NetWorker Recover window. The files have passed their browse period, so they are no longer included in the online file index. Consequently, the files do not appear in the NetWorker Recover window. The user wants to know if the files still exist and whether or not the system administrator can recover them.

Before the search can begin, the user needs to provide the administrator with the name of the *filesystems*, the *filenames*, and the dates of the save sets to be recovered.

To recover the needed files, the system administrator needs to do the following:

1. Open the Save Set Recover window by choosing Recover from the Save Set menu.
2. Enter the name of the client system in the Client field, or click on the arrow button to select a client from the Client scrolling list.

Caution: If the system administrator cannot find the client that needs to have files recovered in the Client scrolling list, it means the client is backed up by another NetWorker server. Open the Change Server dialog box from the Server menu and select the correct server.

3. Select the save set in the Save Sets scrolling list that contains the appropriate filesystem. Once selected, different versions of the save set appear in the Instances scrolling list.
4. Select the save sets with the correct dates from the Instances scrolling list.
5. Click the Details button to see how many files are in the save set. Click Ok when finished with the Save Set Recover Details window.

The system administrator continues the recover process by opening the Save Set Recover Status dialog box.

1. Click the Recover button. The Save Set Recover Status dialog box appears.
2. Enter the entire *pathname* of each file to be recovered in the Paths to recover field. The user does not need to recover the entire save set, only a number of files. Press the [Enter] key to add the file to the Paths to recover scrolling list.

3. Click the Volumes button to open the Save Set Recover Volumes Needed window to determine the backup volumes needed for recovery.
4. Click the Options button to open the Save Set Recover Options dialog box. Enter the full *pathname* of the home directory for the user in the Relocate recovered data to field.
5. Select Rename recovered file or Discard recovered file from the Duplicate file resolution choices to prevent any existing files from being overwritten with recovered files. Click the Always prompt button to be notified every time NetWorker finds an existing file with the same name. Click Ok to apply the choices and to close the dialog box.

If you are uncertain about which files you want, select the Rename recovered file option to get all of the files back during recovery.

6. Begin recovery by clicking the Start button in the Save Set Recovery Status window.
As the recovery process begins, messages that provide information on the progress of the recover appear in the Status display.

Example 2—Recovering a Disk Partition

An engineering company developed software for a federal government project with the hopes of signing a lucrative contract. The project was cancelled because an agreement could not be reached on the terms of the contract. The engineering team was uncertain about when or if the project would resume. Consequently, the engineering team had the company's MIS personnel back up the project files and remove the original data from the disk.

Six months later, an agreement was reached over the terms of the contract. The project has been rescheduled, and the engineering team needs to retrieve the data so it can finish developing the project software.

The MIS personnel need to follow these steps to recover the data back to the original location on the disk partition:

1. Properly prepare the NetWorker client disk partition where the data was originally located.
2. Open the Save Set Recover window by choosing Recover from the Save Set menu.
3. Enter the name of the NetWorker client where the data was originally located in the Client field.

4. Select the save set containing the original project software from the Save Set scrolling list. Different versions of the save set appear in the Instances scrolling list.
Typically a save set is comprised of data from a single disk partition because NetWorker does not automatically cross over mount points when backing up data.
5. Choose the most recent full backup and any incremental backups performed since the last full backup from the Instances scrolling list. The incrementals are necessary because the engineering team needs any changes to the software since the last full backup.
6. Click the Details button for one save set at a time to determine its size and the number of files it contains. The Save Set Recover Details dialog box appears. Click Ok to close the window.

The MIS personnel continue the recovery process by opening the Save Set Recover Status window.

1. Click the Recover button. The Save Set Recover Status window appears. The Save Sets scrolling list contains the save sets selected for recovery. Do not enter a specific file or directory name in the Paths to recover field because the entire save set is needed for recovery.
2. Click the Volumes button to determine the volume(s) needed for recovery. The Volumes Required for Recovery dialog box appears. Click Ok to close the window.
3. Click the Options button. The Save Set Recover Options dialog box appears.
4. Enter the full *pathname* of the disk partition in the Relocate recovered data to field where the files will be recovered. Leave the field blank if you want to recover the files to their original location.
5. Select Overwrite existing file from the Duplicate file resolution choices to overwrite any files that still exist on the original disk partition.
6. Deselect the Always prompt button. There is no need to be notified every time NetWorker finds an existing file with the same name.
7. Click Ok to apply the choices and to close the Save Set Recover Options dialog box.
8. Begin recovery by clicking the Start button in the Save Set Recovery Status window.
As the recovery process begins, messages appear in the Status display that provide information on the progress of the recovery.

Shortcut

Use this section if you are an experienced NetWorker user or after you have reviewed the save set recover examples in this chapter.

Follow these shortcut steps to recover a save set:

1. Choose the Recover command from the Save Set menu.
2. Enter or select the client system whose save sets you want to recover.
3. Select the versions of the save sets you want to recover from the Instances scrolling list. You will probably want the last full and any subsequent incrementals to recover the latest version of the files.
4. Click the Recover button. The Save Set Recover Status dialog box appears. If you need to relocate the files, click the Options button, and enter a pathname in the Relocate recovered data to field, then click Ok. If you want to recover a specific directory or file, enter the complete *pathname* in the File to recover field. If you do not enter a *pathname*, the entire save set is recovered.
5. Click the Start button to begin recovery. Unless you choose to relocate the recovered save set, the data is restored to its original location. You will be notified if any file conflicts occur.

Many variations exist for recovering save sets. The steps above are the minimal number of steps required for recovering a save set.

Understanding the Save Set Recover Feature

Files recovered with the save set recover feature do not include index entries for the file. Consequently, the file index is not repopulated with old entries. This has the distinct advantage of reducing the amount of disk space required to maintain the online file index. NetWorker generates new index entries the next time it backs up the recovered save sets.

Preserving Access Control Lists

NetWorker preserves your operating system access control list when it backs up and recovers files. NetWorker restores the access control list as long as you recover a save set to the same type of filesystem from which it was backed up.

Caution: If you do not recover a save set to the same type of filesystem from which it was backed up, NetWorker does not preserve the access control list.

Recovering Nonbrowsable Data

Save set recovery provides an effective method for retrieving data whose index entries are no longer browsable in the online file index.

All data backed up by NetWorker has both a *browse* and a *retention* policy. The browse policy determines how long index entries for backed-up files remain available in the file index. The retention policy determines how long a file is tracked in the media index and recoverable from the backup media. You can retrieve data from a backup volume as long as it has not been recycled (relabelled).

For more information on browse and retention policies, see “Creating Index Policies” on page 142.

You do not need to keep index entries online any longer than necessary because you can easily retrieve the data directly from the backup volume with the save set recover feature. Index entries take up disk space on your server, so by removing them sooner, you free up more disk space and keep your online indexes from growing too large.

Tip: The **scanner** command also recovers data from media whose entries have been removed from the file index, as long as all parts of the save set are intact and the storage media has not been relabeled. With the **scanner** command you have the choice of scanning in data only, data and index entries, or just index entries. For more information, see the scanner(1M) reference page.

Recovering Large Amounts of Data

Save set recover helps streamline the recovery process for large amounts of data needed to repair a damaged disk or disk partition.

System administrators who are familiar with their system disks can easily determine which save sets they need to recover a disk partition or an entire disk. Typically, a save set is comprised of data from a single disk partition because NetWorker does not automatically cross over mount points while backing up data. Consequently, often only a single save set is required to restore a disk partition, and depending upon the setup of your disk, you may need only two or three save sets to restore the entire disk.

If the index entries are still in the file index you do not need to recover by save set to retrieve data. However, when you are restoring large amounts of data, save set recovery is a more straightforward process. NetWorker simply needs to locate the individual save sets to accomplish the recovery, rather than having to search through and select every single file and directory from the file index.

For more information about the NetWorker Recover program, refer to the *IRIX NetWorker User's Guide*.

Cloning Save Sets

Cloning save sets provides an ideal solution for making identical copies of selected data for storing offsite, sending data to another location, or verifying backed-up data.

Besides cloning individual save sets, NetWorker also clones entire backup volumes. See the section "Cloning Backup Volumes" on page 78.

Cloning data provides additional security for the data on your network by creating identical copies of your backed-up files. You can automatically clone save sets immediately after a backup completes, or manually on an as-needed basis.

The save set cloning feature provides you with a high degree of flexibility for determining which data to copy. You can clone a single save set or configure NetWorker to clone only full backups for an entire group.

The original backup volume is the *source* volume, and the volume you clone the data to is the *destination* volume. The destination volume must be different from the source volume, and it must belong to a clone pool. NetWorker tracks cloned save sets individually and marks them as cloned in the media index. NetWorker indicates whether save sets are "clones" or "has clones" in the Volumes and Save Set Clone windows.

You must have two backup devices to clone data. Use one backup device for the source volume and use the other for the destination volume. If you have an autochanger (jukebox) with two devices, NetWorker mounts the volumes required for cloning automatically. If you have standalone devices, you need to mount the volumes manually.

NetWorker determines which volumes are required, both destination and source, and displays a message in the Messages display to inform you which volumes to mount. If you want to keep the cloned data separate, you may need to premount a blank destination volume.

Cloning save sets requires the following:

- NetWorker server edition, network edition, or power edition
- multiple backup devices
- Active devices value in the Server window set to 2 or more
- source and destination volumes

For information on cloning save sets automatically, see “Automatic Cloning” on page 129.

Navigating the Save Set Clone Window

This section describes the contents of the Save Set Clone window, its associated windows, and how to use them.

Tip: Use the Save Set Clone window to clone save sets manually. Use the Groups window to enable automatic cloning. See “Configuring Backup Groups” on page 159 for more information about automatic cloning.

Save Set Clone Window

Use the Save Set Clone window for browsing and selecting individual save sets for cloning. To open the Save Set Clone window, choose Clone from the Save Set menu, shown in Figure 5-6.

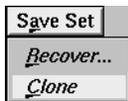


Figure 5-6 Save Set Clone Menu

The Save Set Clone window appears, as shown in Figure 5-7.

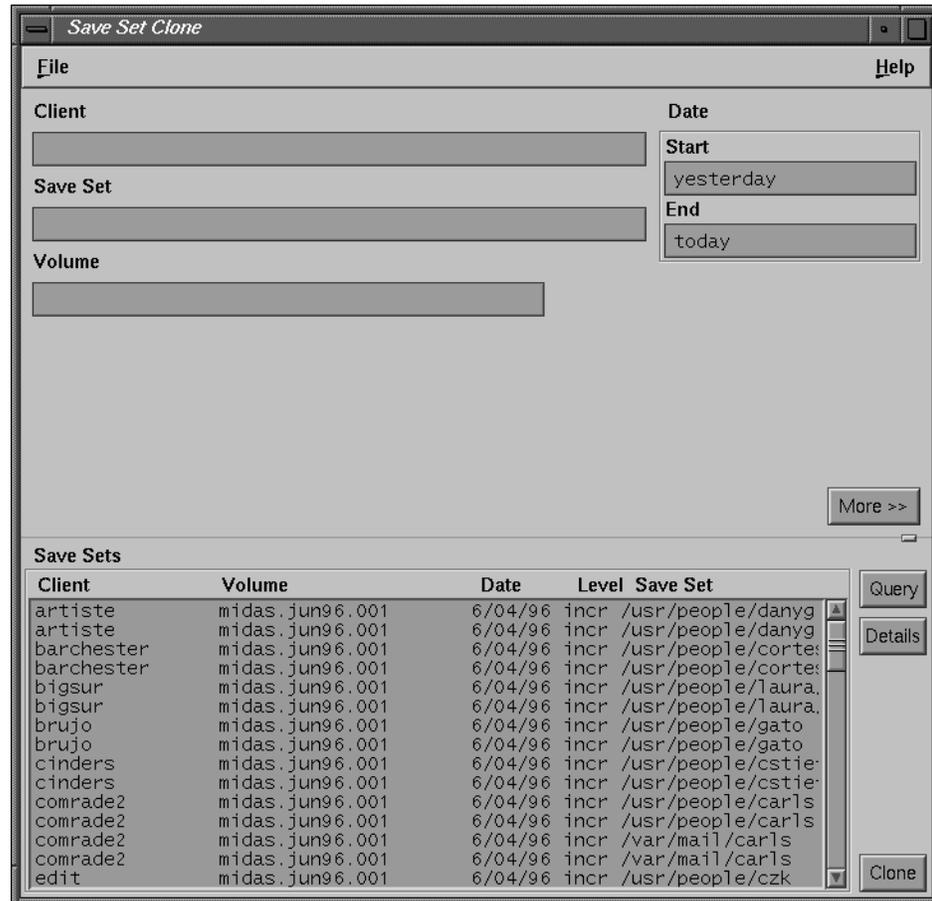


Figure 5-7 Save Set Clone Window

The upper portion of the window contains fields for selecting the save sets you want to browse and select for cloning. You can fine-tune the save sets you select by entering information in as many fields as necessary. The more selection fields you use, the faster NetWorker locates and displays save sets for browsing. If you do not enter a value in the Client field, NetWorker displays all client save sets on the selected server.

When you click the Query button, the lower portion of the window displays the save sets selected for browsing. The lower portion of the window also contains buttons for viewing save set details and for cloning save sets.

The list below describes the portions of the Save Set Clone window used for selecting save sets to browse for cloning:

- Client field—enter the client name whose save sets you want to browse. If you do not enter a client name, NetWorker locates and displays all of the client save sets.
- Save Set field—enter the name of the save set you want to browse.
- Volume field—enter the name of the volume whose save sets you want to browse.
- Date field— enter the range of dates when the save set was created. By default, NetWorker has yesterday in the Start field and today in the End field. Enter the date when you want to begin the search in the Start field and the date when you want to end the search in the End field. You may enter the dates in a variety of ways:
 - completely written out—October 1, 1995
 - numerically—10/1/95
 - with a phrase, for example—two weeks ago

Open the Volumes window to locate information for the Save Set, Volume, and Date fields.

- More button—click for additional browsing choices in the Save Set Clone window. The window changes by adding more choices for selecting criteria for browsing. See “Additional Browsing Choices” on page 119 for more information.

The list below describes the portions of the Save Set Clone window used for selecting and cloning save sets:

- Query button—click to display the save sets chosen for cloning in the Save Sets scrolling list.
Caution: The save sets selected for browsing do not appear in the Save Sets scrolling list until you click the Query button.
- Save Sets scrolling list—displays the save sets currently selected for browsing, including the Client name, Volume name, Date, Level, and Save Set name.
To select a save set for cloning, click once. To deselect a save set, click again.
- Details button—click to open the Save Set Cloning Details window to view more details about a specific save set before cloning.
- Clone button—click to clone the save sets you selected in the Save Set scrolling list. The Save Set Clone Status window appears. You may clone an unlimited number of save sets in one session.

Save Set Clone Details Window

The Save Set Clone Details window contains more detailed information about the save set selected in the Save Sets scrolling list. Select only one save set at a time to view its details. The window displays information for the last save set selected. Information in this window helps prepare for cloning selected save sets. The Save Set Clone Details window is shown in Figure 5-8.

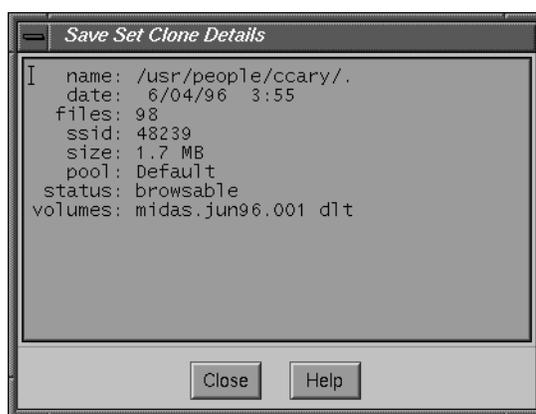


Figure 5-8 Save Set Clone Details Window

The Save Set Clone Details window displays the following items.

- name—save set name.
- date—day and time the save set was backed up.
- files—number of files in the save set. The number of files could be zero if the save set was an incremental backup and the files have not changed since the last backup.
- ssid (save set ID)—internal number that NetWorker assigns to the save set to help track the data.
- size—size of the save set.
- pool—name of the pool to which the backup volume belongs.
- status—status of the save set; also indicates if a save set is a clone or has clones.
- volume—backup volume name and location.

Save Set Clone Status Dialog Box

The Save Set Clone Status dialog box displays the save set selected for cloning. Use this dialog box to enter the name of the clone pool to which the data will be cloned, and then to view the progress of the cloning operation. The Save Set Clone Status dialog box is shown in Figure 5-9.

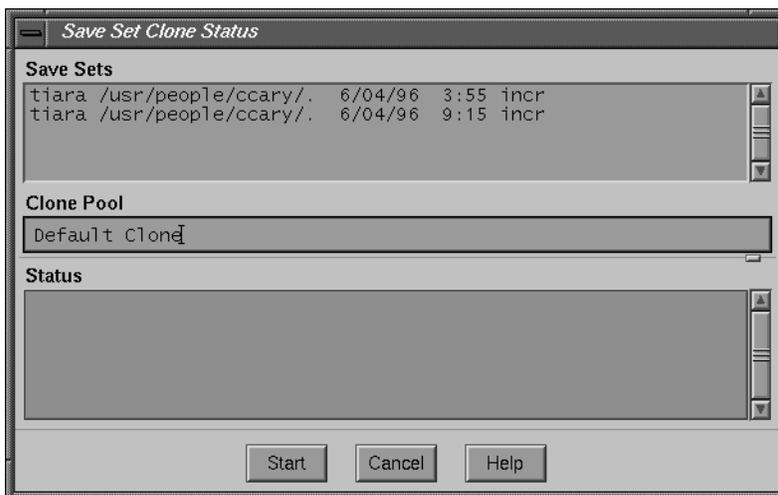


Figure 5-9 Save Set Clone Status Dialog Box

Caution: Volumes used for cloned data must belong to a clone pool. Use either the preconfigured clone pools, or a clone pool you have created yourself. For details, see the section “Using Clone Pools” on page 130.

The Save Set Clone Status dialog box contains the following items:

- Save Sets scrolling list—displays save sets selected for cloning.
- Clone Pool field—enter the name of the clone pool. If you do not specify a clone pool, NetWorker uses the appropriate clone pool for the pool type assigned to the volume. For example, if you are cloning backup volumes and leave Clone Pool blank, NetWorker automatically selects the Default Clone pool.
- Status scrolling list—displays the status of the cloning operation.
- Start button—click to start the cloning operation.

Additional Browsing Choices

To view additional choices for selecting save sets for cloning, click the More button in the Save Set Clone window. The Save Set Clone window changes to display additional fields, as shown in Figure 5-10. To return the window to the previous (simpler) view, click the Less button.

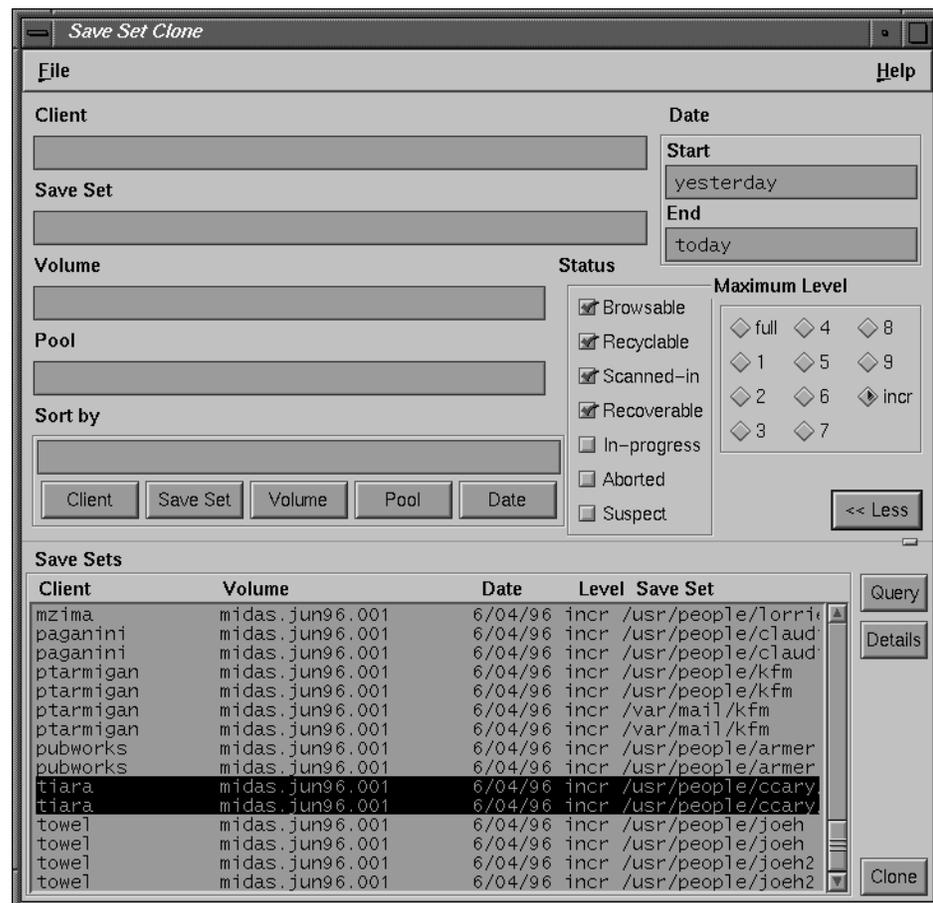


Figure 5-10 Save Set Clone Window (More)

Additional choices for browsing save sets to select for cloning include the following.

- Pool field—enter the name of the pool whose save sets you want to browse.

- **Status choices**—choose the status of the save sets you want to browse. By default, **Browsable**, **Recyclable**, **Scanned-in**, and **Recoverable** are already chosen. The following list describes the save set status choices:
 - **Browsable**—the save set still has an entry in the online file index.
 - **Recyclable**—all the save sets have passed both the **Browse** and **Retention** policy time periods, and the volume is available for relabeling and overwriting. The data is still recoverable as long as the volume has not been relabeled.
 - **Scanned-in**—the save set was scanned in using the **scanner** command.
 - **Recoverable**—the entry for the save set has been removed from the client file index, but the save set is still recoverable from the media.
 - **In-progress**—the save set is in the process of being backed up.
 - **Aborted**—the save set was aborted manually by the administrator during backup or by a system crash.
 - **Suspect**—a previous recovery attempt might have failed.

The Status choices, by default, already have **Browsable**, **Recyclable**, **Scanned-in**, and **Recoverable** chosen for selecting save sets. If you do not select any Status choices, NetWorker selects all of the client save sets for browsing. You can only browse **In-progress** and **Aborted** save sets; you cannot clone them.

- **Maximum Level choices**—choose the levels of the save sets you want to browse. By default, **incr** is chosen, which searches for all levels of backups. NetWorker searches for and displays all the levels dependent upon your level choice. For example, if you select **Maximum Level 5**, NetWorker will browse and display the following save set levels: **full**, **one**, **two**, **three**, **four**, and **five**. If you select **full**, NetWorker searches for and displays only level **fulls**. You must select at least one level; however, you may not select more than one level at a time.
- **Sort by field**—click one or more of these buttons (shown in Figure 5-11) to sort the save sets listed in the **Save Sets** scrolling list. NetWorker displays the selections you made in the **Sort by field** and sorts the save sets in that order.



Figure 5-11 Sort By Field for Save Set Clone

Examples

This section contains three different examples of how and why you would want to clone save sets.

Example 1—Manually Cloning Save Sets

In this example, a user asked the system administrator to manually clone several save sets that are not included in an automatic cloning schedule.

The user wants to send the data to another company location out of state that also uses NetWorker to back up network-wide data. To accomplish this task, the system administrator must clone the most recent full backup and any incrementals since that last full backup to make sure the most current data is sent.

Note: An offsite location will not have the file and media indexes for the cloned data. Therefore, the company location receiving the cloned data must use the **scanner** command to retrieve the data. See the scanner(1M) reference page for information about **scanner** command line options.

To clone the save set for the user, the administrator must have the following information:

- NetWorker client name
- name of the filesystem or directory
- date of the backup to be cloned

To clone the individual save sets, the system administrator follows these steps:

1. Open the Save Set Clone dialog box by choosing Clone from the Save Set menu.
2. Enter the following: client name in the Client field, filesystem name in the Save Set field, and starting and ending dates to search for in the Date field.
3. Accept the default Status choices: Browsable, Recoverable, Recyclable, and Scanned-in.

After choosing the browsing criteria, the system administrator sorts, views details, and selects the save sets for cloning, following these steps:

1. Click the Query button; the Save Sets scrolling list displays all save sets that meet the criteria entered in the upper portion of the dialog box.
2. Sort the save sets by date by clicking the Date button in the Sort by field to make sure that the correct dates have been selected.

3. Select the save set versions for cloning.
4. Click the Details button to determine the size of the data and which backup volume contains the save sets.
5. Click the Clone button. The Save Set Status Clone dialog box appears.
6. Enter the name of the clone pool to which the backup volumes will belong. Click the Clone button in the Save Set Status Clone dialog box. The Status field displays the progress of the cloning process.

Using the Save Set Clone window is useful for occasionally cloning individual save sets or for regularly cloning a small number of save sets.

Example 2—Automatically Cloning One Client

In this example, a system administrator wants to clone data automatically on a daily basis for a new system, *Shakespeare*, that manages electronic documentation for the company. *Shakespeare* performs an incremental backup every night of the week, except on Friday, when it performs a full backup.

To automatically clone backups for *Shakespeare*, the system administrator must accomplish the following tasks:

- Create a group named *Documents* to which the client *Shakespeare* will belong.
- Make the appropriate selections for the client *Shakespeare*, which includes assigning it to the group *Documents*.

To automatically clone data for *Shakespeare*, the system administrator follows these steps:

1. Open the Groups window by choosing Groups from the Customize menu.
2. Create the group *Documents*. See “Configuring Backup Groups” on page 159 for more information.
3. Select the appropriate choices in the Groups window. The following are critical choices:
 - Autostart choices—Enabled
 - Clones choices—Yes
 - Clone pool field—Default Clone
4. Open the Clients window by choosing Client Setup from the Clients menu.

5. Select the appropriate choices for *Shakespeare*. The following are critical choices:
 - Schedule field—Full Every Friday
 - Group choices—Documents

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

Example 3—Automatically Cloning Several Clients

This example illustrates the level of control you can achieve cloning data with NetWorker. You can create an unlimited number of automatic cloning configurations with the aid of the Groups, Schedules, and Clients windows.

In this example, the system administrator wants to automatically clone the data for a small group of NetWorker clients that have valuable accounting information. The administrator wants to clone only full backups to store in an offsite vault. Two different groups and schedules need to be created and applied to two identical sets of individual clients in the Clients window.

Below 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. Create one schedule named *Acct Incr* and another named *Acct Full Clone*. By creating two schedules and selecting them for the appropriate set of clients, you avoid cloning all of the incremental backups.

Table 5-1 Sample Acct Incremental Schedule

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

Table 5-2 Sample Acct Full Clone Schedule

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

- Create two groups: one group named *Acct Incr* and another named *Acct Full Clone*.

Note: Cloning is applied to a group. Consequently, two groups are necessary: the *Acct Full Clone* group for performing full backups and clones, and the *Acct Incr* group for performing incremental backups.

- Add each client name twice to the Clients window. Two names are needed because only one schedule can be applied to a NetWorker client at one time. One entry for the client name uses the *Acct Incr* schedule, and the other entry uses the *Acct Full Clone* schedule.
- Select the appropriate group and schedule for each set of clients, using the Clients window.

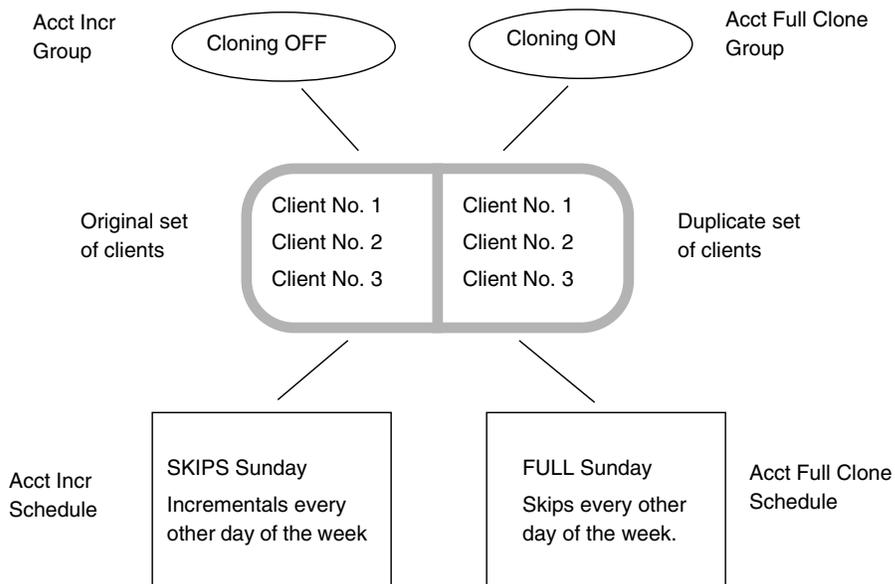


Figure 5-12 Automatic Cloning Scheme

Automatically clone the data by following these steps:

1. Open the Schedules window by choosing Schedules from the Customize menu. (For complete instructions on creating schedules, see “Setting Up Backup Schedules” on page 131.)

The *Acct Incr* schedule (see Figure 5-13) performs incremental backups every day of the week except Sunday, when it skips the backup.

Name:

Period: Week Month

September 1996							October 1996						
S	M	Tu	W	Th	F	S	S	M	Tu	W	Th	F	S
1/s	2/i	3/i	4/i	5/i	6/i	7/i			1/i	2/i	3/i	4/i	5/i
8/s	9/i	10/i	11/i	12/i	13/i	14/i	6/s	7/i	8/i	9/i	10/i	11/i	12/i
15/s	16/i	17/i	18/i	19/i	20/i	21/i	13/s	14/i	15/i	16/i	17/i	18/i	19/i
22/s	23/i	24/i	25/i	26/i	27/i	28/i	20/s	21/i	22/i	23/i	24/i	25/i	26/i
29/s	30/i						27/s	28/i	29/i	30/i	31/i		

Figure 5-13 Acct Incr Schedule

The *Acct Full Clone* schedule (see Figure 5-14) skips a backup six days of the week, but performs a full backup and clones the data automatically on Sunday.

Name:

Period: Week Month

September 1996							October 1996						
S	M	Tu	W	Th	F	S	S	M	Tu	W	Th	F	S
1/f	2/s	3/s	4/s	5/s	6/s	7/s			1/s	2/s	3/s	4/s	5/s
8/f	9/s	10/s	11/s	12/s	13/s	14/s	6/f	7/s	8/s	9/s	10/s	11/s	12/s
15/f	16/s	17/s	18/s	19/s	20/s	21/s	13/f	14/s	15/s	16/s	17/s	18/s	19/s
22/f	23/s	24/s	25/s	26/s	27/s	28/s	20/f	21/s	22/s	23/s	24/s	25/s	26/s
29/f	30/s						27/f	28/s	29/s	30/s	31/s		

Figure 5-14 Acct Full Clone Schedule

- After creating the two schedules, choose Groups from the Customize menu to open the Groups window.
- Create the first group, *Acct Incr*. For detailed instructions on creating a group see "Creating a New Group" on page 161.

Click the buttons for the following critical choices:

- Autostart choices—Enabled
- Clones choices—No

The window with the *Acct Incr* group looks similar to Figure 5-15.

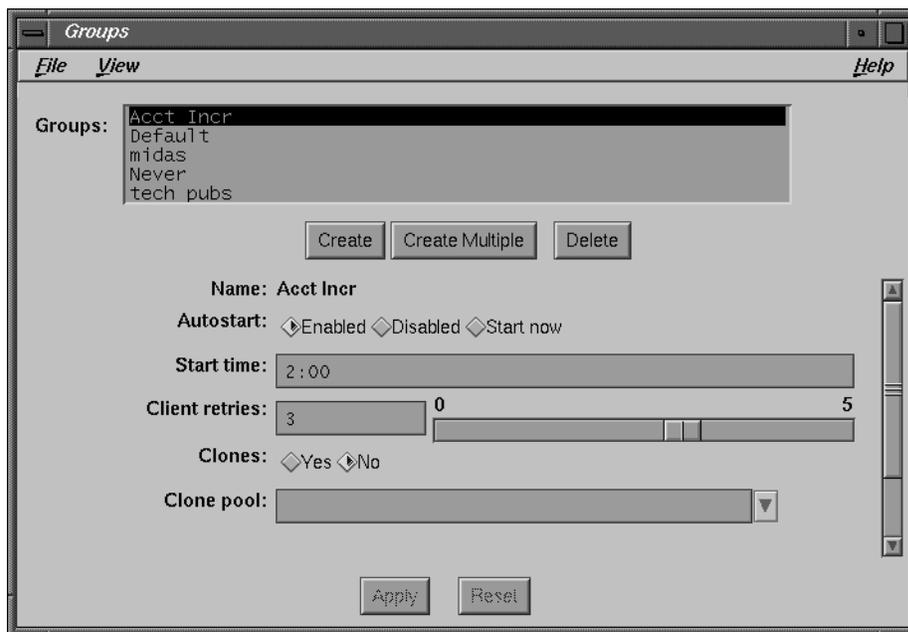


Figure 5-15 Acct Incr Groups Window

4. Create the second group, *Acct Full Clone*. Click the buttons for the following critical choices:
 - Autostart choices—Enabled
 - Clones choices—Yes
 - Clone pool field—Default Clone

Tip: Using two different entries for Start time for the backups helps reduce network traffic.

The *Acct Full Clone* group looks similar to Figure 5-16.

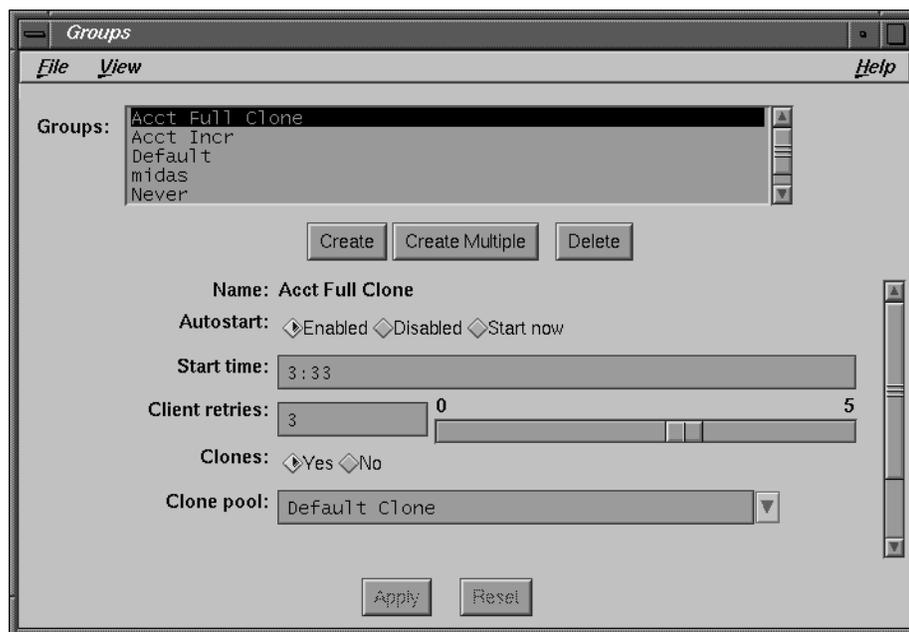


Figure 5-16 Acct Full Clone Groups Window

Next, you create two sets of identical clients and select their schedules and groups:

1. Open the Clients window by choosing Client Setup from the Clients menu. (For complete instructions, see “Adding a New Client” on page 30.)
2. Select the first set of clients, one at a time, and assign them to the Acct Incr group. Make the other appropriate selections in the Clients window. Click the buttons for the following critical choices:
 - Schedule scrolling list—Acct Incr
 - Group choices—Acct Incr
3. Click the Apply button after making the choices for each client.
4. Select the duplicate set of clients, one at a time, and assign them to the Acct Full Clone group. Make the other appropriate selections in the Clients window. Click the buttons for the following critical choices:
 - Schedule scrolling list—Acct Full Clone
 - Group choices—Acct Full Clone

5. Click the Apply button after making the choices for each client.

The previous set of instructions is just one example of how a system administrator might fulfill automatic cloning requirements for clients.

Shortcut

Use this section if you are an experienced NetWorker user or after you review the example of manually cloning save sets.

To clone a save set, follow these steps:

1. Choose Clone from the Save Set menu. The Save Set Clone window appears.
2. Enter or select the criteria you want to use for browsing save sets to select for cloning.
3. Click the Query button; the Save Set scrolling list displays all of the save sets that meet the criteria you entered.
4. Select the save sets you want to clone from the Save Set scrolling list.
5. Click the Clone button. The Save Set Clone Status dialog box appears.
6. Click the Start button in the Save Set Clone Status dialog box.

Understanding the Cloning Save Sets Feature

Cloning save sets allows you to create identical copies of backed-up data quickly and easily. You may want identical copies of your data for storing offsite, or for sharing valuable engineering data with another company location.

NetWorker assures the safety of your data by copying clones of the save set to a different backup volume. Source and destination media can be of different types and capacities.

Cloned save sets span the destination volumes just as they span source volumes. Save sets that fit on one backup volume may spill over to two backup volumes, simply because the destination volume is a different size or has a different capacity. The opposite is also possible, where the original save sets required two volumes and the cloned save sets require only one volume.

By default, save sets on the source and destination volumes have the same *retention* policies. However, you can override the preselected retention policy by applying the *man* mode to the cloned volume in the Volumes window. When you apply the *man* mode to a volume, it means that you have to change the volume manually to the recyclable mode before NetWorker will overwrite the backup volume.

To help minimize network traffic, NetWorker clones save sets after a backup is complete. Allow time to complete the cloning process before scheduling another backup. By providing adequate time to complete either backing up or cloning of data, you help reduce network traffic.

NetWorker creates index entries for the cloned save sets and volumes in the online media index (Volumes scrolling list). However, both the original and cloned volumes use the same set of entries in the online file index. If you eliminate either the cloned or original volume from the media index, NetWorker will not let you remove the index entries from the file index as long as one of the volumes still exists.

Cloning data helps verify the integrity of backed-up data. You receive a message telling you which save sets could not be read, if any, after cloning has completed. If NetWorker performed an automatic clone, you receive a savegroup completion notification by e-mail. If NetWorker performed a manual clone, you receive a message in the Status portion of the Save Set Clone Status dialog box.

Typically, it takes the same amount of time to back up and clone data. The cloning process may take less time, simply because the tape drive does not have to wait for the client disk to send data. CPU speed and the operations taking place on the NetWorker server affect the time required for cloning data. However, generally it is the backup device that controls the speed with which data is backed up or cloned.

Tip: For example, cloning data to a 5 GB tape on an 8 mm tape drive, backing up at a rate of 500 KB per second, should take approximately 2.75 hours.

Automatic Cloning

Use the Schedules and Groups windows to determine the number and kinds of save sets to clone automatically.

The most important thing to remember about automatic cloning is that it is enabled and set up using the Groups window. Select Yes to clone data for a group, and enter or select the name of the Clone pool.

Follow these suggestions for setting up your groups, with the aid of different backup schedules, to accomplish automatic cloning:

- Set up a group with only those clients whose data you want to clone every day of the week.
- Set up a group whose clients perform level full backups and clones one day a week and another group whose clients perform incremental backups every other day of the week.
- Set up a group whose clients back up and clone specific save sets every day of the week.

NetWorker allows you to create a variety of backup schedules and groups that help in setting up automatic cloning. You can easily accomplish automatic cloning of selected save sets by following the suggestions above, or by developing your own configurations.

Using Clone Pools

The volumes chosen for cloning must belong to a clone pool. Assigning cloned volumes and nonclone volumes to separate pools helps you manage and identify cloned and noncloned data.

NetWorker provides the following preconfigured clone pools: Archive Clone, Default Clone, Migration Clone, and PC Archive Clone. These clone pools are listed as choices in the Pool type scrolling list in the Pools window. The Pools window also allows you to create your own clone pools if you do not want to use the preconfigured clone pools. Use descriptive names for the clone pools you create to help identify the data they contain.

For complete instructions on creating pools, see “Using Volume Pools” on page 85.

Customizing NetWorker Backups

This chapter covers the options available for customizing your NetWorker backups. By using NetWorker's advanced configuration capabilities, you have the opportunity to tailor your backups to meet the specific needs of your network. The commands for configuring your NetWorker backups are located in the Customize menu.

NetWorker provides preconfigured settings so you can immediately back up your NetWorker clients. Customizing backups is not required, but is provided as an option so you can create your own unique configurations. See "Preconfigured Selections" on page 246 for information on preconfigured settings.

The following are tasks for customizing NetWorker:

- setting up backup schedules
- creating index policies
- configuring backup groups
- using directives and label templates
- setting up event notification

Setting Up Backup Schedules

The individual backup schedule chosen for each client determines the level and schedule of the backup. The information in this section describes how to create your own customized backup schedule in the Schedules window, which you then apply to individual clients in the Clients window.

NetWorker provides five preconfigured schedules: Default, Full on 1st of Month, Quarterly, Full on 1st Friday of Month, and Full Every Friday. They are described under "Preconfigured Backup Schedules" on page 246. You can make changes to the existing preconfigured backup schedules, but you cannot change their *names*. If you want to change the name of a schedule, you must delete it and re-create it, using a new schedule name.

Navigating the Schedules Window

Create new backup schedules for your NetWorker clients in the Schedules window. To open the Schedules window, choose Schedules from the Customize menu. The Schedules window appears, as shown in Figure 6-1.

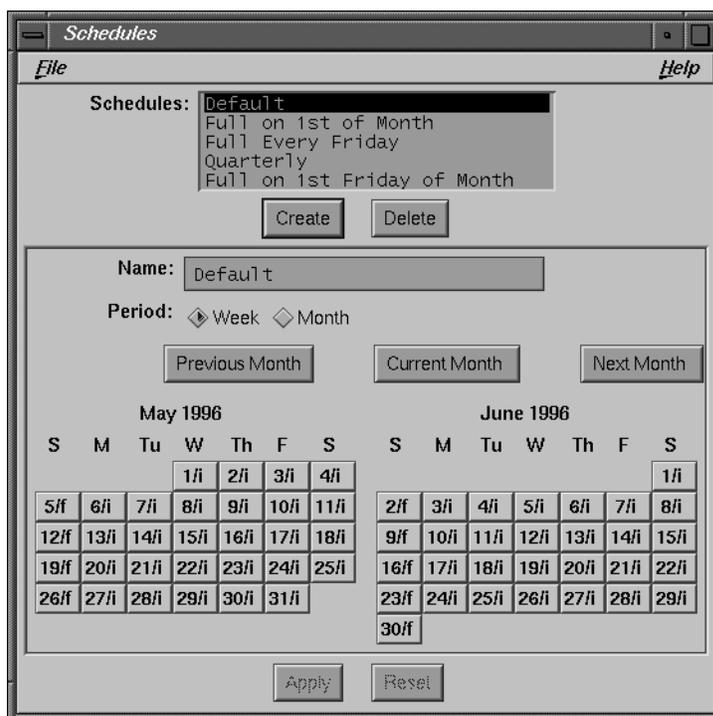


Figure 6-1 Schedules Window

The Schedules window contains the following items:

- Schedules scrolling list—displays the names of the backup schedules known to the NetWorker server, including the preconfigured schedules provided with NetWorker and any new schedules that you create.
- Name field—displays the name of the currently selected schedule. Use the Name field to enter the name of a new schedule after you click the Create button.
- Period choices—let you set the schedule to apply to a weekly or monthly period. Week is the default selection.

When you select **Week** and set up a schedule, the backup level you select is applied to that day of the week for all weeks in the calendar (for example, level fulls every Sunday).

When you select **Month** and set up a schedule, the backup level you select is applied to that day of the month for all months in the calendar (for example, level fulls on the fifteenth of each month).

- **Previous Month** button—click to view the backup schedule for the previous month in the calendar.
- **Current Month** button—click to view the backup schedule for this month and next month in the calendar.
- **Next Month** button—click to view the backup schedule for the upcoming month in the calendar.

Note: Each time you click the **Previous Month** button, NetWorker displays the previous two months in the calendar year; each time you click the **Next Month** button, NetWorker displays the next two months in the calendar year.

- **calendars**—display the day of the month and the backup level scheduled for each day. The date on each calendar day is followed by a forward slash (/) and one of the following numbers or characters indicating the level and type of scheduled backup for that day:
 - **f**—full backup.
 - **1-9**—level one through nine backups.
 - **i**—incremental backup.
 - **s**—skip scheduled backup.

An asterisk (*) next to a backup level means you have forced a different level (an override) for that day. For example, you scheduled a full backup on every Monday of the month, but the second Monday is a holiday. You can force a “skip backup” for the second Monday and force a full backup on Tuesday instead.

Before setting up a backup schedule, decide which schedule best fits the needs of your network and learn about the backup levels NetWorker supports.

Examples

This section illustrates setting up backup schedules with two examples. Read through the examples before you try to set up any new schedules on your own.

To create a new schedule, name the schedule and select your choices in the Schedules window. Then apply the schedule to a client in the Clients window.

Example 1—A Weekly Schedule

This schedule is suitable for NetWorker clients with files that change often. You want to do a full backup every Monday because a weekly full backup simplifies the disaster recovery process. You also want to do a full backup on Mondays because someone will be available to change the media and monitor the server. The other days of the week, incremental backups protect the files that have changed since the last backup. You are going to name the backup schedule “Monday Full” to remind yourself that this schedule does a full backup every week, on Mondays.

To set up the “Monday Full” schedule, follow these steps:

1. Open the Schedules window by choosing Schedules from the Customize menu.
2. Click the Create button.
3. Enter “Monday Full” in the Name field.
4. Click the Week button, if it is not already selected.
5. Select any Sunday in the calendar for the either month and click the primary mouse button to pop up the menu of backup levels, as shown in Figure 6-2.

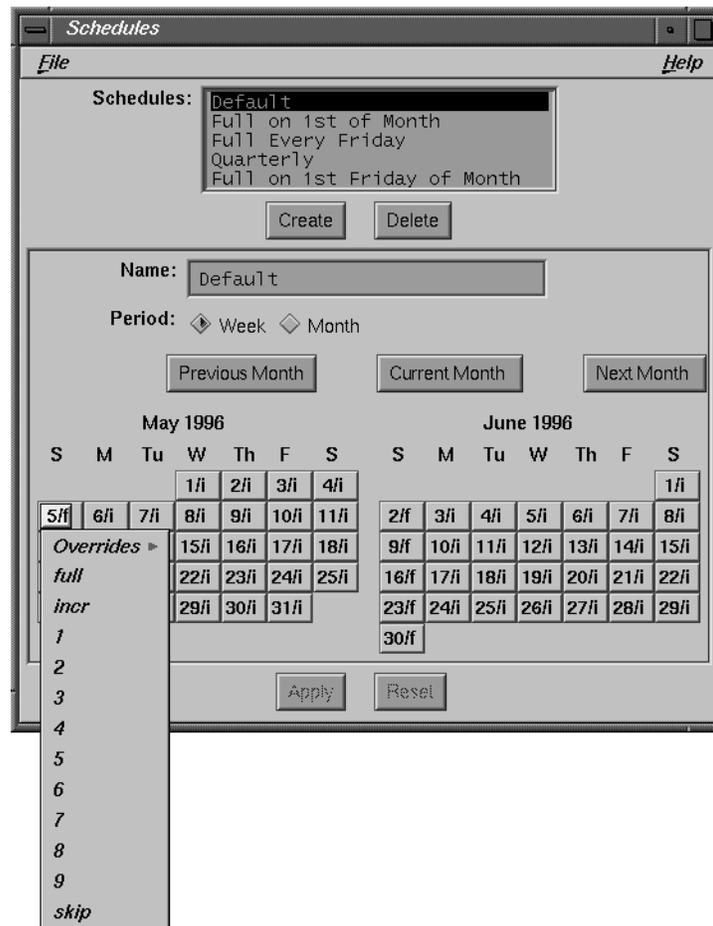


Figure 6-2 Schedules Window, Backup Level Menu

6. Choose *incr* from the backup level popup menu.

Notice all the Sundays of the calendar change to a backup level *i*, which signifies incremental backup.

Next, set the calendar to do a full backup every Monday:

1. Select any Monday in the calendar for the current month and pop up the backup level menu.
2. Select full from the backup level popup menu.
3. Click Apply to save the schedule.

Note: If you see the error message “user *user_name* needs to be on administrator’s list,” it means you do not have permission to make changes to NetWorker configurations.

See “Adding or Changing Administrators” in Chapter 4 for more information.

All the Mondays in both calendars display “f” next to the date, indicating a full backup takes place instead of the incremental backup. When you scroll through all the months, you see how the schedule set up for the current month is maintained throughout the entire calendar.

Example 2—A Monthly Schedule

Use a monthly schedule for clients with files that do not change often. These clients need a full backup performed only once a month.

In this schedule, a full backup takes place on the first day of the month. A level 5 backup takes place in the middle of the month, backing up all files that have changed since the full backup. All other days of the month, an incremental backup protects daily changes. Name this schedule “Monthly Full” to remind you that this schedule does a full backup only once a month. Skip the level 5 backup with an override on November 16 because it is a company holiday; instead force the backup on November 17.

To set up the “Monthly Full” schedule, follow these steps:

1. Choose Schedules from the Customize menu to open the Schedules window.
2. Click the Create button.
3. Enter “Monthly Full” in the Name field.
4. Click the Month button, if it is not already selected.

Notice the months are already set up to do a full backup on the first day.

Note: When you set up one month in the Schedules window, that setup applies to all calendar months, and carries over from year to year. In the same way, if you set up a weekly schedule, it applies for all calendar weeks. This allows you to plan all the weeks or months in one session, without having to repeat the process every week or month.

Set the month for a level 5 backup on the sixteenth day:

1. Select the sixteenth day (16/i) of the current month.
2. Select "5" from the backup level popup menu.

Scroll through the months, and notice how the monthly schedule is carried through every month—a full backup on the first day, a level 5 on the sixteenth day, and incremental backups on the other days of the month.

Next, override the scheduled level 5 backup on November 16 with a level "skip," to skip the backup for that day, and schedule the level 5 backup for November 17:

1. Use the Previous Month or Next Month buttons to display the November calendar.
2. Select the sixteenth day in the November calendar.
3. Pop up the backup level menu, then select Overrides.
4. Slide the cursor to the right to display the Overrides menu.

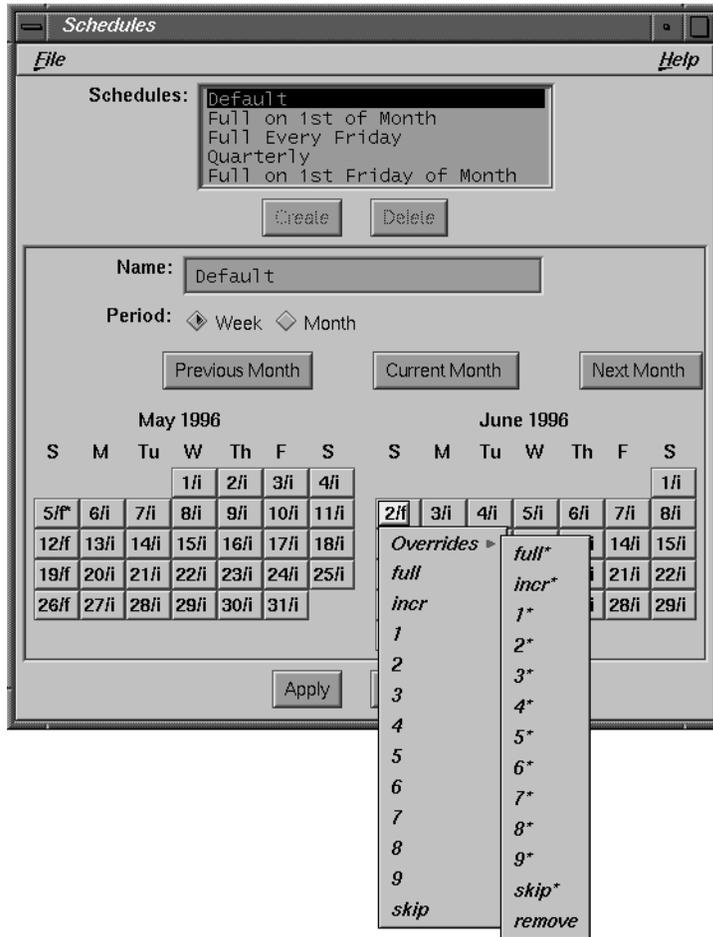


Figure 6-3 Schedules Window, Level Override Menu

5. Select skip* from the Overrides menu.
6. Select the seventeenth day in the November calendar.
7. Select 5* from the Overrides menu.
8. Click Apply to save the schedule.

The date on November 16 displays an s*, indicating you have forced a “skip backup” instruction for that day. The date on November 17 displays a 5* to indicate a level 5 backup will take place instead of the incremental. Overrides must be scheduled explicitly each time. An override does not automatically repeat within a monthly or weekly period.

Shortcut

Use this shortcut section if you are an experienced NetWorker user, or you have reviewed the examples preceding this section.

To create a backup schedule, follow these steps:

1. Choose Schedules from the Customize menu to open the Schedules window.
2. Click the Create button in the Schedules window.
3. Enter a name for the new schedule in the Name field.
4. Choose a weekly or monthly period for the schedule.
5. Place the cursor on a calendar day, and select it to pop up the backup level menu. Set the backup level for each day (or week, or month) by selecting one of the following choices:
 - full for a full backup
 - incr for an incremental backup
 - skip to skip a backup
 - any of the numbers 1 through 9 for a level backup
 - Overrides from the backup level menu to bring up the Overrides menu. Choose an override backup level to apply to the selected calendar date.
 - remove from the Overrides menu to remove an existing override
6. Click the Apply button to create the new schedule when you are through setting backup levels.

You may set up as many NetWorker backup schedules as you like, as long as you give each one a unique name. To assign a backup schedule to a client, you must use the Clients window. See Chapter 2, “Configuring and Monitoring Clients” for more information.

Understanding Backup Schedules

NetWorker makes it easy to set up your backup schedules. However, deciding which backup schedule best fits your requirements requires some planning.

When creating backup schedules, consider the following items.

- How long do you want to keep the backed-up data?
- How many versions of the data do you want to maintain?
- How much data do you have to back up?
- How many backup volumes do you want to use?
- How much time do you have to complete a network-wide backup?
- Do you want to use just a few backup volumes to recover from an entire disk crash?

For example:

Suppose the capacity of an 8 mm tape cartridge is about 5 gigabytes and the maximum transfer rate is around 500 kilobytes per second. Based on these figures, it would take more than 5 hours to back up 10 gigabytes of data. To back up a network with a lot of data, you may use more than one schedule to stagger the full backups over several days.

Think about how many backup volumes you want to keep; this number will depend on how often the data changes and how long you want to keep the files available for recovery.

If you run only incremental backups every night, you will need more backup volumes to fully recover from a disk crash. If a site has 10 gigabytes of data and 5% of all the data is modified each day, that means you need to back up 500 megabytes of data every day. At 400 kilobytes per second, 500 megabytes takes about 25 minutes to back up and fills about one tenth of an 8 mm tape cartridge. If you hold on to backups for three months, you will need to maintain twelve 8 mm tapes.

You must also determine a policy for recovering files. For example, if the users expect to recover *any version* of a lost file for at least three months, maintain *all* the backup volumes for a three-month period. On the other hand, if the users expect to recover only the *latest* version of a lost file, use level 1-9 backups to decrease the quantity of backup volumes you need to maintain. For more information on policies see, "Creating Index Policies" on page 142.

Selecting the Backup Levels

NetWorker supports four kinds of backup levels:

- 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 backup level. Each backup level is represented by numbers 1 through 9, where 1 represents the fullest backup and 9 represents the most minimal backup.
- Incremental—backs up files that have changed since the last backup, regardless of the level.
- Skip—skips the scheduled backup. You may want to skip a backup on a holiday if you know that no one will be available to change or add more backup volumes.

If you do not need to maintain every version of a backed-up file online, you may use a backup scheme that includes occasional full backups followed by level 1-9 and incremental backups.

Different backup levels allow you to trade off the number of backup volumes and amount of time required to complete a backup with the number of backup volumes and amount of time required to recover from a disk crash.

Figure 6-4 illustrates how backup levels work.

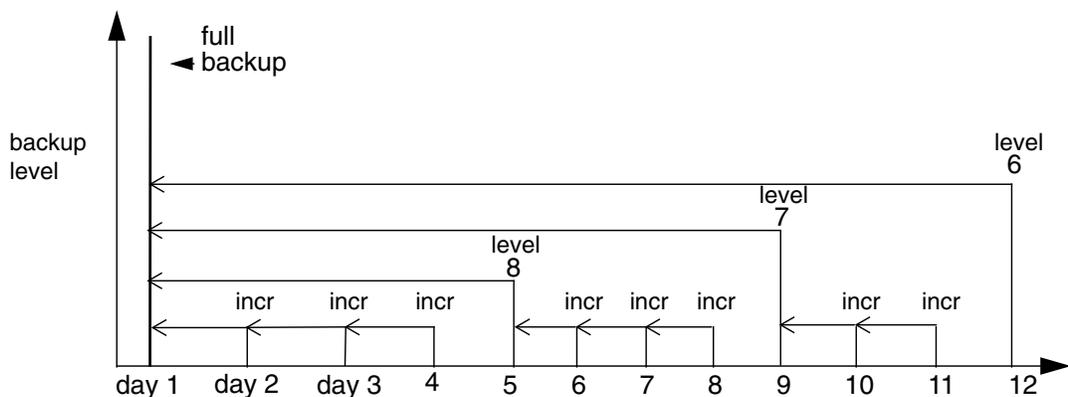


Figure 6-4 Backup Levels, Full and Incremental

Assume you use a new backup volume for each daily backup. On day 1, a full backup occurs. On day 2, the incremental backs up everything that has changed since the full backup. On day 3, the incremental backs up everything that has changed since day 2, and on day 4, everything that has changed since day 3. At this point, you have 4 backup volumes. To recover from a disk crash, you need all four of them—the one with the full backup (day1), and all the volumes with incremental backups.

On day 5, the level 8 backs up everything that has changed since the full backup. You no longer need the data on the backup volumes from day 2, 3, or 4. To do a full recovery from a disk crash, you only need two backup volumes: the full backup volume and the level 8 backup volume.

On day 9, the level 7 backs up everything that has changed since the full backup. You still need only two backup volumes to recover a disk: the full backup and the level 7.

Level 1-9 backups help you maintain control of your pool of backup volumes. Carefully planning your backup strategy allows you to recover everything on a disk with a minimum number of backup volumes.

Tip: You 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. See the section “Using Directives” in this chapter.

Creating Index Policies

This section describes how to create policies that automatically manage the online indexes. You can also manage the indexes manually with the *Indexes and Volumes* windows. See “Manually Managing the Online Indexes” on page 36 for a description of manual index management and index policy concepts.

When NetWorker starts a backup, it creates entries for the saved files in the online indexes. NetWorker maintains two types of indexes: a file index and a media index. The file index contains entries for the individual files, which are backed up in groups as save sets. The media index stores entries for the backup volumes that contain the backed-up save sets.

The media index does not recognize individual files, but instead recognizes the save sets to which the files belong. NetWorker maintains one file index per client and one media

index per NetWorker server. NetWorker uses the indexes as databases to locate files marked for recovery.

As these indexes grow, they take up more disk space. NetWorker uses browse and retention policies to manage and reduce the size of the online indexes. Use the Policies window to create browse policies for managing the file index entries and retention policies for managing the media index entries. A policy (preconfigured or one you create yourself) can be used as either a browse or a retention policy.

Apply the policies to individual clients in the Clients window by selecting them from the Browse and Retention policies scrolling lists.

NetWorker provides the following preconfigured policies: Decade, Month, Quarter, Week, and Year. See the section “Preconfigured Policies” on page 249 for more information on preconfigured policies.

Navigating the Policies Window

Use the Policies window to create new policies for your browse and retention policies.

Choose Policies from the Customize menu to open the Policies window. The Policies window appears, as shown in Figure 6-5.

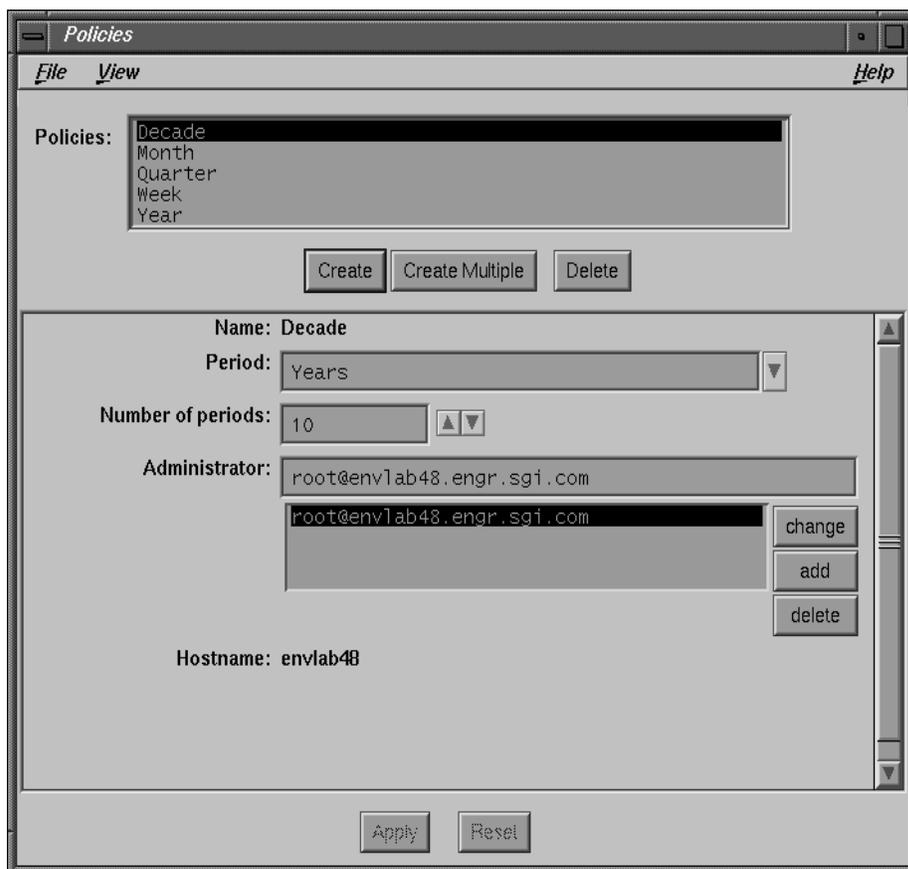


Figure 6-5 Policies Window

The Policies window contains the following items:

- Policies scrolling list—displays the names of all the policies known to the NetWorker server. When a policy is highlighted in the scrolling list, its name, period, and number of periods appear in the lower section of the window.
- Name field—displays the name of the currently selected policy. Enter the name of a new policy in the Name field after you click the Create button.
- Period field—displays the period of time you want the policy to apply. Use the arrow button to display the Periods scrolling list choices.

- **Number of periods field**—displays the number of periods applied to the policy. Use the arrow buttons to select the number of periods you want.

For example: A choice of **Months** for the **Period** and **3** for the **Number of Periods** means NetWorker automatically removes index entries from the online file index that are older than three months.

Creating a New Policy

To create a new policy, follow these steps:

1. Click the **Create** button.
2. Enter a name for the policy in the **Name** field.
3. Select a **Period** for the policy.
4. Select the **Number of periods** for the policy by clicking the arrow button to select from the scrolling list or by entering a number in the field.
5. Click the **Apply** button to apply your selections. After you click **Apply**, the newly created policy appears in the **Policies** window scrolling list and in the **Clients** window as both a **Browse** policy and a **Retention** policy.

If you make a mistake, click the **Reset** button to revert to the previous selections.

You can enter any name for the policy that you want. Using the time period for policy names helps you remember the length of the policy. For example, when you specify **Quarter** as a browse policy for a client, it is easy to tell that the entries are kept in the online file index for three months.

Deleting a Policy

Follow these steps to delete a policy:

1. Select the policy you want from the scrolling list.
2. Click the **Delete** button. A warning message appears, as shown in Figure 6-6.
3. Click **OK** to delete the policy, or click **Cancel** if you change your mind.

Note: You cannot delete the preconfigured NetWorker policies, or policies currently applied to a client.

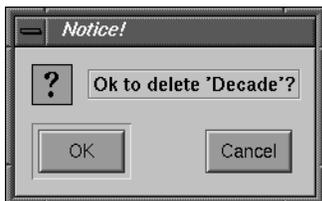


Figure 6-6 Notice to Delete Policy

Changing a Policy

Follow these steps to change a policy:

1. Select the policy in the scrolling list.
2. Select the Period and Number of periods for the policy and make the changes. However, you cannot change the name of the policy.
3. Click the Apply button to change the policy.

NetWorker does not allow a client browse policy to exceed its retention policy. This means that an index entry for a save set must be removed from the file index before it can be marked recyclable or removed from the media index.

Understanding Index Policies

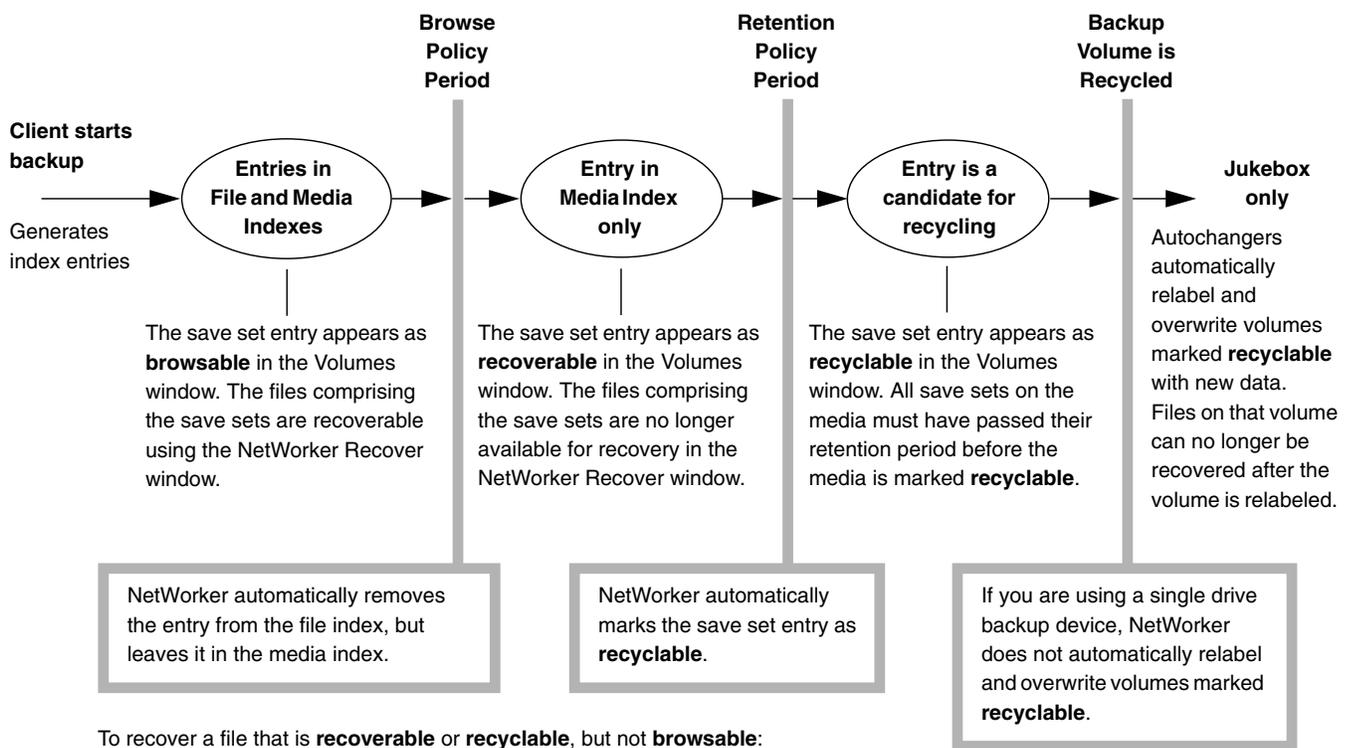
Every backup adds entries to the indexes, so the indexes require management in order to control their size and contents. NetWorker automatically manages the contents of the indexes by using the policies you chose for each client. You can conserve disk space by using automatic index policies to remove old index entries.

Each client requires two separate policies: a browse policy and a retention policy. NetWorker compares the backup date of the entries in the indexes with the time period specified by the policies. The browse policy determines how long entries for your files remain in the online file index and thus browsable in the NetWorker Recover window. File index entries older than the browse policy plus one cycle are automatically removed from the online file index, reducing disk space requirements. The retention policy determines how long entries for your volumes are retained in the media index and are available for recovery. Media index entries older than the retention policy plus one cycle

are marked as recyclable in the media index, making the backup volume available for relabeling and overwriting with new data. The files on the volumes are deleted when you overwrite the volume with new data.

If you create new policies for automatic index management using the Policies window, the policies appear as selections for browse and retention policies in the Clients window. You determine the amount of time files remain browsable and recoverable for a client.

Figure 6-7 illustrates the interaction between index entries and policies.



To recover a file that is **recoverable** or **recyclable**, but not **browsable**:

1. Determine which save set contains the files you want to recover, and use the Save Set Recover window to recover that save set. Using the Save Set Recover window to recover files only recovers the data, not the index entries.
2. If you want to recover the index entries as well, use the **scanner** command to recover files from a backup volume and to recover index entries for those files.

Figure 6-7 Automatic Index Management

Using Directives

Directives can reduce the amount of data you back up, possibly even eliminating the need to change backup volumes on the days you perform a full backup. A directive contains instructions to assist the backup process. Sometimes directives require execution of special programs called Application Specific Modules (ASMs), which are NetWorker modules that perform specific operations on files. See the `uasm(1M)` reference page for more information on ASMs.

NetWorker is shipped with the following preconfigured directives: DOS standard, Index, NetWare standard, NT standard, NT with compression, Unix standard, and Unix with compression. Each set of directives covers the most important and most useful backup instructions.

For more details about preconfiguration, see “Preconfigured Directives” on page 255.

Navigating the Directives Window

Use the Directives window to change preconfigured directives or create new ones. To open the Directives window, choose Directives from the Customize menu. The Directives window appears, as shown in Figure 6-8.

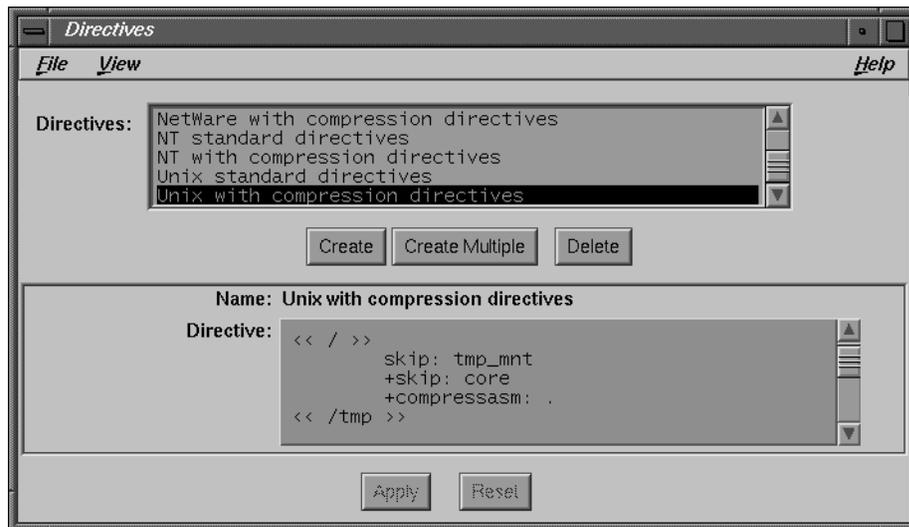


Figure 6-8 Directives Window

The Directives window contains the following items:

- Directives scrolling list—displays the names of the directives known to the NetWorker server.
- Name field—displays the name of the currently selected directive. After you click the Create button, enter the name of a new directive in the Name field.
- Directive field—shows the syntax for the currently selected set of directives. Each directive is preceded by the directory to which it applies (for example, << /tmp >>).

Adding or Changing Directives

NetWorker backs up your data effectively with the existing preconfigured directives. However, you can create new directives or change existing ones.

To create a set of directives or modify existing instructions, follow these steps:

1. Click the Create button to add a new directive or select an existing directive from the Directives scrolling list.
2. Type your instructions in the Directive field.
3. Click the Apply button, after you add or change a directive, to apply the change.

For more information about directives, including algorithms and the correct syntax used to apply them, see the `nsr(1M)` reference page.

Understanding Directives

Suppose your NetWorker client is a fast machine, but your network is slow or busy. You may want to *compress* client files in order to use fewer backup volumes and less network bandwidth backing up its files. Using the **compressasm** directive can reduce the space consumed on a backup volume by as much as 50% (actual savings may vary). If you use **compressasm** on all files, a full backup of 8 GB may fit on a single backup volume.

Tip: If your backup hardware compresses data, it is wasteful to use the NetWorker **compressasm** directive and device compression at the same time.

Use the *skip* directive to omit 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 have NetWorker skip backing up *core* files and object (*.o) files.

Using Label Templates

NetWorker generates labels for backup volumes according to the rules of a *label template*. Label templates provide a method for consistently naming and labeling your backup volumes. You can use the preconfigured label templates or create custom label templates using the Label Templates window.

NetWorker also provides a facility for affixing your autochanger volumes with bar code labels. See “Understanding Labeling With Bar Codes” on page 182 for more information.

This section describes how to use label templates, introduces preconfigured templates, describes how to create a custom label template, and provides examples for your use.

NetWorker uses label templates to label and keep track of your backup volumes. All of the backup volumes belong to volume pools that require corresponding label templates. For more information on pools, see the section “Using Volume Pools” on page 85.

It does not matter to NetWorker how the volumes are labeled as long as each one has a unique name. It tracks the backups and maps them to the backup volumes, so you do not have to remember which backups are on which volumes. When NetWorker needs a backup volume for recovery, it requests it by name in the Pending display of the NetWorker Administrator window.

NetWorker provides the following preconfigured label templates: Archive, Archive Clone, Default, Default Clone, Full, Migration, Migration Clone, NonFull, Offsite, PC Archive, PC Archive Clone, and Two Sided. These preconfigured label templates allow you to easily begin labeling your backup volumes without having to create any templates yourself.

NetWorker also provides preconfigured volume pools with names corresponding to the preconfigured label templates. The preconfigured volume pools automatically use the preconfigured label templates of the same name. The number range for all of the preconfigured label templates starts at “001” and ends with “999” to allow for expansion of the volume pools.

For more information about preconfigured label templates, see “Preconfigured Label Templates” on page 253.

Choosing a Label Template

Before you choose a label template, you must first choose the device that contains media ready for labeling from the Devices display in the NetWorker Administrator window. Next, choose a label template in the Pools window, then start the labeling process by choosing Label from the Media menu.

For a label template to appear as a choice in the Label window, the associated pool must be enabled in the Pools window (see Figure 6-9).

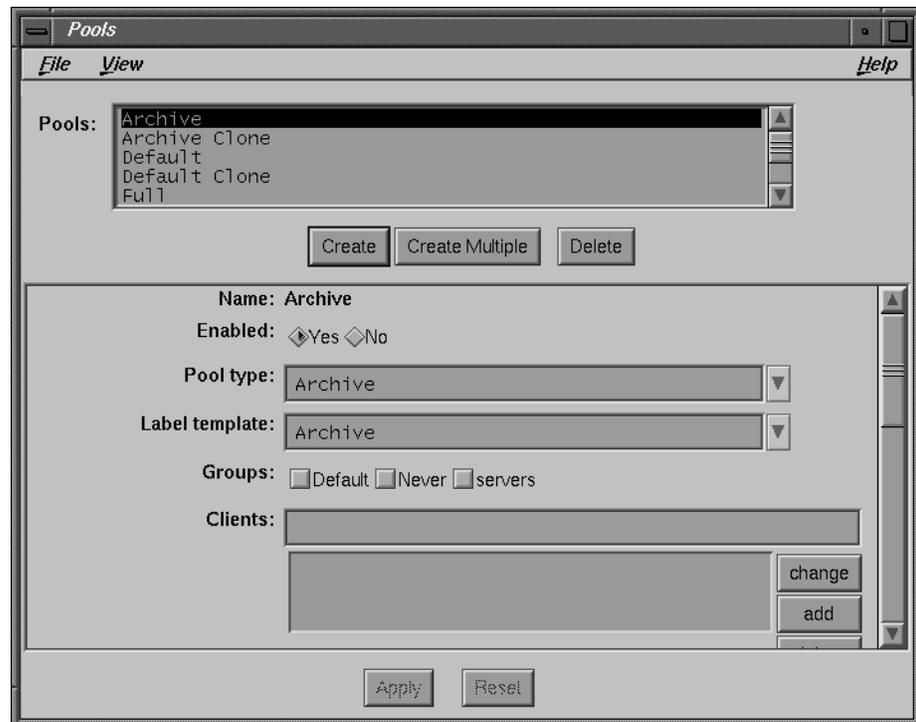


Figure 6-9 Pools Window

Tip: Whether you use a preconfigured or custom label template, remember to attach an external label to each volume with the correct label template name.

Navigating the Label Templates Window

This section explains how to create custom label templates using the Label Templates window.

Open the Label Templates window by choosing Label Templates from the Customize menu. The Label Templates window appears, as shown in Figure 6-10.

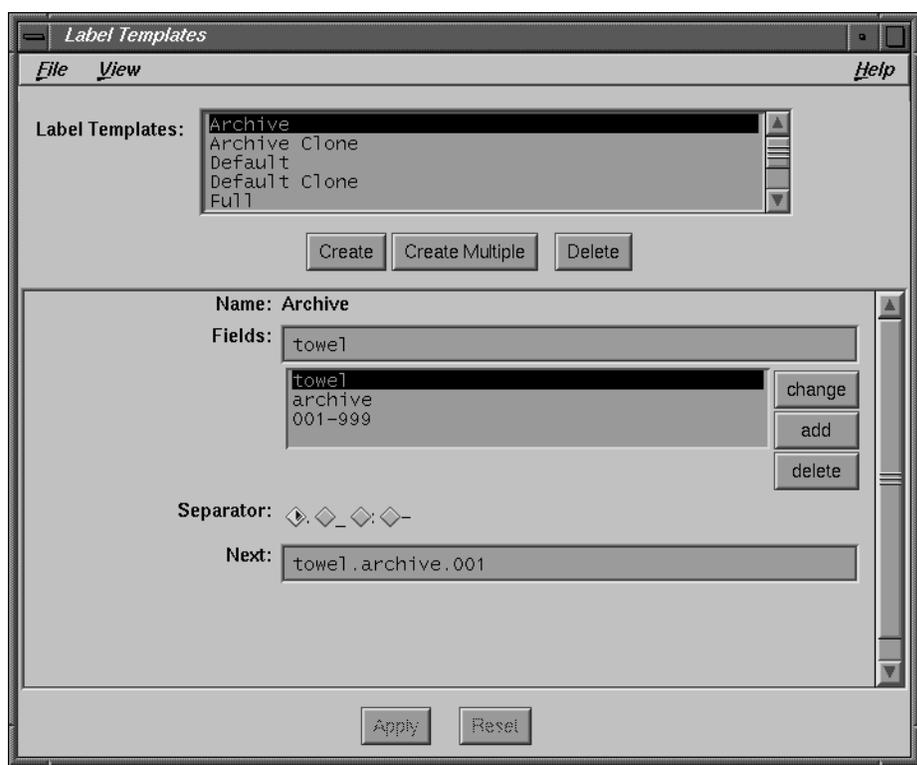


Figure 6-10 Label Templates Window

The Label Templates window contains the following items:

- Label Templates scrolling list—displays the label templates known to the NetWorker server. The preconfigured label templates and any templates you create appear in this list.

- Name field—displays the name of the selected label template. Use the Name field to enter the name of a new label template after you click the Create button.
- Fields field—use for entering the fields that make up a label template. The scrolling list displays the components of the label template. Each component appears as one line in the display. There are four types of components:
 - range of numbers—for example, *001-999*
 - range of lowercase letters—*aa-zz*
 - range of uppercase letters—*AA-ZZ*
 - character string—for example, *Accounting, 001*

If you enter a range of numbers or characters, the label number increments with each new label. For example:

first label Engineering.001

second label Engineering.002

third label Engineering.003

Each range includes a starting value, a dash (-), and an ending value. The starting value and the ending value must have the same number of characters. For example, use “aaa-zzz,” not “aa-zzz.” (This rule does not apply to a list of strings or words.) You may 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.

- Separator choices—use the period, dash, colon, or underscore to separate each component of the label template. If you do not select a separator choice, the label components will not have separators (for example, AA00aa).
- Next field—displays the next label NetWorker will use to label the backup volume according to the template. If you want to force a label or start the labeling scheme at a certain point, enter the starting label in the Next field. NetWorker continues to generate labels from that point on, according to the rules of the template. If you want NetWorker to generate the first name for you, leave the Next field blank.

Examples

Use the examples in this section to help you create a custom label template that identifies your media in a logical and meaningful order.

If your label template matches your backup media storage system, it helps you organize and locate the media in storage. For example, if you store backup media in bins or racks, include a place for the bin or rack number in the label template.

It is helpful if label template names match the name of the corresponding pool. Here is an example:

Name	Label	Pool Name
AcctFull	AcctFull.001	Accounting
	AcctFull.002	
	AcctFull.003	
EngTest	EngTest.001	Engineering Test
	EngTest.002	
	EngTest.003	

Example 1—A Storage-Oriented Template

In this example, assume you have only one NetWorker server and one autochanger. Your backup media is magnetic tape, arranged in numerical order, and organized in three racks. You want to create a scheme that is simple so anyone can easily store and find the backup volumes on the rack shelves.

Figure 6-11 shows three racks, named “1,” “2,” and “3.” Each rack has five shelves; the top shelf is “1” and the bottom shelf is “5.” Each shelf holds 100 tapes, left to right.

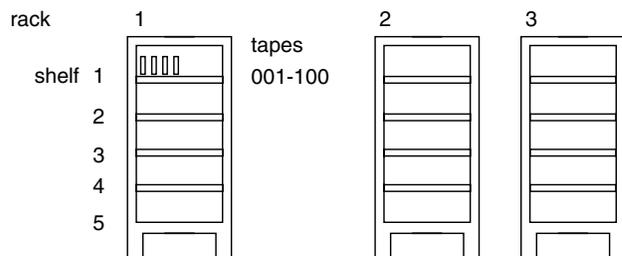


Figure 6-11 Rack Storage of Backup Volumes

Using this labeling scheme, a tape labeled “2-3-54” will be stored in the second rack, on the third shelf, positioned between the fifty-third and fifty-fifth tapes on the shelf. NetWorker labels the tapes in the autochanger consecutively according to the label template.

NetWorker starts with the tapes in the first rack, on the first shelf, and numbers the tapes 1-1-001 to 1-1-100. Then, the tapes on the second shelf of the first rack are labeled, numbered from 1-2-001 to 1-2-100, and so forth, until the last tape in the third rack is labeled "3-5-100." If you want to add a fourth rack to your tape storage system, you can easily do so by changing the label template.

Tip: 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 to create a template for only 10 tapes and run out of labels.

Follow these steps to create the storage oriented label template:

1. Choose Label Templates from the Customize menu.
2. In the Label Templates window that appears, click the Create button.
3. Enter the name of the new labeling scheme in the Name field; for example: "Engineering backups."
4. Enter the component of the first field into Fields. This field indicates the rack number, so enter "1-3" (without the quotes).
5. Click the add button.

Note: You must enter the same number of characters for all numbers in the label template. For example, use: "01-20" not "1-20" or "1-9" not "01-9" templates.

Add the second and third components of the label template:

1. Enter the second component of the label template into Fields. This component represents the shelf number, so enter "1-5" in the field.
2. Click the add button.
3. Enter the third component of the label template, the tape position numbers. Enter "001-100" into Fields.
4. Click the add button.
5. Click the dash (-) Separator button to choose the separator between components.
6. Click the Apply button to apply the new template to the NetWorker server.

The Label Templates window now looks similar to Figure 6-12.

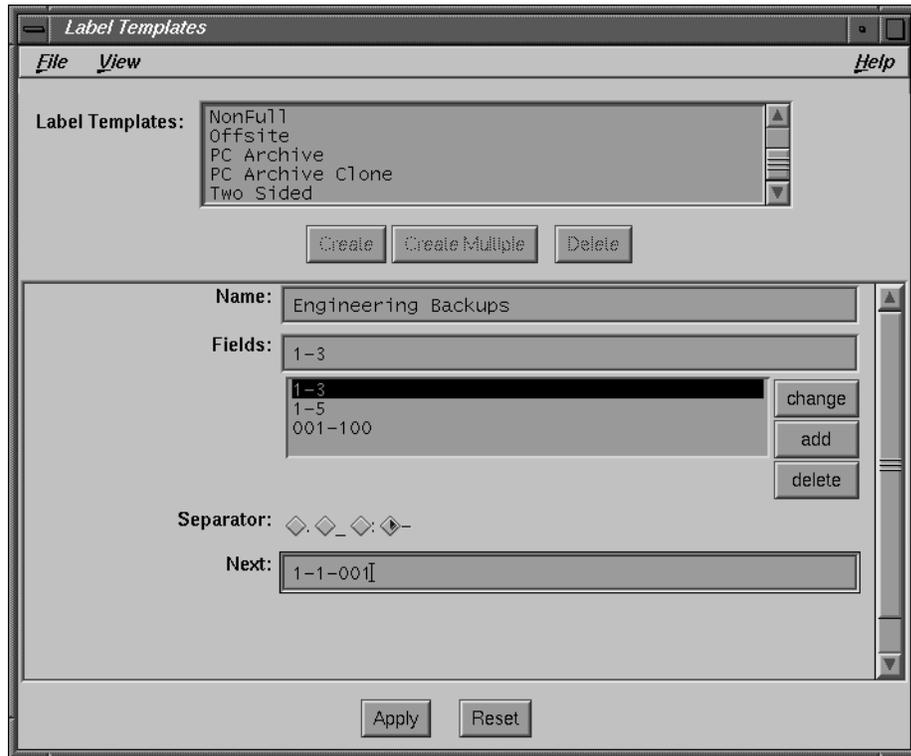


Figure 6-12 Label Templates, Storage-Oriented

Your new template appears in the Label Templates scrolling list, template components appear in the Fields scrolling list, and the Next field displays the label of the next tape.

Example 2—A Sequential Label Template

This example uses a sequential labeling scheme. In this scheme, there is no storage system pattern to follow—the labels are simply a way for NetWorker to identify the backup volumes. This label template generates lots of labels—you will probably never “run out” of labels for your backup volumes.

The label template in this example has two fields: “aa-zz” and “00-99.” It generates 67,600 different labels (26^2 times 10^2).

Table 6-1 shows an example of how NetWorker generates labels.

Table 6-1 Label Template Example

aa.00, aa.01, ... aa.98, aa.99,
ab.00, ab.01, ... ab.98, ab.99,
...
az.00, az.01 ... az.98, az.99,
ba.00, ba.01, ... bz.98, ba.99,
...
zz.00, zz.01, ... zz.98, zz.99

Follow these steps to create a sequential label template:

1. Choose Label Templates from the Customize menu.
2. In the Label Templates window that appears, click the Create button.
3. Enter the labeling scheme into the Name field, for example, "Personnel backups."
4. Enter the component of the first field into Fields. In this example, this field contains pairs of character strings, so enter "aa-zz" into the field.
5. Click the add button.

Add the second component of the label template:

1. Enter the second component of the label template in Fields. This field contains double digits, so enter "00-99" in the field.
2. Click the add button.
3. Click the period (.) Separator button to choose the separator between components.
4. Click the Apply button to apply the new template to the NetWorker server.

This template is useful if you do not want to attach any meaning to the labels, but want to be able to generate lots of them over the years without devising a new template.

Example 3—A Two-Sided Media Label Template

If you have an optical backup device, you can use the Two Sided preconfigured template or create a custom label template that accommodates two-sided media. Label templates for two-sided media have *a* and *b* fields to designate the two sides of the optical disk.

Follow these steps to create a label template for two-sided media:

1. Choose Label Templates from the Customize menu.
2. In the Label Templates window, click the Create button.
3. Enter a name for this template in the Name field (for example, "TwoSidedLabel").
4. Add the following components to the Fields list: "server_name," "001-999," and "a-b" where "a-b" distinguishes the two sides of the media.
5. Click the period (.) Separator button to choose the separator between components.

The Label Templates window now looks similar to Figure 6-13.

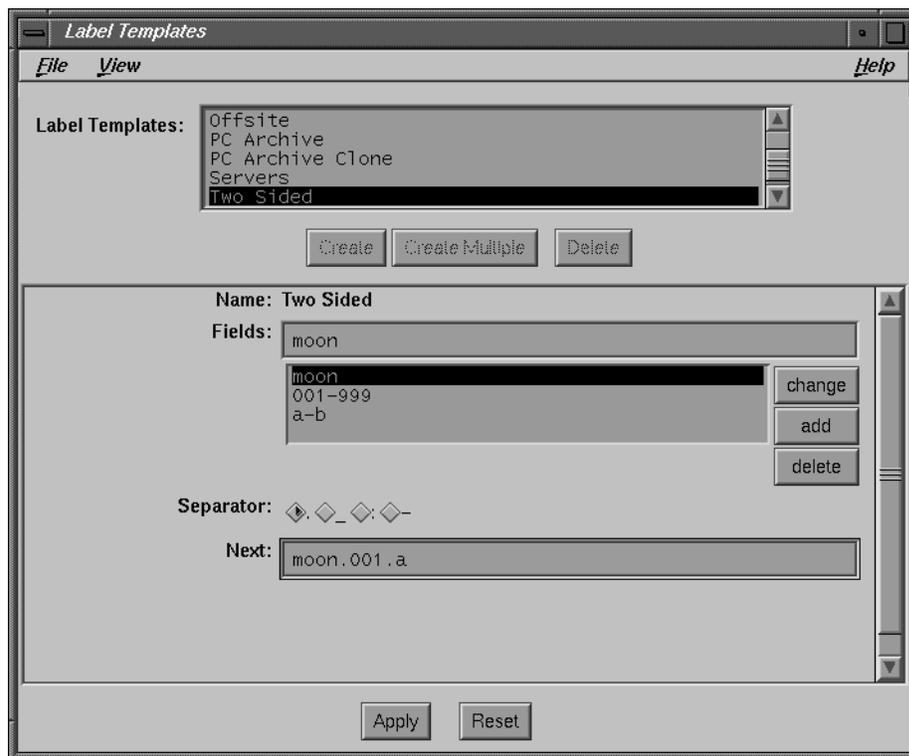


Figure 6-13 Label Templates, Two-Sided

6. Click Apply to save and apply the label template.

Using the example shown, NetWorker labels the first side of the optical disk as *moon.001.a* and the second side as *moon.001.b*.

When you create a label template, the order in which you enter each component of the template is important. NetWorker applies each component in a left-to-right order, starting with the first one you enter. NetWorker displays the ordered components in the Fields scrolling list, with the first component at the top.

Configuring Backup Groups

NetWorker allows you to bundle clients into backup *groups* so you can easily start selected groups of clients backing up at different times. When you create the different groups, be sure to stagger their backup times to avoid overloading the server.

NetWorker is shipped with a Default group to which NetWorker automatically adds all of the clients. You must enable the Default group before it will automatically back up at 3:33 a.m. You are permitted to make changes to the Default group, but you may not delete it.

If you have a large number of clients, consider creating several groups with different backup times to help reduce network traffic. For example, start the backup of the engineering department machines at four o'clock in the morning, and all the other clients on the network at midnight. You can have any number of backup groups on your NetWorker server. The clients in each backup group start their automatic backups according to the start time of the group.

Note: Create groups in the Groups window. Select groups for individual clients in the Clients window.

Navigating the Groups Window

This section explains how to create new groups using the Groups window.

Open the Groups window by choosing Groups from the Customize menu. The Groups window appears, as shown in Figure 6-14.

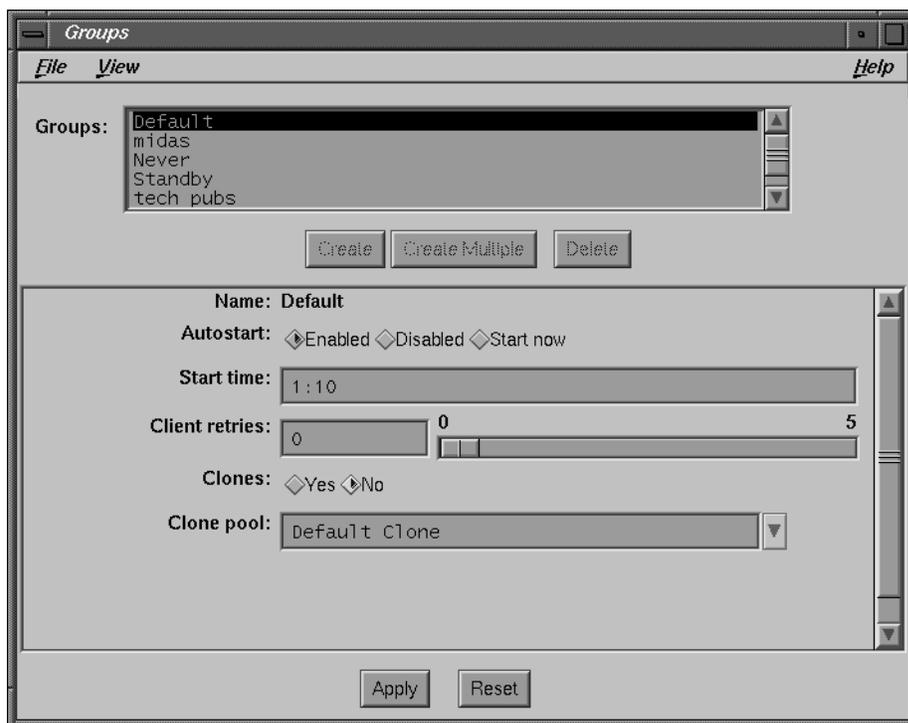


Figure 6-14 Groups Window

- Groups scrolling list—displays all backup groups known to the NetWorker server. When you first open the Groups window, you see the preconfigured Default group in the Groups scrolling list. When you select a name in the scrolling list, it also appears in the Name field along with its backup time in the Start time field.
- Name field—displays the name of the group selected in the Groups scrolling list. Use this field to enter the name of a new group after you click the Create button.
- Autostart choices—allow you to disable, enable, or immediately start a network-wide backup. The Default group comes disabled. To allow NetWorker to back up your clients in the Default group, you must select Enabled from the Autostart choices. The following lists the available Autostart choices:
 - Enabled—allows network-wide backup of the highlighted group to start at the scheduled time
 - Disabled—prevents network-wide backup of a group from starting

- Start now—immediately starts a network-wide backup of a group, instead of using the scheduled start time
- Start time field—displays the time NetWorker starts the backup of the clients in the group selected in the Groups scrolling list. The Default group is preconfigured to back up at 3:33 am.
- Client retries field—the number of times NetWorker will attempt to back up a failed client during a scheduled backup. Use the sliding bar or type a number in the field. For example, if you enter the number four and a client fails to back up, NetWorker will attempt to back up that client as many as four times.
- Clones choices—select Yes to automatically clone the data being backed up for a group. Select No if you do not want to clone the data being backed up in a group. In order to clone data, NetWorker must support concurrent devices.
- Clone pool field—enter the name of a pool or use the arrow button to select the clone pool from the scrolling list to which cloned data is backed up.
- Migration clone pool field—enter the name of a pool or use the arrow button to select the migration clone pool from the scrolling list to which cloned migration data is backed up.

Creating a New Group

You may choose to add all of your clients to the Default group. However, that may not be an efficient way to accomplish backing up client data across the network.

To create a new group and assign a backup start time to the group, follow these steps:

1. Open the Groups window by choosing Groups from the Customize menu.
2. Click the Create button, and enter the name of the new group in the Name field. Choose a name that describes the systems in the group, either by department or schedule.
3. Select Enabled from the Autostart choices.
4. Delete the time in the Start time field and enter a new time, using 24-hour clock convention. (For example, enter 10 p.m. as 22:00.)

Tip: When you select start times for the different groups, make sure you schedule them far enough apart so that one group has sufficient time to complete backup before the next group begins.

5. Use the Client retries sliding bar to set a value for the number of times NetWorker will attempt to back up a failed client.
6. Select one of the Clones choices. If you click Yes, then select or add the name of the clone pool in the Clone pool field.
7. Click the Apply button to apply the changes.

The name of the group appears in the Groups scrolling list. If you make a mistake, click the Reset button to clear your changes and start over.

If this error message appears, it means that you do not have permission to make configuration changes:

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

See "Adding or Changing Administrators" on page 52 for more information.

After you create a backup group and set the start time, the group appears as a check box in the Clients window. To assign clients to the backup group, click the check box displayed for the group in the Clients window. See the section "Configuring Clients" on page 29 for more information.

Retrying a Failed Client Backup

NetWorker normally attempts to back up a client in a group once during a scheduled backup. If NetWorker cannot back up a client in a group (because the client machine is down, for example), it will not try again unless you set a value for Client retries.

Use the Client retries sliding bar (Figure 6-15) in the Clients window to control the number of times NetWorker should keep trying to back up failed clients.



Figure 6-15 Client Retries Slider Bar

1. Drag the sliding bar to the right or left to set the number of client retries.
The number of times you set for NetWorker to attempt backing up a failed client appears in the Client retries field.
2. Click the Apply button.

Changing a Backup Time

To change the start time for a backup group, follow these steps:

1. Open the Groups window.
2. Select the group for which you want to change the backup time from the Groups scrolling list.
3. Delete the time in the Start time field and enter a new time, using 24-hour clock convention. (For example, enter 11 p.m. as 23:00.)
4. Click the Apply button to apply the changes.

If you make a mistake, click the Reset button to clear your changes and start over.

Deleting a Backup Group

NetWorker does not let you delete a backup group while clients still belong to it. Before you try to delete a group, check the Clients window to make sure the group does not have clients. The Default group cannot be deleted.

To delete a group, follow these steps:

1. Open the Groups window.
2. Select the name of the group you want to delete from the Groups scrolling list.
3. Click the Delete button. A dialog box appears to confirm your request to delete the group.

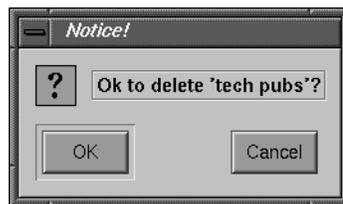


Figure 6-16 Notice to Delete Group

4. Click Ok to delete the group.

If the group does not have clients, it is deleted and its name disappears from the Groups scrolling list. If the group you are trying to delete still has clients, the following warning appears:

```
The "groupname" group is still used by one or more clients.
```

Setting Up Event Notification

System administrators need to know about events taking place in the NetWorker environment. To keep you informed of backup activity, NetWorker displays messages on the system console (*/dev/console*) or mails them to you electronically.

These notices alert you to important NetWorker events. If a group of clients did not complete a nightly backup, for example, you are informed of this by the savegroup completion notices in your electronic mail.

NetWorker provides preconfigured notifications for the following events:

- Bootstrap—prints information needed for recovering from a disk crash, including the server index, media index, and *nsr.res* file.
- Cleaning cartridge—reports cleaning cartridge status.
- Device cleaning—reports device cleaning status.
- Index size—suggests it is time to reduce the online index.
- Migration—reports file migration status. Refer to the Legato documentation on Hierarchical Storage Management for information on using NetWorker HSM.
- Registration—requests that you properly register your NetWorker products.
- Savegroup completion—reports on the success of a scheduled backup.
- Tape mount request 1, 2, and 3—reports that NetWorker is requesting mounting of a backup volume.

Note: For more information on the preconfigured notifications, see the section “Preconfigured Notifications” on page 255.

Navigating the Notifications Window

This section describes the contents of the Notifications window and how to use it for creating new event notifications.

Open the Notifications window by choosing Notifications from the Customize menu. The Notifications window appears, as shown in Figure 6-16.

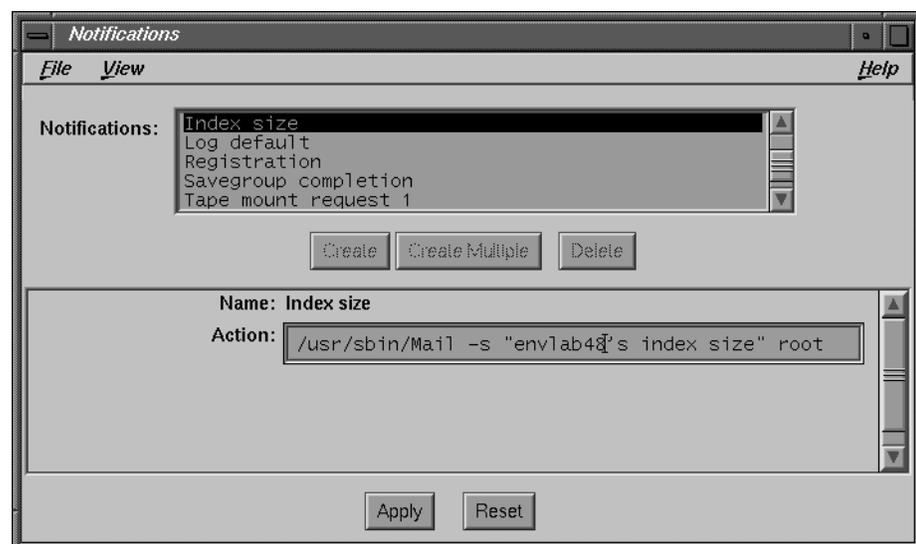


Figure 6-17 Notifications Window

The Notifications window contains the following items:

- Notifications scrolling list—displays the names of all the notifications known to the NetWorker server. Information about a selected notification is displayed in the lower half of the window.
- Name field—displays the name of the currently selected event notification. Use this field to enter the name of a new notification after clicking the Create button.
- Action field—displays what type of action NetWorker takes with the notice. For example, NetWorker can use electronic mail to notify individuals that a volume needs to be mounted in the server. The action can be any UNIX command. You can change the Action field to customize how NetWorker posts notifications.

Creating a New Notification

To create a new notification, follow these steps:

1. Open the Notifications window by choosing Notifications from the Customize menu.
2. Click the Create button, and enter the name of the new notification in the Name field.
3. Enter the code or commands in the Action field.
4. Click the Apply button.

Tip: Choose Details from the View menu if you want to use the expert mode options for creating a more detailed notification. All event choices are enabled, by default. Click to disable event choices you do not want.

Summary

The NetWorker Customize menu provides access to powerful and versatile tools for altering backup configurations. You can create highly individual and specific methods for backing up data on your network. On the other hand, except for enabling backup groups in the Groups window, some administrators seldom have any need to touch the Customize menu.

Autochanger Software Module

This chapter provides information for configuring and operating NetWorker with your autochanger. See Chapter 4, “Managing Media and Backup Devices,” for additional information pertinent to all backup devices.

The software displays the term “jukebox” so this chapter also uses the term “ jukebox” to refer to an autochanger. Both terms refer to a variety of backup devices, including carousels, libraries, near-line storage, datawheels, and autoloaders.

Configuring Jukeboxes

Jukeboxes automate the task of loading, mounting, and labeling backup media. Before NetWorker can back up to a jukebox, you must do the following:

- Install and enable NetWorker server edition, NetWorker network edition, or NetWorker power edition
- Enable the NetWorker Autochanger Software Module.
- Load and label your backup volumes.

Several jukebox configuration changes occur when you install and enable the Autochanger Software Module. After installation, you can make changes to the available slots, devices, bar code choices, cleaning devices, and auto media management settings. If you want to install additional jukeboxes on your NetWorker server, see the *IRIX NetWorker Installation Guide*.

Note: You cannot add or create jukebox resources using the NetWorker Administrator program. Instead, you must modify jukeboxes previously installed and configured using the **jb_config** program.

Use the Jukeboxes window to see the names of the autochangers NetWorker recognizes, and to remove autochangers no longer in use.

Navigating the Jukeboxes Window

The Jukeboxes window displays the configurations selected when you installed the Autochanger Software Module. After installing, enabling, and configuring the jukebox, you can make changes to the following in the Jukeboxes window: Available slots, Devices, Bar code reader, Match bar code labels, and Auto media management.

Open the Jukeboxes window by choosing Jukeboxes from the Media menu. The Jukeboxes window appears, as shown in Figure 7-1.

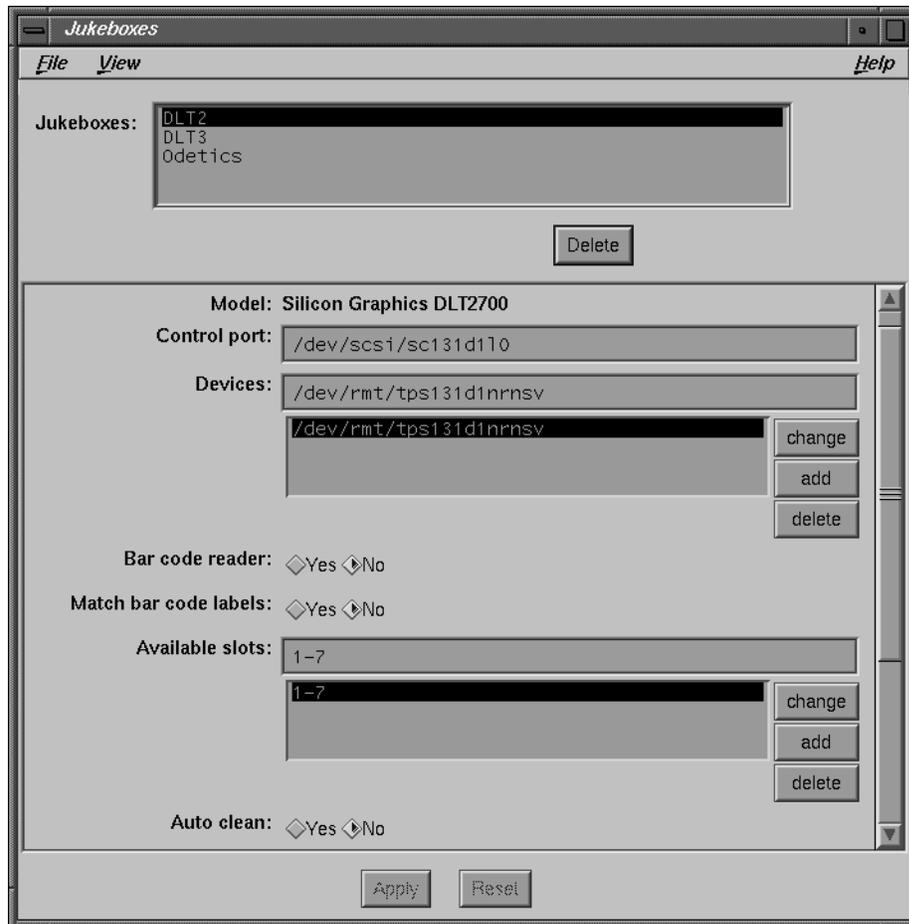


Figure 7-1 Jukeboxes Window

The following jukebox options appear in the Jukeboxes window. You determine all of these settings when you install and enable the Autochanger Software Module.

- Jukeboxes scrolling list—displays the names of the jukeboxes known to the NetWorker server.
- Name field—displays the name of the jukebox currently selected in the Jukeboxes scrolling list.
- Model choices—displays the model or type of the currently selected jukebox.
- Control port field—displays the system *pathname* of the device that controls the jukebox loading mechanism. The loading mechanism loads and unloads the backup volumes into and out of the device(s) in the jukebox.
- Devices field—displays the system *pathname* of the device that controls the media drive currently selected from the Devices scrolling list. The scrolling list displays the system *pathnames* for all of the devices located in the jukebox. Devices inside a jukebox, like single-media devices, are configured using the Devices window. Configure the jukebox devices before using the Jukeboxes window to complete the jukebox configuration.

Use the change, add, and delete buttons to change, add, or delete the highlighted device *pathname* in the Devices field in the Jukeboxes window. See “Navigating the Devices Window” on page 96 for more information about the Devices window.

Devices inside of a jukebox are designated by (J) next to their *pathname* in the Devices display of the NetWorker Administrator window, as shown in Figure 7-2.



Figure 7-2 Jukebox Devices Field

- Bar code reader choices—provides choices for labeling your backup volumes with bar code labels.

For more information on bar code labels, see the section “Labeling With Bar Codes” on page 180.

- Match bar code labels choices—provides choices for matching the bar code label to the volume label.

If your jukebox device lacks a bar code reader, select No for the Bar code reader and Match bar code labels choices.

- Available slots field—displays the available slot range currently selected in the scrolling list. NetWorker uses any backup volumes in this range to back up the data from the clients on the network. The Available slots correspond to the numbered slots available for holding backup volumes inside the jukebox. Typically, the slots are identified by a range of numbers. Entering just a single number—for example, 10—restricts access to just that slot (slot “10”). Entering the number “10” is not the same as entering “1 through 10.”
- Auto clean choices—use to enable or disable the automatic cleaning function. When Auto clean is enabled and Cleaning interval is defined, NetWorker automatically cleans the devices in the jukebox and sends a “Device cleaned” notification, indicating that the cleaning process is complete. Auto clean is disabled, by default.

When you disable Auto clean, NetWorker ignores the value of Cleaning slots and uses the slots for backup volumes. NetWorker does not automatically clean the devices in the jukebox. If the Cleaning interval field contains a value, NetWorker sends a “Device cleaning required” notification message, indicating that the device needs cleaning. Disable Auto clean to manually load a cleaning cartridge in the jukebox.

- Cleaning slots field—reserves a specified slot or range of slots in the jukebox for the cleaning cartridge(s). NetWorker automatically selects the correct slots for cleaning when you configure the jukebox.

A single positive integer reserves the specified slot; two positive integers separated by a hyphen reserve a specified range of slots. If this field is blank or Auto clean is disabled, no slot is reserved for a cleaning cartridge.

- Default cleanings field—determines the number of times NetWorker uses a new cleaning cartridge mounted in a jukebox. NetWorker automatically selects the correct number of default cleanings when you first configure the jukebox.

When the cartridge has been used the specified number of times, NetWorker sends a “Cleaning cartridge expired” notification message, indicating the cartridge is used up and needs replacing.

- Auto media management choices—gives NetWorker exclusive control over the media loaded in a jukebox. If you select Yes, NetWorker automatically labels, mounts, and writes to volumes it considers unlabeled. If you select No, NetWorker ignores the unlabeled media and does not consider it for backup.

Adding More Jukebox Devices

To add more devices to a jukebox, you must first configure the NetWorker server to recognize the added devices. Then you must enter pathnames for the jukebox devices in the Jukeboxes window so NetWorker recognizes that these devices belong to the jukebox.

Tip: After you add a device pathname to the Devices window, you must also add it to the Devices field in the Jukeboxes window. NetWorker must associate the device in the jukebox with the jukebox name. See “Navigating the Devices Window” on page 96 for a description of the Devices window.

To add a device to a jukebox, follow these steps:

1. Choose Devices from the Media menu.
2. Click the Create button in the Devices window, shown in Figure 7-3.

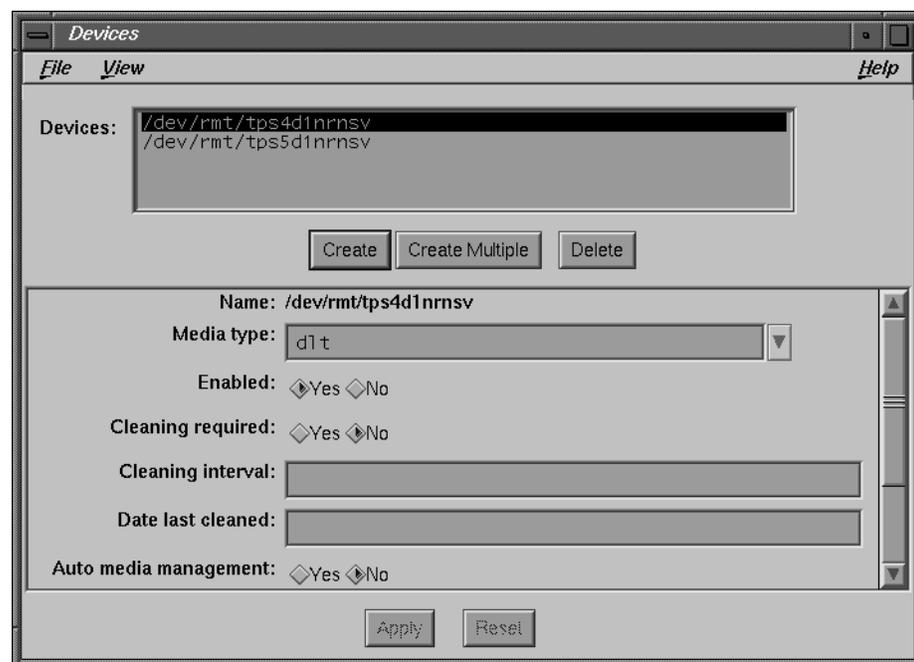


Figure 7-3 Devices Window

3. Enter the jukebox device pathname in the Name field.

4. Select the appropriate Media type choice for the device you are adding to the jukebox. By default, Yes is selected for the Enabled choices.
5. Enter a value for Cleaning interval. This field accepts any combination of numbers and time periods, with or without a space between: 2day, 2 days; 3week, 3 weeks; 6month, 6 months; and so on.

NetWorker keeps track of the time elapsed since the last cleaning and determines when cleaning is required. A blank value in Cleaning interval disables the automatic cleaning function.

6. Click the Apply button.

Next, enter the jukebox devices in the Jukeboxes window:

1. Choose Jukeboxes from the Media menu.
2. In the Jukeboxes window, enter the jukebox device pathname in the Devices field, as shown in Figure 7-4.

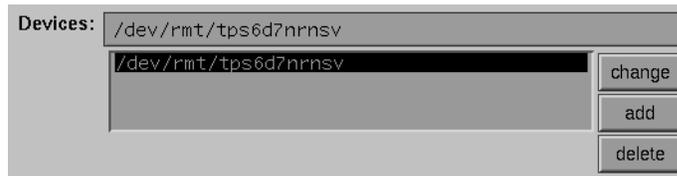


Figure 7-4 Pathname in Jukebox Devices Field

3. Click the add button to the right of the Devices scrolling list.
4. Click the Apply button.

Tip: If a jukebox has more than one device, the device *pathnames* must be listed in the same order as their physical location in the jukebox. If you are unsure of their order, refer to the jukebox hardware manual.

For an alternate method of determining the names of devices within a jukebox, see "Testing the Device Driver Installation" in the *Installation and Maintenance Guide*.

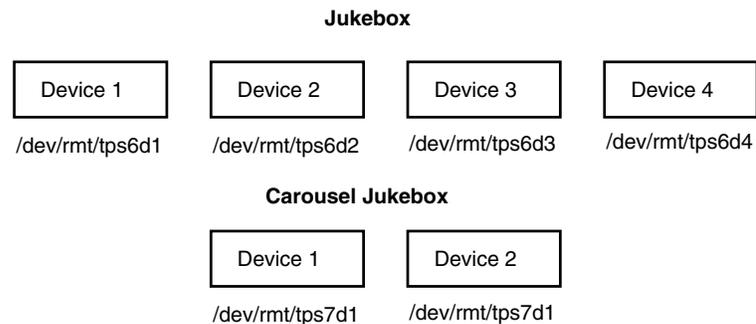


Figure 7-5 Jukeboxes With Multiple Devices

Changing Available Slots

The Available slots feature controls exactly which backup volumes NetWorker uses for backing up data. NetWorker uses *all* of the volumes in the jukebox for recovers. However, you can control which volumes NetWorker automatically selects for backups by designating a range of available slots in the jukebox.

For example, you may designate slots 1 through 7 for your NetWorker backups in a jukebox that contains 10 slots, as shown in Figure 7-6. The entries can be a range of slot numbers or a single slot number.

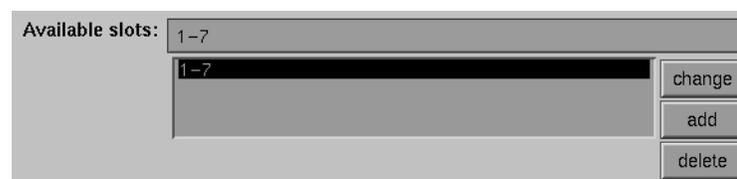


Figure 7-6 Jukeboxes Available Slots Field

To change the available slots:

1. Enter the range of available slots into the Available slots field.
2. Click the change button. The slot numbers appear in the Available slots scrolling list.
3. Click the Apply button in the Jukeboxes window.

With two-sided media, the number of available slots is always one half the number of labels (or sides of the optical disks). For example, if you have 32 optical disks, labeled *atlas.001.a* to *atlas.032.b*, you have a total of 64 labels (and 64 sides), but only 32 slots.

Tip: Make sure you have placed backup volumes in all the available slots of the jukebox so NetWorker can proceed uninterrupted with an automatic backup.

Cleaning Jukebox Devices

Maintaining a backup device in good working order requires periodic cleaning. NetWorker provides automatic cleaning of devices located in a jukebox; NetWorker does not support automatic cleaning for standalone devices.

Use both the Jukeboxes and Devices windows to make the necessary selections for automatically cleaning your jukebox devices. The choices specific to the jukebox appear in the Jukeboxes dialog box where you enable and disable the automatic cleaning feature. Select the cleaning slot(s) for the cartridge(s).

The functions specific to the devices located in the jukebox appear the in the Devices window where you are notified when a device needs cleaning, the date the device was last cleaned, and decide how often a device should be cleaned.

NetWorker cleans devices only before mounting or after unmounting a backup volume to ensure that there will not be any interference with other jukebox operations.

Cleaning cartridge support provides the notification messages listed in Table 7-1 to keep you informed about cleaning cartridge operations:

Table 7-1 Cleaning Cartridge Notifications

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 enabled and no usable cleaning cartridge is available.
Cleaning cartridge expired	Cleaning cartridge has been used the specified number of times.

Check the documentation from your jukebox manufacturer for recommendations on the frequency and cleaning methods for your jukebox devices.

Automatically Managing Media

Auto media management further automates the backup procedure by freeing the administrator from the task of prelabeling volumes before backup. When Auto media management is enabled in the Jukebox dialog box, NetWorker assumes that the backup volumes loaded in the jukebox may be handled exclusively as NetWorker media. Backup volumes that appear unlabeled are considered blank and are automatically labeled, mounted, and overwritten with new data. NetWorker uses backup volumes that it considers unlabeled only if it cannot locate a writable volume or a volume ready for recycling.

NetWorker does not recognize the following and considers these volumes unlabeled:

- Volumes without a label.
- Volumes labeled with something other than a NetWorker label.
- Volumes with NetWorker labels that were written in a density different than the device in which it is currently loaded. For example, if you performed NetWorker backups on an older 8 mm tape drive, a newer 8 mm tape drive may not be able to read the volumes from the older device because a different density was used to write the data.

Caution: Be careful about sharing NetWorker volumes between different backup devices if you have Auto media management enabled. You could potentially relabel and overwrite valuable data.

The Auto media management feature is located in both the Jukebox and Devices windows. For jukeboxes, you only need to enable Auto media management in the Jukeboxes window. NetWorker does not allow you to enable Auto media management in the Devices window if the device is located in a jukebox.

Note: If you back up to a standalone device and want to use these automated features, see “Using the Auto Media Management Feature” on page 99.

When Auto media management is enabled, NetWorker assumes that the jukebox is exclusive to NetWorker and handles the backup volumes loaded in the device as NetWorker media. Do not enable this feature if the jukebox or device is a shared resource, because NetWorker assumes it can label and overwrite any media it considers unlabeled. However, if the jukebox is a shared device, you may use the available slot range to restrict the volumes NetWorker uses for backup. If you place or store in the jukebox other backup volumes that you do not want NetWorker to use, place them in slots that are not included in the available slot range.

Operating Jukeboxes

This section contains information for performing jukebox operations, using other choices located in the Media menu. You may need to perform some of these tasks daily and others less frequently, depending on a variety of factors. For example, the frequency with which you perform these tasks depends upon the number of backup volumes you label at one time, how often you change the contents of your jukebox, and how often you need to manually mount volumes for backup.

Jukebox Mounting Dialog Box

To mount a backup volume in a jukebox, you must first make sure the correct jukebox device is selected in the NetWorker Administrator window. This is especially important if you have the NetWorker server configured to use more than one device.

Open the Jukebox Mounting dialog box (shown in Figure 7-7) by selecting a jukebox from the Devices display list in the NetWorker Administrator window, then clicking the Mount speedbar button or choosing Mount from the Media menu.

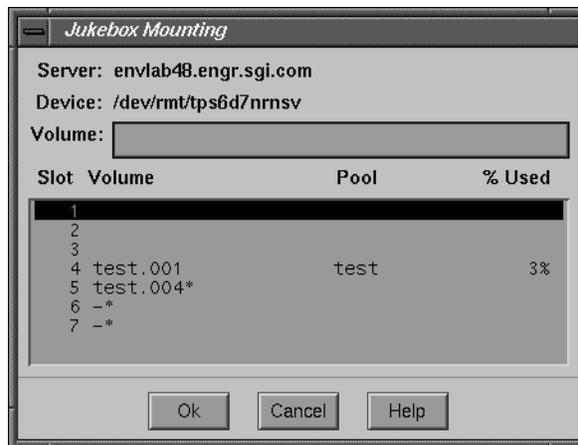


Figure 7-7 Jukebox Mounting Dialog Box

The Jukebox Mounting dialog box displays the name of the NetWorker server and the system *pathname* of the selected jukebox device. The Volume field displays the name of the currently selected backup volume.

The scrolling list displays the slot numbers (Slot), the label of the volume stored in the slot (Volume), the pool (Pool) to which the volume belongs, and the percentage of the capacity filled on the backup volume (% Used). Use this information to determine which volumes to mount in the jukebox. You may not want to mount full volumes or volumes near capacity. To mount the selected backup volume in the jukebox, first select it from the list, then click the Ok button.

Tip: Although jukebox devices are identified by (J) next to their pathname in the NetWorker Administrator window, it is a good idea to place an adhesive label on the jukebox identifying its device pathname(s). This is to remind you which device pathname(s) belong to the jukebox. This practice is especially helpful with multiple jukeboxes

Jukebox Labeling Dialog Box

Use the Jukebox Labeling dialog box to make the necessary choices for labeling the backup volumes in a jukebox. Select a backup device located in a jukebox from the Devices display in the NetWorker Administrator window. Then click the Label speedbar button, or choose Label from the Media menu. The Jukebox Labeling dialog box appears, as shown in Figure 7-8.

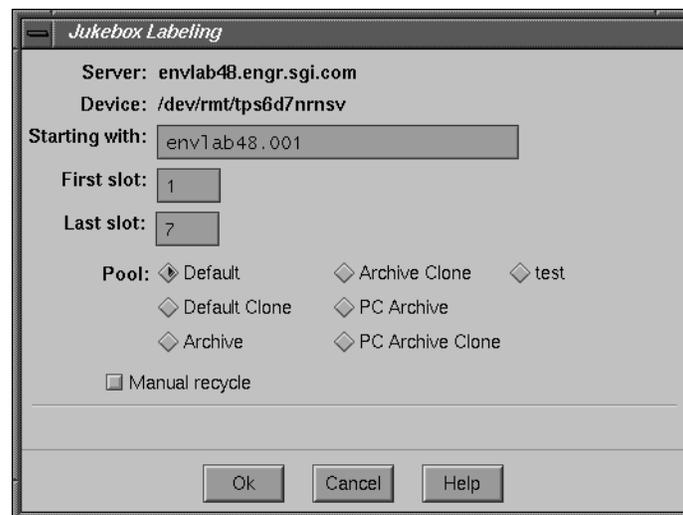


Figure 7-8 Jukebox Labeling Dialog Box

The Jukebox Labeling dialog box displays the name of the NetWorker server and the system *pathname* of the selected jukebox device. NetWorker starts labeling the media in the jukebox with the label displayed in the Starting with field.

The First slot and Last slot fields display the range of slots containing backup volumes that NetWorker labels according to the label template selected in the Pools window. Use these fields to control the range of slots in the jukebox for labeling or updating backup volumes.

Starting a Labeling Session

NetWorker uses the label template currently selected in the Pools window to label backup volumes loaded in the jukebox. Choose a label template from the Pools window or create your own before starting a labeling session.

Tip: It takes NetWorker some time to label the volumes in a jukebox. You may want to perform the labeling process at a time when you do not need to back up or recover files.

To label the backup volumes in a jukebox, follow these steps:

1. Select the jukebox device in the Devices display in the NetWorker Administrator window. A jukebox device is identified by (J) next to its *pathname*.
2. Open the Jukebox Labeling dialog box by selecting the device in the Devices display, and then choosing the Label command from the Media menu or clicking the Label speedbar button.
3. Override the range of slots specified by the label template, if desired, by entering slot numbers into the First slot and Last slot fields in the Jukebox Labeling dialog box.
4. Select a pool from the Pool choices for the volumes being labeled in the jukebox. NetWorker uses the label template associated with that pool to label the volumes.

Note: Only the pools enabled in the Pools window appear in the Jukebox Labeling dialog box.

5. Select the Manual recycle toggle button only if you want to manually recycle the backup volume.

If you select Manual recycle, any previous browse and retention policies chosen for the volume and its save sets will be overridden and the volume must be recycled manually. Use this option if you plan on backing up important data to either cloned or noncloned volumes, and you want to decide when the data and the volumes are ready for recycling.

6. Click the Ok button to begin the labeling process. A confirmation message appears, as shown in Figure 7-9.



Figure 7-9 Notice to OK Jukebox Labels

When a valid NetWorker label already exists on the media you are trying to label, a confirmation message appears. The confirmation is important because it keeps you from accidentally relabeling the media. When a backup volume is relabeled, its contents under the previous label cannot be recovered. When you select OK to confirm, the volumes in the slots are loaded, labeled, and unloaded.

If you label a specific *range* of backup volumes, the Starting with name *must* match the label template. If you label a *single* backup volume, you can use *any* name—it does not have to match the label template. The First and Last slots must be identical when labeling a single volume.

The following examples illustrate labeling a specific range and a single volume.

First example: If you previously labeled six tapes and need to label six more, your Jukebox Labeling dialog box might look like Figure 7-10.

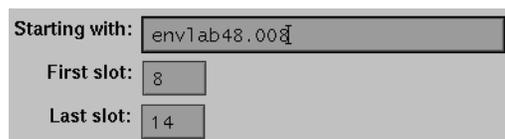


Figure 7-10 Jukebox Labeling More Tapes

Caution: If you remove backup volumes from the jukebox, remember there may still be one or more volumes left in the jukebox devices. Make sure to unmount them by selecting the jukebox devices, then clicking the Unmount speedbar button or choosing Unmount from the Media menu. This prevents you from accidentally relabeling these volumes when you reload the jukebox with new volumes.

Second example: If you label a single backup volume, the First slot and Last slot must be identical, as shown in Figure 7-11.

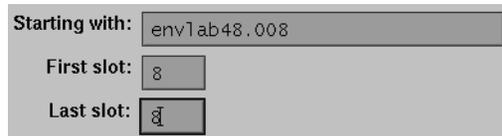


Figure 7-11 Jukebox Labeling One Tape

The Jukebox Labeling dialog box gives you control over volume labeling. You can use a label template to generate names as needed or you can specify a single unique name.

For more information on labeling volumes, see “Using Label Templates” on page 150.

Labeling With Bar Codes

Labeling media with bar code labels provides two distinct advantages: it significantly speeds up volume inventory and provides improved labeling accuracy.

Inventorying with bar code labels is fast and efficient because you do not have to load the volumes into the device. Instead, the jukebox 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 a carousel cartridge.

Bar code labels also provide greater labeling accuracy because the labels are attached to the media prior to being loaded and scanned in the jukebox. Once the jukebox scans the bar code label, NetWorker records and tracks the label in the media index.

Note: NetWorker uses *bar code labels* for only one reason: to inventory backup volumes. NetWorker uses the *volume label* (normally created with a label template) for identifying the volumes required for backup and recovery. However, NetWorker displays both the bar code label and the volume label in the Pending display of the NetWorker Administrator window when requesting a volume. You may also open the Volumes window to view both the volume label and its associated bar code label.

A backup volume must have a volume label, but does not require a bar code label. Use the Jukeboxes window and the Jukebox Labeling dialog box to recognize bar code labels.

To label your NetWorker volumes with bar code labels, follow these steps:

1. Properly apply the bar code labels to your backup volumes.
2. Place the volumes with the bar code labels in the jukebox.
3. Choose Jukeboxes from the Media menu. The Jukeboxes window appears.
4. Click the Yes button to enable the Bar code reader.
5. If you want the volume label (normally created from label templates) to match the bar code label, click Yes in the Match bar code labels choices. If you do not want the volume label to match the bar code label, click No.

In either case, two labels result: the bar code label, and the volume label (which may or may not match the bar code label depending upon the Match bar code labels choice you made).

6. Open the Jukebox Labeling dialog box by choosing Label from the Media menu.

If you chose Yes for the Match bar code labels choices in the Jukeboxes window, the Starting with field in the Jukebox Labeling dialog box displays Bar code label, as shown in Figure 7-12.

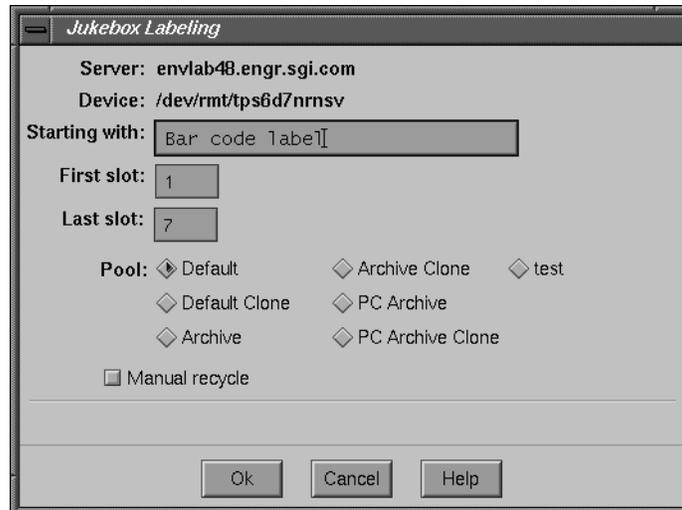


Figure 7-12 Jukebox Labeling for Bar Codes

If you chose No for the Match bar code labels choices in the Jukeboxes window, the Starting with field in the Jukebox Labeling dialog box displays a label template name, as shown under “Jukebox Labeling Dialog Box” on page 177.

Once you select your bar code label choices and label the backup volume, both volume names and bar code labels for the volume appear in the Volume Management window, as shown in Figure 7-13.

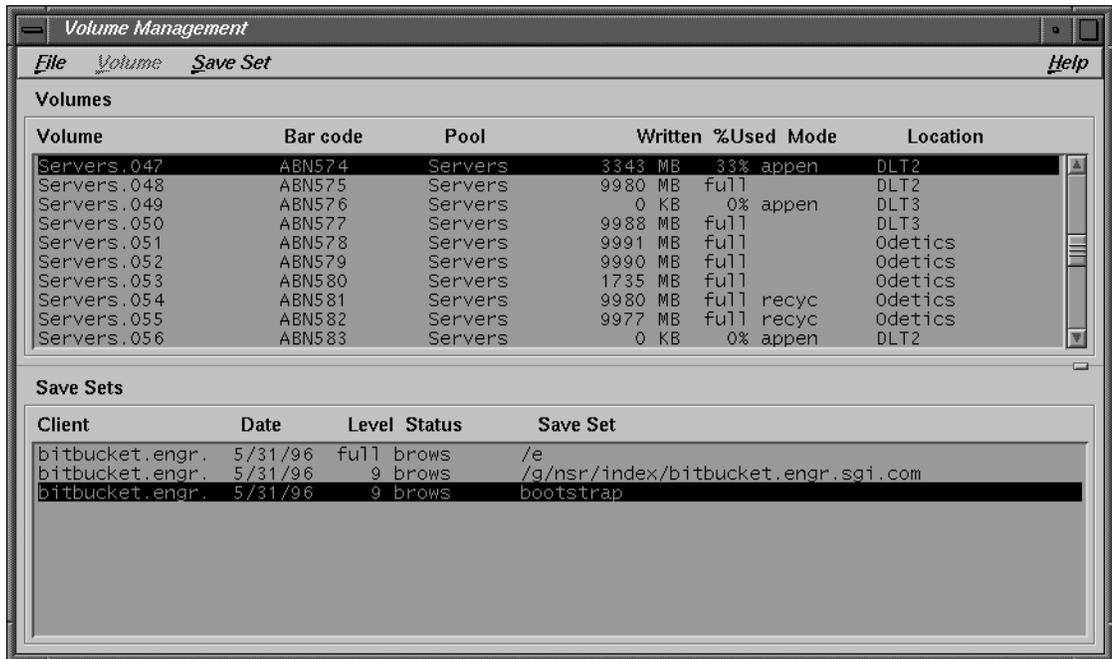


Figure 7-13 Volume Management Window

Understanding Labeling With Bar Codes

The volume label and the bar code label are both pieces of information that NetWorker uses to identify backup volumes. Both labels are recorded in the media index. However, the volume label is also recorded internally on the media. NetWorker uses bar code labels for inventorying volumes and uses the volume labels to identify the volumes needed for backup and recovery. You can require the volume and bar code label to match, by making appropriate selections for bar code labeling in the Jukeboxes window.

Tip: Do not use identical bar code labels for any of your NetWorker 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.

If you try to label a second volume with an identical bar code label, and you have enabled Matching bar code labels in the Jukeboxes window, NetWorker displays an error message and does not allow you to label the second volume. The error message identifies the jukebox slots containing the two volumes with identical labels and the bar code number. To correct the problem, you must either apply a different label and begin the labeling process again, or you must disable Match bar code labels in the Jukeboxes window while labeling the second volume.

Caution: If you choose not to match the volume label with the bar code label, it is advisable to create and attach volume labels on the outside of media volumes.

If you choose to have the labels not match, we suggest you label the volumes in this manner:

1. Attach all of the bar code labels to the media, then load the volumes in the jukebox.
2. Begin the labeling process. Click Yes for the Bar code reader choice, and click No for the Match bar code labels choice in the Jukeboxes window. NetWorker uses the next available label from the label template for the volume name. NetWorker labels the volumes and records both labels in the media index.
3. Open the Volumes window, after NetWorker completes the labeling process, to determine the volume label and bar code label for each volume. Create your own volume labels to attach to the backup volumes.

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 attach the correct volume labels to the volumes with bar code labels.

If you click Yes for the bar code choices and forget to attach a bar code label, you receive an error message informing you that there is no bar code label for that volume.

It is not necessary to label existing backup 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. However, if you have volumes you use often for recovery or for overwriting with new data, it is beneficial to label them with bar code labels. Labeling these volumes with bar codes saves you hours of time inventorying your backup volumes.

If you decide to use bar code labels on your existing backup volumes, you must first apply the bar code labels to the existing volumes. Then, load and mount each volume individually so NetWorker can match the bar code label with the existing volume label.

Tip: You may purchase a variety of bar code labels from a third-party vendor. You can choose numeric, alphanumeric, or a special combination of numbers and characters to meet your labeling needs. You can even 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*.” Labeling instructions for bar code labels are usually provided with your jukebox hardware documentation. If you have questions about bar code labels, contact the hardware manufacturer.

Using a consistent labeling scheme helps you better organize and track your backup volumes. It also aids the inventory process if all the volumes use bar code labels rather than just a limited number of them.

When NetWorker relabels volumes automatically, it reuses the original volume label name. You can only change the label name if you relabel the volumes manually. NetWorker 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.

Jukebox Inventory Dialog Box

Use the Jukebox Inventory dialog box to enter the range of slots you want NetWorker to inventory in your jukebox. It is a good idea to inventory backup volumes in a jukebox when volumes have been moved or added. To bring up the Jukebox Inventory dialog box, choose Inventory from the Media menu.

Inventorying Backup Volumes

To inventory backup volumes in a jukebox, follow these steps:

1. Load the new backup volumes into the jukebox.
2. Select the jukebox device in the Devices panel of the NetWorker Administrator window. Jukebox devices are marked with (J).
3. Choose Inventory from the Media menu.

The Jukebox Inventory dialog box appears, as shown in Figure 7-14.



Figure 7-14 Jukebox Inventory Dialog Box

4. Enter the range of slots you want to inventory in the First and Last Slot fields.
5. Click Ok to begin the inventory. The Jukebox Inventory dialog box displays a message showing the progress of the inventory operation (see Figure 7-15). Since NetWorker mounts each backup volume, reads its label, and unmounts it, the inventory process can take some time to complete.



Figure 7-15 Jukebox Inventory Finished

You may cancel the inventory operation after NetWorker completes inventorying the current volume by clicking Cancel in the Jukebox Inventory dialog box. Click Cancel to close the window when the inventory is complete.

After completing an inventory, NetWorker registers the contents of the jukebox and then proceeds with its network-wide backup and recover services.

Disabling the Element Status Feature

Your jukebox may support the element status feature. This feature, along with the bar code label feature, speeds up the inventory process by eliminating the need to load and read volume labels. NetWorker uses the element status feature to determine whether or not a slot contains a volume. If the volume has a bar code label, NetWorker reads it without loading the volume.

If you perform an inventory on a jukebox containing volumes without bar code labels, NetWorker must load the volume to read the internal label. In this case, the element status feature is not beneficial.

Check the documentation from your jukebox manufacturer for information on features for your jukebox.

If you are not using bar code labels and want to disable the element status option for your jukebox, follow these steps:

1. Choose Jukeboxes from the Media menu in the NetWorker Administrator window.
2. Choose Details from the View menu in the Jukeboxes window.
3. Use the scroll bar to move down and make the Jukebox options field visible.
4. Enter the following option in the Jukebox options field, then click add to add it to the Jukebox options list and disable the element status option:

`!element_status`

The option is added to the Jukebox options list, as shown in Figure 7-16. Click Apply to save the new jukebox settings, or click Reset to restore the previous settings.

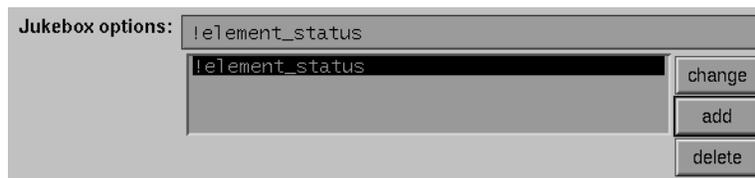


Figure 7-16 Jukebox Options Field

To restore the element status feature, select the setting from the Jukebox options list and click delete, then click Apply.

Understanding the Inventory Process

When NetWorker labels the contents of a jukebox, it registers the location of the backup volumes in the jukebox slots when it assigns the volume label. As long as you do not change the volumes in the jukebox after labeling them, NetWorker can access the volumes because each volume label is assigned to a specific slot.

However, if you change the contents of the jukebox without performing the labeling process, or if you move volumes into new slots, you must inform NetWorker that the jukebox 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 jukebox, NetWorker reads the label of each volume and records its slot number. For example, if you have more than one job pack for a jukebox, you must inventory each time you remove one job pack and load another one into the jukebox, if you do not label the volumes in the new job pack.

NetWorker permits devices to read bar code labels to speed up the inventory process. Silicon Graphics strongly recommends that you use bar code labels if you have a large number of backup volumes or change your jukebox contents often. See “Labeling With Bar Codes” on page 180 for more information.

Checking Jukebox Notifications

NetWorker uses electronic mail to send notices about NetWorker events. The software uses the three Tape mount request notices to inform the system administrator that the jukebox needs attention. The notification Tape mount request 1 is undefined so you can create your own notification message.

The following situations require attention:

- The backup volumes in the jukebox are 90% full.
- The jukebox needs more backup volumes to continue.
- The jukebox has a mechanical problem.
- The jukebox device needs cleaning.
- The cleaning cartridge in the jukebox needs replacing.

Choose Notifications from the Customize menu to open the Notifications dialog box. The Notifications window contains tape mount and device cleaning request notifications pertaining to jukebox operations. See “Setting Up Event Notification” on page 164 for more information.

NetWorker automatically mounts the correct volume so that the recovery proceeds as long as the volume is loaded in the jukebox. If NetWorker requires backup volumes for a recovery that are not loaded in the jukebox, you receive a notice in the Pending window of the NetWorker Administrator program.

After you correct a jukebox problem, you may need to mount a backup volume before continuing backup or recovery. Check the Pending display in the NetWorker Administrator window for messages.

Suggestions for Operating Jukeboxes

This section provides additional suggestions to help you use your jukebox and backup media effectively and reliably.

Recycling Versus Adding More Backup Volumes

NetWorker can save on volumes marked appen (appendable) in the Volume Management window. If the volumes inside a jukebox are marked full, they cannot receive additional backups. You can do one of the following with volumes marked full:

- If you need to keep the volumes for long-term storage, remove the full volumes and replace them with new media.
- If you do not need data on the full volumes, you can manually change the mode to recyc in the Volume Management window and NetWorker will overwrite data with new backups, yet maintain the existing labels. This is the only instance where you do not need to relabel a volume to make it eligible to be overwritten with new data.

Note: The mode of a volume is automatically changed to recyclable when all of the save sets on the volume have passed the time period specified by its retention policy.

There are advantages to both recycling media and adding more media to a pool. By recycling, you reuse the same backup 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 volume pool instead of recycling it. For example, a jukebox may need new backup volumes every three months if the company policy is to maintain the backups for a year. In this case, you have to keep adding new media to the pool until you can recycle the backup volumes that contain expired or old backups.

Using the Jukebox Cartridge Access Port

A cartridge access port (CAP) allows you to deposit and withdraw media cartridges in a jukebox without opening the door to the jukebox. Each time you open the door of a jukebox to add or remove media, you invalidate the status of the jukebox. You then need to reinventory the contents of the jukebox so that NetWorker can track the backup media. The inventory process can take a long time to complete.

This CAP feature is useful because you can add (deposit) and remove (withdraw) backup volumes in a jukebox without having to reinventory the jukebox.

When you use the CAP to deposit or withdraw backup volumes, NetWorker does not automatically take inventory, read bar code labels, or locate empty slots in the jukebox. Use the Jukebox Inventory dialog box and Jukeboxes window for these tasks.

Depositing a Cartridge

To deposit a cartridge using the CAP, follow these steps:

1. Press the button on the front of the jukebox to move the cartridge holder forward and open the CAP.
2. Place the cartridge in the holder.
3. Press the button again to move the media into the jukebox and close the CAP.
4. At the NetWorker server, enter the **nsrjb** command at the system prompt, replacing *slot* with the slot number for the cartridge, and *volume* with the label name of the backup volume:

```
nsrjb -d -s slot volume
```

If you do not specify a volume name with the **nsrjb** command, NetWorker marks the slot with a hyphen and asterisk (-*) to indicate that the volume in that slot is unknown.

You need to inventory the slot with this command:

```
nsrjb -I -S slot
```

To verify that the backup volume was deposited in the correct slot, click the Mount button in the NetWorker Administrator window.

Withdrawing a Cartridge

To withdraw a backup volume from a specific slot in the jukebox, follow these steps:

1. At the NetWorker server, enter the **nsrjb** command at the system prompt, replacing *slot* with the slot number for the cartridge, and *volume* with the label name of the backup volume:

```
nsrjb -w { -S slot | volume | slot volume }
```

2. When the media cartridge reaches the CAP, press the button to open the CAP.
3. Remove the backup volume and close the CAP.

Using Pools With a Jukebox

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 in the Jukebox Labeling dialog box. This ensures that your full backups are located in a consecutive range of slots in the jukebox, which makes it easy for you to remove all of the volumes at the same time.

Calibrating the Devices

Check the jukebox manufacturer's documentation for information about the frequency and method for calibrating the loading mechanism for the jukebox device.

Archive Application

This chapter provides information specific to configuring and operating NetWorker with the Archive Application. It also includes explanations that compare and contrast the different methods used for protecting data across a network.

The NetWorker archive server provides file archiving and retrieval services to a range of client machines. It is packaged as an add-on extension to existing NetWorker backup servers, and uses the same license mechanism as NetWorker.

The NetWorker archive client can be any machine on a network that employs archive services provided by the archive server. Clients may be enabled for backups, for archives, or for both.

Data archiving is the process of taking a snapshot of files or directories as they reside on primary media (usually disk) at a given point in time. The snapshot image is typically stored on removable media, such as tape or optical disc. Once the snapshot is safely stored on removable media, related files may be deleted to conserve space on disk.

Navigating the Windows

The following sections explain the features and use of the various Archive windows:

- Clients window—to enable archive services for a client.
- Archive Requests window—to schedule the archiving of a client.
- Archive Request Control window—to monitor scheduled archive requests.
- Archive Request Details window—to see the progress of a recent archive.

Clients Window

Before archiving can occur, you must configure NetWorker to recognize archive clients. Choose Client Setup from the Clients menu to open the Clients window (see Figure 8-1).

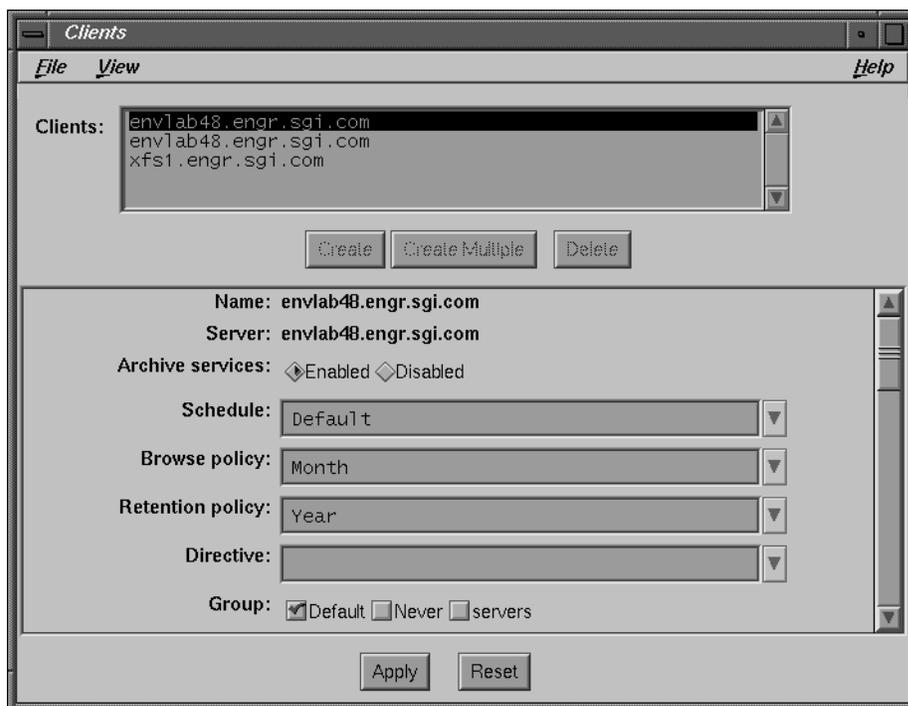


Figure 8-1 Enable Archive in Clients Window

The Archive services choices enable or disable archives for the currently selected client in the Clients list. To allow archives for the client, click the Enabled choice for Archive services. If archive services remain Disabled, the client will not be able to perform an archive. When you enable archive services for one client, you also enable other clients of the same name on that server.

Find the Archive users list at the bottom of the scrolling section of the Clients window. To allow users on the client to perform manual archives, enter their *username* into the Archive users list. To schedule an archive request of an entire workstation, *root* (or equivalent) must be on the Archive users list for that client, or *root@client* must be in the Administrator list for the server.

For a complete description of the Clients window, see "Navigating the Clients Window" on page 23.

Archive Requests Window

To schedule an archive on a one-time basis, choose Archive Requests from the Customize menu. The Archive Requests window appears, as shown in Figure 8-2.

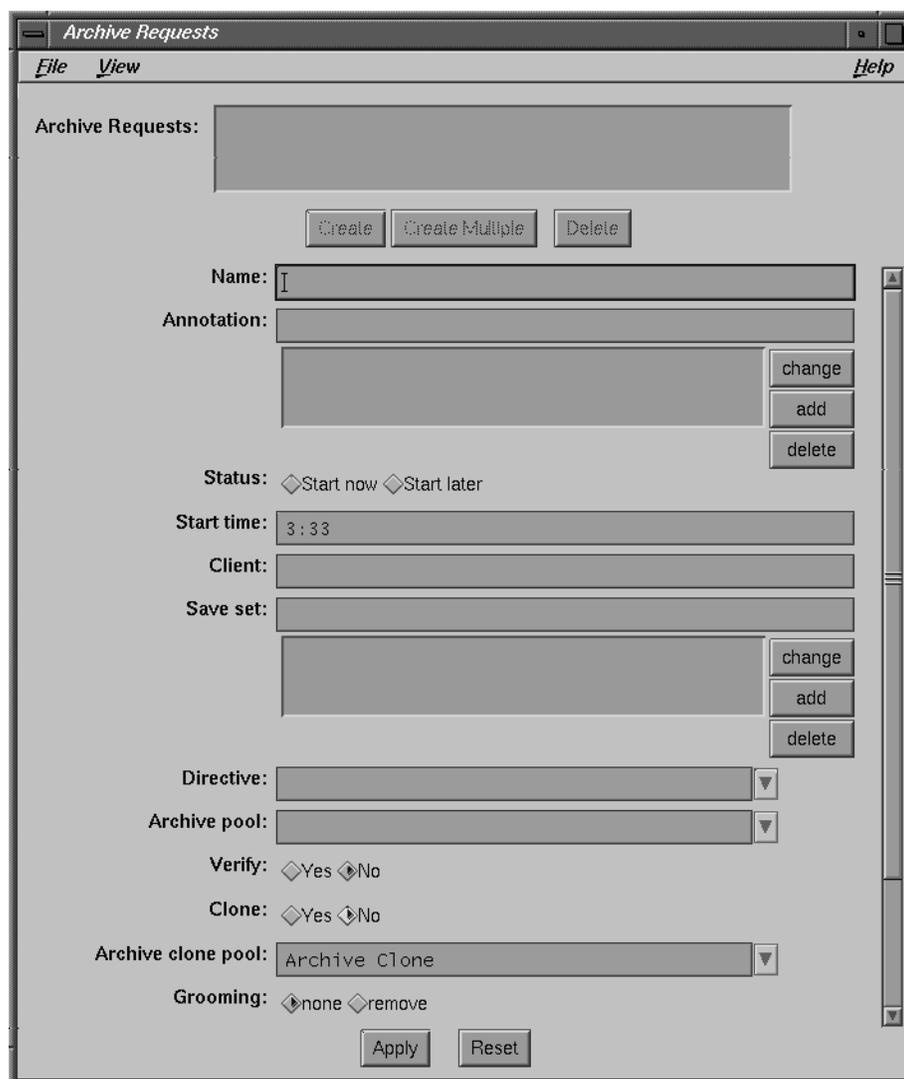


Figure 8-2 Archive Requests Window

The Archive Requests scrolling list displays archives previously requested on this server. You can add new archive requests and delete old ones using the Create and Delete buttons.

The scrolling panel displays the following information:

- Name—displays the archive name.
 - Annotation—a mixed-case comment string that you provide for every archive request (limited to 1024 characters).
 - Status—indicates when to begin the archive. Start now means archiving begins when you click the Apply button. Start later means archiving begins as indicated in the Start time field.
 - Start time—gives the next time to begin, based on a 24-hour clock.
 - Client—displays the *hostname* for the archive client. To request an archive for the server, enter its name in the Client field.
 - Save set—specifies *pathnames* of directories or files to archive.
 - Directive—specifies the backup method, usually Unix standard or Unix with compression.
 - Archive pool—specifies the volume pool to which archives should be sorted. The default volume pool is Archive.
 - Verify—indicates whether or not to automatically check the integrity of data archived on tape.
 - Clone—indicates whether or not to automatically clone the archive volumes for extra security.
 - Archive clone pool—specifies the volume pool for archive clones. The default archive clone pool is Archive Clone.
 - Grooming—indicates whether or not client files and directories should be removed after archiving is complete.
- Caution:** If you use directives that include instructions to skip files, do not enable the Grooming option. Grooming occurs after a file has been archived. If a file is skipped, it cannot be groomed and will cause the archive to fail.
- Archive completion—contains an optional command to execute after archiving is complete, for example */usr/ucb/mail*.

Archive Request Control Window

To see the results of an archive request, or to start or disable a scheduled archive, choose Archive Request Control from the Server menu. The Archive Request Control window appears, as shown in Figure 8-3.

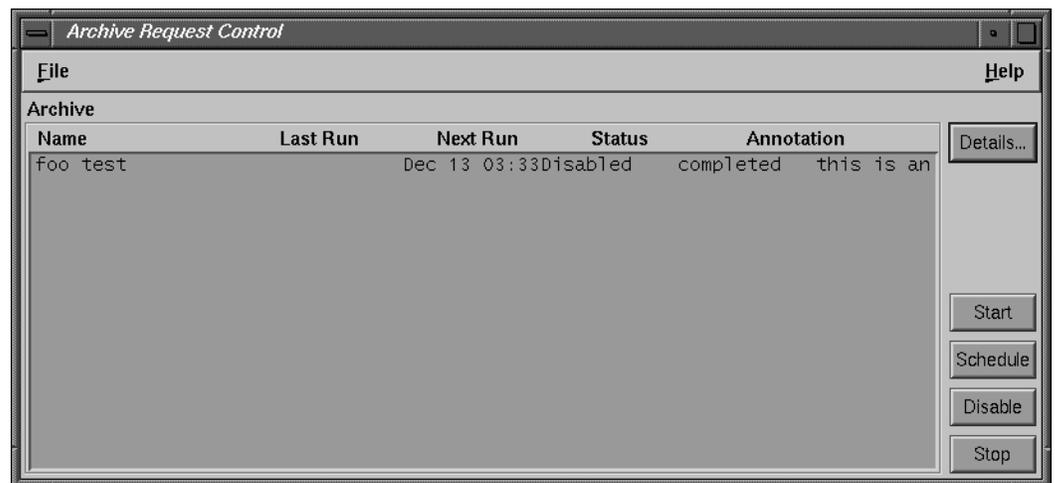


Figure 8-3 Archive Request Control Window

The most recently run archive request is initially highlighted in the Archive list. Usually it is Disabled, since scheduled archives run only once. To see details about when this archive completed and how successfully it ran, click the Details button.

If you want details about, or control over, a different archive request, select the one you want from the Archive list.

To initiate an archive request immediately, click the Start button. This has the same effect as selecting a Status of Start now for that archive request in the Archive Requests window.

To reschedule an archive request, click the Schedule button and enter a new starting time. This has the same effect as selecting Start later and specifying a new Start time in the Archive Requests window.

To disable an archive request that you scheduled here or in the Archive Requests window, click the Disable button. To halt an archive in progress, click the Stop button.

Archive Request Details

To see the progress of a recent archive (the one currently selected), click the Details button inside the Archive Request Control window. The Archive Request Details window appears, as shown in Figure 8-4.

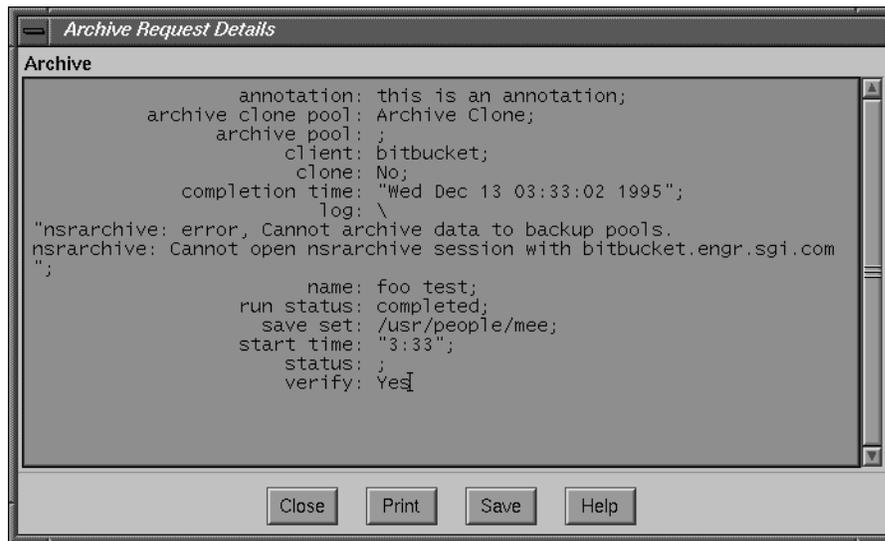


Figure 8-4 Archive Request Details Window

This window provides information about the completion of an archive request. Some of the information in this window, such as annotation, archive clone pool, archive pool, client, clone, name, save set, start time, status, and verify, also appears in the Archive Requests window.

The completion time field displays when the archive finished. Its duration is the difference between this and start time. The log field shows messages generated by the archive. The run status field shows the outcome of the archive request, either completed, failed, or partial.

Tip: For more information about failed archives, see the log file, usually */nsr/logs/daemon.log* on the server.

Archive Example

Suppose you must shut down and remove the workstation (and *hostname*) of someone who has left the company. It would be wise to archive the system data first, in case its filesystems contain essential files you need to access later.

This section gives an example of how to schedule and run an archive request.

Creating an Archive Client

Only registered archive clients can use the archive facility. To create an archive client, follow these steps:

1. Choose Client Setup from the Clients menu. The Clients window appears, as shown on page 191.
2. Click the Create button; the Clients window changes.
3. Enter the *hostname* of the workstation in the client Name field.
4. Click Enabled after Archive services to allow archives for this client.
5. If you want to permit users on the archive client machine to use archive and retrieve, scroll to the bottom of the Clients window and add their user names to the Archive users field.

(It is unlikely that you would want to allow manual archives and retrieves on a workstation about to be shut down.)

Machine *hostname* is now a registered archive client. However, an archive will not take place until you request one.

Tip: If you want to allow archives on the server, make sure that archives are enabled for the server as a client of itself.

Making an Archive Request

Valid archive users may request archives manually using the **nwarchive** command. However, a manual archive often takes a long time. To avoid overloading a busy network with an archive request during the day, schedule it late at night when the network has less traffic.

To make an archive request, follow these steps:

1. Choose Archive Requests from the Customize menu.
The Archive Requests window appears, as shown in Figure 8-2.
2. Click the Create button; the Archive Requests window changes.
3. Enter the Name you want to assign to this archive request, and a brief Annotation to remind you of the purpose for the archive.
4. Click Start later for the Status choice to schedule this archive for that night. Enter the Start time you want, or accept the default starting time of 3:33 a.m.
5. Enter the archive client machine *hostname* in the Client field, and the *pathname(s)* you want to archive in the Save set field.
6. Specify a custom Archive pool, or accept the default Archive volume pool.
7. Click Yes for the Verify choice to check that the archived data was saved correctly.
8. If you want to make a duplicate copy of this archive volume, click Yes for the Clone choice. Accept the default Archive Clone pool.
9. To remove files and directories from disk after archiving them to tape, click remove for the Grooming choice.
10. To be notified when the archive completes, enter a command into the Archive completion field, for example:

```
/usr/ucb/mail -s archive_request admin@host
```
11. Click the Apply button to activate your changes.

You have now requested an archive of a client machine *hostname* to begin that night at 3:33 a.m.

Checking the Archive Request

The next morning, check the outcome of the archive. If you set up an Archive completion notice, look for an e-mail message containing a log of the archive request.

To check details of the archive, follow these steps:

1. Choose Archive Request Control from the Server menu.
The Archive Request Control window appears, as shown in Figure 8-3.

2. Click the Details button.

The Archive Request Details window appears, as shown in Figure 8-4.

3. If this window shows the archive completed successfully, you can safely reconfigure the ex-employee's machine.

If the archive failed, you can reschedule it (see the following section.)

You may use the Archive Request Control window to start, schedule, disable, or stop another archive.

Rescheduling the Archive Request

Suppose that the archive of *hostname* did not complete last night because, for example, an energy-conscious employee turned off the computer. You decide to reschedule for the next night.

To reschedule the archive, follow these steps:

1. Choose Archive Request Control from the Server menu. The Archive Request Control window appears, as shown on page 195. Make sure that the archive request you want is highlighted.
2. Click the Schedule button. The Archive Request Schedule window appears, as shown in Figure 8-5.



Figure 8-5 Archive Request Schedule Window

3. Enter a new starting time in the Schedule Archive field, using the 24-hour clock, and click Ok.

The archive request executes again that night, at the time you specified. If you change your mind and want to discontinue the archive, click the Disable button in the Archive Request Control window.

Clone and Verify

NetWorker contains two preconfigured volume pools for use with archiving: Archive and Archive Clone. If you want to create a new volume pool for archives, follow these steps:

1. Choose Pools from the Media menu.
2. Click the Create button in the Pools window.
3. Fill out all the fields according to your needs.

Tip: Make sure the Pool type is set to Archive, and the Store index entries field is set to No. These two traits distinguish archive pools from backup pools.

4. Click the Apply button to activate the new volume pool.

Now when you schedule a new archive request, you may use the new archive volume pool you created. If you choose to clone archive data, also create a new archive clone pool. NetWorker will write archive data only to an archive volume and archive clone data only to an archive clone volume.

If you have already made an archive and want to make a clone of it, follow these steps:

1. Choose Clone from the Save set menu.
2. Enter criteria for locating save sets in the Save Set Clone window.

Tip: Click the More button and type **Archive** in the Pool field.

3. Click the Query button to see save sets matching your criteria.
4. Select the save set you wish to clone and click the Clone button.
5. Click the Start button in the Save Set Status Clone window to activate the save set clone.

To verify data already archived on an archive volume, you have two alternatives:

- Clone the archive data by choosing Clone from the Save set menu. During cloning, the original archive data will be verified as the save set is copied from one volume to another. When you are done, you may reuse the cloned volume.
- Determine the save set ID of the archived data, for example by searching the NetWorker Archive Retrieve window. Then run the following command at the system prompt, substituting the save set ID of the archived data for *ssid* (this is how NetWorker verifies archived data):

```
# nsrretrieve -n -S ssid
```

Both alternatives verify the integrity of the data, but do not actually compare archived data with data on disk. Chapters 6 and 7 in this manual provide more information about volume pools and save set cloning.

Archiving Shortcut

Shortcut: To schedule and run an archive request, follow this general procedure:

1. Create and enable an archive client using the Clients window.
2. Using the Archive Requests window, fill in all or most of the fields with your preferences.
Name and Client are mandatory, as is the Start time if you Start later. Save set defaults to empty, Directive and Archive pool to the default, Verify and Clone to No, and Grooming to none.
3. Check the archive status in the Archive Request Details window.
4. To reschedule an archive, bring up the Archive Request Schedule window. To discontinue a scheduled archive request, click the Disable button. To start and stop an archive request (for example, to test it), click the Start and Stop buttons.

Understanding the Archive Feature

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 of differential saves, or incremental saves.

Note: Archives are not recorded in the online file index, so they are not affected by the browse policy. This feature helps conserve disk space.

Retrieve is similar to recover, except that it works with archive save sets instead of backup save sets. Since archived files are not recorded in the file index, the user interface for retrieve is based on save sets, rather than on a directory hierarchy.

Users on the NetWorker administrator list have permission to configure archive services. These users, and users registered on the Archive users list in the Clients window, have permission to use the archiving and retrieval facilities. Registered users may archive any file for which they have read permission.

Anyone can browse archive save sets—that is, look at information in the media index. However, you may retrieve only those files that you own, unless you are *superuser*, in which case you may retrieve any file.

Tip: If you want to overwrite an archive tape, first make sure nobody will ever need the data again. Then relabel the volume, as you would a backup volume.

Archive Functions

This section describes three archive functions: archiving, verification, and retrieval.

Data Archiving

Data archiving can be performed by end users or by the system administrator. Registered clients can perform manual archives that start right away. System administrators can perform manual archives or schedule an archive to take place anytime during the next 24 hours. For example, the best time to perform a large archive might be in the middle of the night.

Both users and administrators can request an extra copy of their archive save set, called a *clone*. NetWorker Archive employs the volume pools feature to separate backup volumes from archive volumes and archive volumes from archive clone volumes.

After archiving is complete, users and administrators are given the option of deleting archived files and directories or leaving them in place. This option is called *grooming*. Grooming helps conserve disk space after a project is finished.

This chapter describes the Archive Application's graphical user interface. For details about the command-line interface, see the `nsrarchive(1M)` reference page.

Data Verification

Since archived files are often deleted from the system, NetWorker provides an extra measure of security to ensure archived data are correct. NetWorker verifies two ways.

- media verification—NetWorker checks the archive volume to ensure that it is writable and contains no bad spots.
- data verification—NetWorker reads data from the archive volume as if doing a retrieve, but does not actually write any archived data back to disk.

If a volume is suspect, or if there are problems with the data on the tape, NetWorker issues a warning and suspends grooming.

If you decide to groom files, we recommend you also select Verify or Clone to avoid deleting improperly archived files. Figure 8-6 shows the ordering of archive functions.

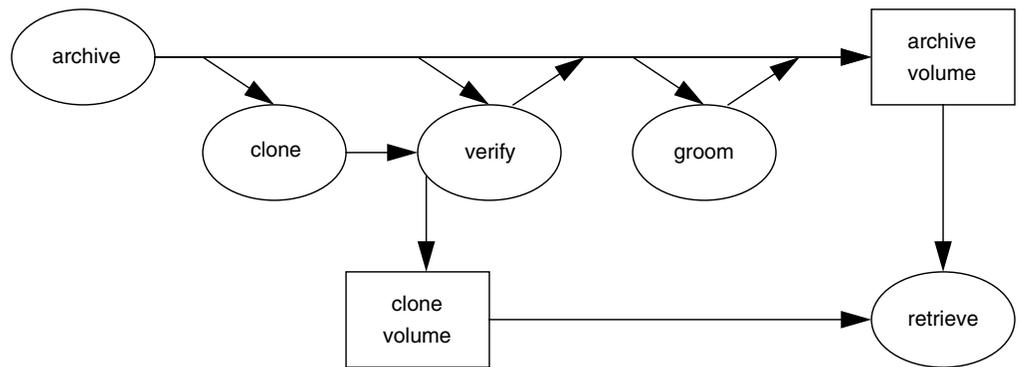


Figure 8-6 Archive Process Flow Chart

Data Retrieval

When you use NetWorker Retrieve, the NetWorker Retrieve window displays archived save sets for the selected server, listed by client name. You can retrieve a save set only if you have administrator or archive user privileges for that server, are the owner of files in the save set, or are *root*.

It is possible to search for specific archives and to alter the sort order of archive save sets in the viewing list. Refer to Chapter 5, "Archiving and Retrieving Files" in the *IRIX NetWorker User's Guide* for more details on retrieving archived files.

When the user picks an archive save set to retrieve, and the administrator ensures that the relevant archive volume (or a clone of that volume) is mounted on the NetWorker server, retrieval can begin.

Retrieved save sets can be relocated, renamed, or allowed to overwrite existing files of the same name, as with the NetWorker recover feature.

Methods for Protecting Data

Data backup is the process of storing copies of files and directories from local disk onto removable media, usually tapes. These copies can be recovered in case the original files are lost or damaged. The system administrator usually schedules daily backups. Any new files, or files that changed since the last backup, are copied to tape so they can be restored on disk if necessary.

Archiving is normally performed on data associated with specific projects, rather than on an entire system. Unlike data backup, end users usually archive their files as needed, so a network-wide archiving policy is not required. Archives, unlike backups, are not associated with a level (full, differential, or incremental).

When users archive project data, they can choose to delete the files from the system disk automatically to conserve space. In this case, archived files need to be placed on long-lasting archive media.

Hierarchical storage management (HSM) is a data management strategy where data is automatically migrated from one storage medium to another, based on a set of rules. The rule most often employed is *access rate*—the longer a file is inactive, the more likely it is to migrate.

Storage hierarchy is usually governed by the cost of storage for each media. The benefit of HSM is that it provides users with a seemingly infinite storage capacity, at the lowest possible cost.

The principal goals of backup, archiving, and HSM are as follows:

- The goal of backup is to protect data against accidental loss or damage. Backups should be reliable and efficient.
- The goal of archiving is to conserve online storage space. Storage media must be durable, safe, and reliable.
- The goal of HSM is to conserve network storage resources. Migration and recall must be automatic and reliable.

Recovering From a Disk Crash

This chapter contains procedures to follow to prepare for and recover from a major disk crash. The four different types of disk crash situations in this chapter are summarized below:

- The primary disk, which contains the operating system and NetWorker binaries, is damaged. This can apply to a client system or a NetWorker server.
- The secondary disk, which contains other filesystems, is damaged. This can apply to a client system or a NetWorker server.
- The NetWorker server disk, which contains the online indexes (the */nsr* filesystem), is damaged. You have to recover the indexes before using NetWorker to recover any filesystems.
- The NetWorker fileserver was destroyed. You have to recover everything to a new NetWorker server.

If a primary disk suffers a head crash, you may need to replace the disk, boot from mini-root, format and partition the disk, re-install the operating system and NetWorker binaries, and then recover affected filesystems. In this case, before using NetWorker to recover the data to the disk, consult the system administration manuals you used to set up your fileserver for the first time.

If a secondary disk suffers a head crash, its recovery procedure is simpler, since you do not have to reinstall the operating system and NetWorker binaries.

Preparing for a Crash

The ultimate disaster for a system is to lose all the files on its disk. Most sites back up their file servers daily as a preventive measure. If a primary system disk suffers a crash, you can rebuild its filesystems with NetWorker, after you reinstall the operating system (if necessary).

If the NetWorker server filesystem or disk that contains */nsr* and/or any contents of */nsr* linked to another location is destroyed, the recovery procedure involves an extra step; you must recover the online indexes for the server as well as the server filesystems. The */nsr* directory on the server contains one index for each client, including an index for the server as a client of itself.

If your NetWorker server was destroyed (in a fire, for example) you need to replace it with another machine. You may do this as long as you do the following:

- Name the replacement server with the same *hostname* as the original NetWorker server, if possible.
- Reinstall NetWorker using the same directory locations for the online indexes as in the original installation.
- Reregister the new NetWorker server.

Once you understand the procedure for a disaster recovery, make sure you have carefully thought of a disaster recovery plan for your site. If possible, test the ability to recover from a disaster at your site.

If you set up your network and enabled NetWorker for scheduled network-wide backups, you are well prepared for a disaster. Every time NetWorker backs up a group of clients, it also backs up all the online indexes for those clients. This backup includes the file index, media database, and NetWorker configuration files for the server itself (which contain entries for the client indexes, should these require recovery). The media database, NetWorker configuration files, and part of the server index are saved to a special save set named *bootstrap*. The save set identification numbers (*ssid*) for all recent bootstraps are sent to a default printer, providing a hard copy for your records.

Silicon Graphics recommends that you take these additional precautionary steps to help you recover from a possible crash:

- Keep a file containing hard copies of the bootstrap records. Place these daily sheets of paper in a three-ring binder or a file folder.
- Make a hard copy record of the disks, partition sizes, and mount points for the server and any clients that have a local hard disk. This information makes a future recovery much smoother for you.
- Save your enabler code certificate and the authorization code for the NetWorker server.

File the Bootstrap Information

NetWorker sends a record of the bootstrap save sets to your default printer, so you have a piece of paper with the dates, locations, and save set ID numbers needed for disaster recovery.

If you ever need to recover the server online indexes, the information on this piece of paper can save you a great deal of time. Save this information in a safe place.

The information sent to the printer looks similar to this:

```
December 13 03:35 1996 server's bootstrap information Page 1
  date      time      level      ssid  file  record  volume
12/12/96 18:23:23   full      12096   10     0    vol.001
12/13/96  3:34:46     9         12099   13     0    vol.001
```

NetWorker prints all the bootstrap save sets for the past month. The bootstrap save set may span more than one backup volume. The file and record numbers are used to find the associated save set quickly.

You can also manually back up the NetWorker server indexes by using the **savegrp -O** command. Using this command also sends the bootstrap information to a printer. For example:

```
# savegrp -O -c servername
```

To use the **savegrp** command, you must be *root* on the NetWorker server.

File the Disk Information

Use the disk information command (**df**) to find out how the NetWorker server disks are partitioned and mounted. Use the disk volume header command (**prtvtoc**), or the logical volume manager (**xlv_mgr**) to print disk partitioning information. Do the same for any NetWorker clients that have local hard disks. (On old SunOS[®] systems, use the **dkinfo** command to show disk partitioning. On AIX[®] systems, use the **lslv** command to show disk partitioning.)

Do this for each system NetWorker backs up, unless the systems are consistent in disk and filesystem layout. Print and file this information in case you ever have to recover from a disk crash.

For example, the **df** information looks similar to this:

```
% df -k
Filesystem                Type  kbytes    use    avail  %use Mounted on
/dev/root                  xfs   39287     29074   10213   75  /
/dev/dsk/xdv/user2        xfs  3826740  3735108   91632   98  /usr
/dev/dsk/xdv/e            xfs 35224840 34471480  753360   98  /e
/dev/dsk/xdv/d            xfs 1953920  1708676  245244   88  /d
/dev/dsk/xdv/g            xfs 35222600 32929604 2292996   94  /g
/dev/dsk/xdv/b            xfs 35228840 31598192 3630648   90  /b
/dev/dsk/xdv/h            xfs 35224840 33678664 1546176   96  /h
```

The following **prtvtoc** command example gives you information about how a root disk is partitioned for an IRIX system. The device name is the “raw” device corresponding to the device name used for the output from the **df** command.

```
# prtvtoc
Printing label for root disk
* /dev/rdisk/dks0d1s0 (bootfile "/unix")
* 512 bytes/sector
* 104 sectors/track
* 5 tracks/cylinder
* 5 spare blocks/cylinder
* 4019 cylinders
* 5 cylinders occupied by header
* 4014 accessible cylinders
*
* No space unallocated to partitions
Partition  Type  Fs  Start: sec  (cyl)  Size: sec  (cyl)  Mount
0          xfs  yes  2575  ( 5)  1905500  (3700)  /
1          raw
8          volhdr  0  ( 0)  2575  ( 5)
10         volume  0  ( 0)  2069785  (4019)
```

If a disk is destroyed in a head crash, you will be able to restore it and recover the filesystems to their original state, using the hard copy information from these disk information commands. At a minimum, you need to have partitions large enough to hold all the recovered data.

Using Save Set Recover or Recover

You have a choice between using the save set recover feature or the normal recovery procedure to recover filesystems after disk failure. This section describes the advantages of each method.

During backup, NetWorker multiplexes different filesystems simultaneously to the backup media. By recovering multiple filesystems at the same time, NetWorker has the opportunity to de-multiplex the filesystem save sets from the same backup volume in parallel, thus reading each backup volume only once. This is why you should recover multiple filesystems at the same time if it is practical for you to do so. You can do this by marking different save set points (using the save set recover feature) or different mount points (using the normal recover procedure).

With release 4.1 and higher, the NetWorker server can recover several save sets from the same backup volume in parallel, eliminating the need to read the same backup volume several times during a recovery.

An advantage to using the save set recover feature is that you spend less time browsing and marking filesystems. With normal recovery, an entry for each individual file is accessed in the NetWorker file index to reconstruct an accurate view of the filesystem. It takes time to “pick” the most recent versions of files from the tape. With save set recover, individual file browsing is bypassed and entire save sets are recovered in one step. If the browse policy has expired, save set recover is the only way to recover a file using the GUI. For information about using save set recover, see Chapter 5, “Recovering and Cloning Save Sets.”

Caution: Whenever you have to recover the primary disk (for example, *root*), do so in single-user mode from the system console, not multi-user from the X Window system. Before starting this procedure, make sure all filesystems are mounted.

There are two ways to use the save set recover feature:

- Run the **nwadmin** program and choose the **Recover** command from the **Save Set** menu. This opens the **Save Set Recover** window.
- Run **recover** with the **-S ssid** option from the system prompt. This is the way to recover if you are not using the X window system. See the **mminfo(1M)** reference page for instructions on how to find the save set(s) you want to recover. See the **recover(1M)** reference page for instructions on using the **recover -S** command.

To use the Save Set Recover window for recovering an entire filesystem, you need to pay attention to the backup levels of the save sets you are marking for a recovery. For each save set you wish to recover, you need the last full backup followed by the most recent level backups of the save set to bring the system back to the state it was in before the system crash.

A level 0 backup is a full backup. Other levels after a full are represented by ascending numbers, or by the letter *i* for incremental saves. In other words, a level 3 backs up more data than a level 5. The following illustration shows you several backups over time. The bars represent the level backups since the last full backup.

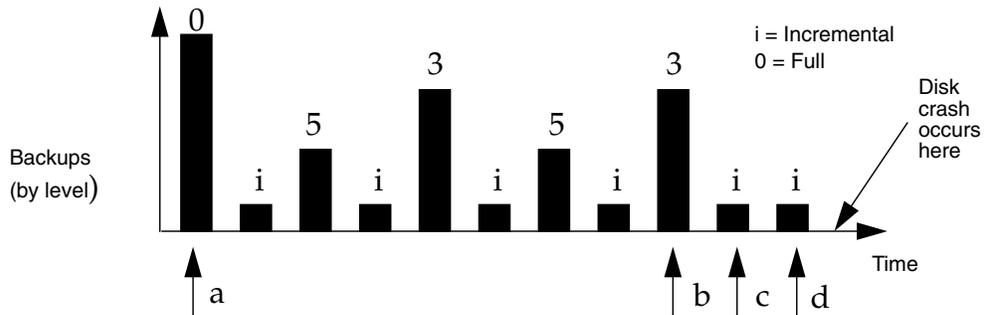


Figure A-1 Backups Needed for Disaster Recovery

The arrows in the illustration point to the save sets you would need in order to recover your filesystems from the disk crash.

By using save set recover, it is possible to recover files that were deleted between backups. For example, if file F existed at time a, but was deleted prior to time b, file F will be recovered. This may require more disk space than is available if large files or a significant number of deleted files are recovered. Normal recover does not recover the deleted files, and the filesystem will be restored exactly as it was at time d with no superfluous files recovered.

For example, suppose you wish to recover the save set `/xlv0` in the Save Set Recover window, as shown in Figure A-2.

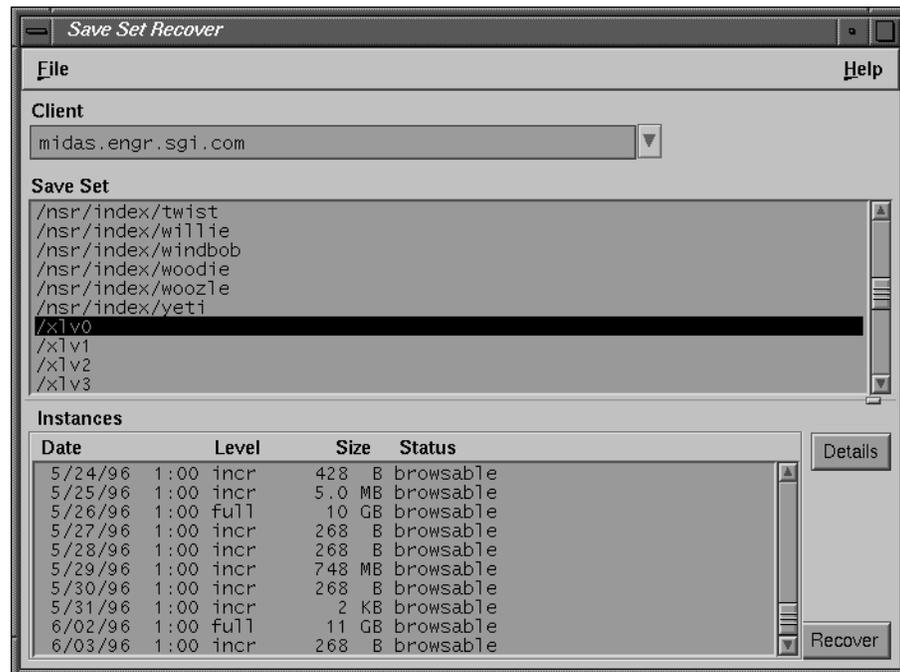


Figure A-2 Save Set Recover Window

The Instances display shows the backup history of */xlv0*. To recover */xlv0*, you need to mark the most recent full backup of */xlv0* and the most recent level 9 to restore */xlv0* to the state it was in before the system crash. In other words, you need to mark the appropriate save set levels for each filesystem you are recovering. As indicated above, you may recover files you were not expecting and do not need. Just delete these files.

Normally during disaster recovery, you want to force the recover program to overwrite existing files. Using save set recover, this is even more important, since the same file may be recovered multiple times with each successive version coming from a later save set. Since each file in each save set is recovered with save set recover, files or directories previously deleted or renamed (with the **mv** command) between backups are still recovered. These files or directories need to be manually deleted. You may even run out of space during the recovery if there are too many instances of previously deleted files or directories recovered by save set recover.

The save set recover feature reads each save set in its entirety during recovery. If you have to recover many save sets, the “normal” recover may work better than save set recover. However, if you need only the last full backup of a save set, save set recover is a better method. By contrast, if you use the Recover command from **nwrecover**, you will recover filesystems from the last backup before the crash, and NetWorker will read the minimum amount of tape to recover the files. This method has these disadvantages:

- You spend more time browsing and marking files for recovery, since NetWorker needs to find an entry for each file in the index to add all the files in one filesystem.
- You could run out of swap space if the list of filesystems to recover grows too large.

In summary, save set recover allows faster disaster recovery in these cases:

- If you can determine the correct save sets to recover.
- If there are only a few save sets to recover for each filesystem.
- If you are not recovering the NetWorker indexes.
- If recovering extra files is acceptable (not an issue with full backups).

Use the normal recovery procedure in these cases:

- If you cannot determine which save sets to recover.
- If you need to recover files from many backups to restore the filesystem acceptably.
- If you are recovering the NetWorker indexes.
- If recovering extra files is unacceptable, or you are recovering incremental backups.

Recovering a Secondary Disk

This section provides an example of how to recover a secondary disk using NetWorker. The example may apply to either a NetWorker server or a client.

It is impossible to provide step-by-step instructions on how to recover your system from a disaster, since every site is unique. The examples in this chapter are designed to give you general principles of recovering a primary or secondary disk, and to demonstrate the procedure. They are meant to be examples only, not instructions.

The example in Figure A-3 assumes the primary disk is still operational, so the system is up and can run NetWorker. However, a secondary disk is lost due to a head crash.

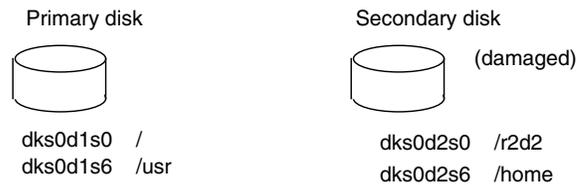


Figure A-3 Secondary Disk Damaged

If the disk is damaged, replace it with a new disk of the same type and size as the old one. You will need a disk large enough to hold all the filesystems to be recovered.

1. Install the replacement disk. Make sure the operating system and kernel recognize the new disk.
2. Label and partition the new disk so that you can recover the filesystems. Use the hard copy of the disk information to remember how large each partition was. (See “File the Disk Information” on page 207.)

Note: If you do not have this information, look at the */etc/fstab* file to find out how the disk was partitioned into filesystems. You will have to guess how much space to give each partition. Since you still have the primary disk, the partition information is available.

3. Make filesystems for each raw partition that you are going to recover and mount the block partition, consulting the hard copy of the **df** output. (NetWorker does not initialize filesystems; it recovers data into existing ones.)

Caution: Using the **mkfs** command destroys the disk contents. Be sure this disk is really destroyed before using **mkfs**.

For example, for an IRIX system:

```
# mkfs /dev/rdisk/dks0d1s0
...
# mount /dev/rdisk/dks0d1s0 /
```

After creating and mounting all the filesystems on the replacement disk, use the NetWorker save set recover feature or the NetWorker Recover window to recover the files. For more detailed information about save set recover, see Chapter 5, “Recovering and Cloning Save Sets.”

Recovering a Primary Disk

In the example shown in Figure A-4, a disk with the operating system or the NetWorker binaries is damaged. You have to first reinstall and reboot the operating system. If necessary, remake and mount all filesystems. Next, reinstall NetWorker from the original software distribution, so you can recover the data.

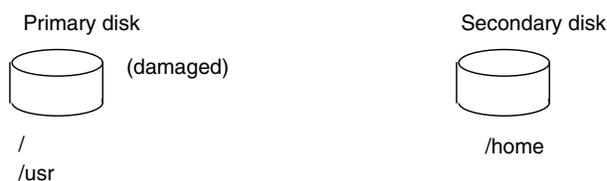


Figure A-4 Primary Disk Damaged

Caution: Anytime you have to recover the primary disk (for example, *root*), do so in single-user mode from the system console, not multiuser from the X Window system. Before starting this procedure, make sure all filesystems are mounted.

After replacing the damaged disk, format it and reinstall the operating system, using the original software distribution. Consult the documentation included with the operating system for instructions on how to reinstall the system. On a NetWorker server, reinstall the NetWorker software using the same paths and directory locations. On NetWorker clients you need access only to the NetWorker binaries. You may run NetWorker from the *nsr_extract* directory or NFS mount the binaries from another system running NetWorker.

Using the original partition information, make filesystems for each partition you are going to recover and mount them. If a filesystem is already created and mounted, you do not need to do this. For example, if you reinstalled */* and */usr*, you do not need to re-create them.

Next, reinstall NetWorker from the distribution media. Since you may have several different versions of NetWorker on the media, use the version number that matches the version of NetWorker you were running before you lost your disk. For example, if you were running NetWorker version 4.2, install version 4.2. The version must be equal to or later than the version used for the backups. Refer to the installation chapters in this guide for detailed instructions.

Tip: You do not need to reload the license enablers. If the */nsr* partition is present, the license enabler is still loaded. If the */nsr* partition was also lost, recovering */nsr* (described in the next section) recovers your license enablers.

There is one other way to access the NetWorker binaries if these were located within one of the damaged filesystems. If there is another system of the same type as the system being recovered on the network which has NetWorker running, you may NFS mount the NetWorker binaries on the damaged system.

For example:

```
# mount venus:/usr/etc /mnt
# /mnt/recover -s server -q
recover> add /
recover> force
recover> recover
```

With the operating system and NetWorker back in place, you are ready to start recovering the remainder of the data lost from the disk.

Use `mmrecov -s server_name` or the `recover -s server_name` command to recover each filesystem on the disk being recovered.

Caution: Always reboot a system after recovering its primary disk.

Recovering /nsr on a NetWorker Server

This section addresses the case where the */nsr* filesystem on a NetWorker server is lost due to a disk crash. The */nsr* filesystem contains the indexes that hold the necessary information to recover the NetWorker clients.

If the server loses its operating system and NetWorker programs, they must be reinstalled first. (See “Recovering a Primary Disk” on page 214.)

The next important step is to recover the server indexes from the backup media, using the `mmrecov` command. The `mmrecov` command asks you for the *bootstrap* save set identification number (*ssid*). If you followed the procedure recommended to prepare for a disk crash, you have a piece of paper with the name of the backup media you need and the bootstrap *ssid*.

In the following example, *ssid* "12102" is the most recent bootstrap backup:

```
December 13 03:35 1996 server's bootstrap information Page 1
  date      time      level      ssid  file  record  volume
12/12/96 18:23:23  full      12096  10    0      vol.001
12/13/96  3:34:46    9         12099  13    0      vol.001
12/14/96  3:35:18    9         12102  16    0      vol.001
```

If you do not have this piece of paper, you can still recover the indexes by finding the *ssid* using the **scanner -B** command. (See "Finding the Bootstrap Save Set ID" on page 219.)

You may need more than one backup media to recover the server indexes. During the recovery, you can use the **nsrwatch** command or open the NetWorker Administrator window to watch for pending messages requesting backup media.

With the OS and NetWorker in place, recover the indexes from backup media:

1. Find the printout of the *bootstrap* save set ID information. You need it for the next two steps.
2. Retrieve the backup media that contains the most recent backup named *bootstrap* and load it into the server's device.
3. Use the **mmrecov** command to extract the contents of the *bootstrap* backup.

For example:

```
# mmrecov
$pathname/mmrecov: Using mars as server
NOTICE: mmrecov is used to recover the NetWorker server's on-line file
and media indexes from media (backup tapes or disks) when either of
the server's on-line file or media index has been lost or damaged.
Note that this command will OVERWRITE the server's
existing on-line file and media indexes. mmrecov is not used to
recover NetWorker clients' on-line indexes; normal recover procedures
may be used for this purpose. See the mmrecov(8) and nsr_crash(8)man
pages for more details.
```

```
What is the name of the tape drive to use [/dev/rmt/tps0d3]? [Enter]
Enter the latest bootstrap save set id []: 12102
Enter starting file number (if known) [0]: 16
Enter starting record number (if known) [0]: 0
```

```
Please insert the volume on which save set id 1148869870 started into
/dev/rmt/tps0d3. When you have done this, press <RETURN>: [Enter]
```

```
Scanning /dev/nrst8 for save set 1148869870; this may take a while...
scanner: scanning 8mm 5GB tape mars.006 on /dev/nrst8
uasm -r /nsr/res/nsr.res
uasm -r /nsr/res/nsrjb.res
uasm -r /nsr/res/
nsrmmdbasm -r /nsr/mm/mmvolume
/nsr/mm/mmvolume: file exists, overwriting
uasm -r /nsr/index/mars/
nsrindexasm -r /nsr/index/mars/db
scanner: ssid 1148869870: scan complete
scanner: ssid 1148869870: 31 KB, 10 files
nsr/index/mars/db: file exists, overwriting
uasm -r /nsr/index/
uasm -r /nsr/mm/
uasm -r /nsr/
uasm -r /
mars: 31 records recovered, 0 discarded.
nsrindexasm: Building indexes for mars...
8mm 5GB tape mars.006 mounted on /dev/nrst8, write protected
```

The bootstrap entry in the on-line index for mars has been recovered. The complete index is now being reconstructed from the various partial indexes which were saved during the normal saves for this server.

```
mars# nsrwatch
```

```
nsrindexasm: Pursuing index pieces of nsr/index/mars/db from mars.
Recovering 2 files into their original locations
Total estimated disk space needed for recover is 11 MB
Requesting 2 files, this may take a while...
```

```
nsrindexasm -r ./db
merging with existing mars index
mars: 25711 records recovered, 0 discarded.
nsrindexasm -r ./db
merging with existing mars index
Received 2 files from NSR server 'mars'
mars: 733 records recovered, 0 discarded.
nsrindexasm: Building indexes for mars...
nsrindexasm: Suppressing duplicate entries in mars - 50 duplicates
discarded.
```

The on-line index for 'mars' is now fully recovered.

Notice how in the example above, the shell prompt appears during the *bootstrap* recovery. You can use the NetWorker Administrator window, or commands such as **nsrwatch** to gauge the progress of the server. Open a new terminal window to monitor the recovery so that the **mmrecov** output does not display on top of the **nsrwatch** output.

Replacing the */nsr/res* Directory

The **mmrecov** command also recovers the */nsr/res* directory, which is used by NetWorker to store configuration information such as the list of NetWorker clients and registration information. Unlike the indexes, the contents of this directory can not be reliably overwritten while NetWorker is running. Therefore, the **mmrecov** program recovers the */nsr/res* directory as */nsr/res.R*.

To complete the recovery of the */nsr/res* directory, shut down NetWorker, move the recovered */nsr/res* directory into its original location, and then restart NetWorker.

Complete these steps after **mmrecov** has finished and this final message appears:

The on-line index for 'server' is now fully recovered.

1. Shut down the NetWorker server using the **networker stop** command:

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

2. Save the original */nsr/res* directory and move the recovered version into the correct location:

```
# cd /nsr
# mv res res.orig
# mv res.R res
```

3. Restart the NetWorker server. When it restarts, the server uses the recovered configuration data:

```
# cd /
# nsrd
```

4. Once you have verified that the NetWorker configuration is correct, you can remove the */nsr/res.orig* directory:

```
# rm -r /nsr/res.orig
```

To recover other filesystems, see “Recovering a Secondary Disk” on page 212.

Finding the Bootstrap Save Set ID

If you did not file a hard copy of the *bootstrap* information, you can still find the save set ID of the most recent *bootstrap* by using the **scanner -B** command. For example:

1. Place the most recent media used for scheduled backups in the server device.
2. Use the **scanner -B** command as *root* to locate the most recent *bootstrap* on the backup media. Replace the device name in the appropriate example with the name of your device:

```
# scanner -B /dev/rmt/tps0d3nrnsv
```

The **scanner -B** command displays information on the latest *bootstrap* save set found on the backup volume, as illustrated below:

```
scanner: scanning 4mm tape ishi001 on /dev/rmt/tps0d4nrnsv
scanner: done with 4mm tape ishi001...
Bootstrap 12099 of 12/13/96 3:34:46 located on volume v.001, file 13.
```

If the *bootstrap* date looks reasonable, run the **mmrecov** command and supply the save set ID and file number displayed by the **scanner** command. Otherwise, use another backup volume to try to find a more recent *bootstrap*.

Recovering to a New Server

This section describes the case where your old NetWorker server is beyond repair, and you wish to recover NetWorker to a new server.

Tip: Things may go more smoothly if the new NetWorker server has the same *hostname* as the old NetWorker server. However, since licensing is based on host ID, you need to relicense the NetWorker server anyhow.

Follow these steps to transfer from one NetWorker server to another:

1. Install the NetWorker software from the original distribution media on the new server.

Note: If you have a jukebox, do not start the NetWorker daemons. Refer to the instructions in this guide for jukebox device driver installation and testing.

2. Find the printout of the *bootstrap* save set ID information from the old server. You need it for the next two steps.

3. Retrieve the backup media that contains the most recent backup named *bootstrap*, and load it into the new server device.
4. Add the name of the old server as an alias for the new server.
5. Use the **mmrecov** command to extract the contents of the *bootstrap* backup.

The NetWorker daemons should start up on the new server and display the following messages:

```
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 may be that you moved the NetWorker
server to a new computer.
new_server syslog: You must reregister these enablers within 15 days
to obtain new codes.
new_server syslog:
new_server syslog: License enabler #xxxxxxx-xxxxxxx-xxxxxxx (NetWorker)
```

Relicense your new NetWorker server. After moving NetWorker from one system to another, you have 15 days to relicense the new server with Silicon Graphics. Follow the instructions in the *IRIX NetWorker Installation Guide*.

Follow these steps after successfully moving your server:

1. Verify that all the clients are included in the scheduled backups.
2. Recover the client filesystems and indexes.
3. Use the Recover window to make sure all the client indexes are visible and therefore “recoverable.”
4. Back up the indexes on the new server by doing a full backup of the new server as soon as possible.

Disaster Recovery With Jukeboxes

This section provides a description of how to use autochangers during disaster recovery. The procedures are similar, but employ jukebox-specific commands.

Follow these steps to perform disaster recovery with a jukebox:

1. Read the disaster recovery procedures listed in the `nsr_crash(1M)` reference page. Perform all steps up to the point where you issue the `mmrecov` command. If only one volume is needed to recover the NetWorker file indexes, follow the instructions for `nsr_crash`.
2. Run the `jb_config` command to add the jukebox.
3. Issue the command `nsrjb -H`. This resets the jukebox for operation. If any volumes are loaded in the media drives, they are moved back to a slot. This operation may take a few minutes to finish.
4. Using the instructions in `nsr_crash`, determine which volume(s) are needed to retrieve the NetWorker file indexes. Load these volume(s) into the jukebox. See the `nsr_crash(1M)` reference page for additional information.
5. Issue the command `nsrjb -I`. This reinventories the jukebox. If you want to speed up this process, issue the command with the `-S` flag and list the slots where you placed the required backup volumes. You must list the slots in order (for example, `nsrjb -I -S 1-3`). If you want to inventory slots out of order (for example, 2, 1, and 4), you must issue the `nsrjb -I -S` command separately for each slot. All the volumes currently loaded in the jukebox will be marked with an asterisk since there is no media database.
6. Run `mmrecov`. When prompted, provide the device name you want to use, the last save set ID, and (if known) the file number and record number.
7. Load the first volume that `mmrecov` requests into the first drive in the autochanger. From a new window, issue the command:

```
nsrjb -l -n -S slot -f device_name
```

where `slot` is the slot in which the first volume is located and `device_name` is the pathname of the first drive.
8. Issue the `nsrjb -u` command after the indexes have been recovered.
9. Shut down NetWorker using the `/etc/init.d/networker stop` command.
10. The `mmrecov` command creates a directory named `/nsr/res.R` and recovers three files into it: `nsr.res`, `nsrjb.res`, and `nsrta.res`. When `mmrecov` completes the recovery, copy the three files contained in `/nsr/res.R` into the `/nsr/res` directory.
11. Restart NetWorker daemons using the `/etc/init.d/networker start` command.

Summary

The following steps summarize what you need to do if a primary or secondary disk is damaged, destroying the filesystems of a NetWorker server or client:

1. Reload and boot the system, if the operating system is lost, using the same *hostname* and disk partitioning, if possible.
2. Replace the damaged disk(s), if necessary; format, partition, make new filesystems, and mount filesystems with the same names as those that were damaged.
3. Reinstall or access the NetWorker binaries if they were lost. On a NetWorker server, reinstall them from the distribution media. On a NetWorker client, extract them or temporarily NFS mount them over the network.
4. Use **mmrecov** to recover the indexes for the NetWorker server if the */nsr* directory is destroyed on the server.
5. Recover the lost filesystems by using the normal recover process or the save set recover feature. Recover the client indexes, using the normal recover process.

Managing the NetWorker Environment

This appendix provides examples and suggestions for you to consider while you are thinking about setting up your NetWorker environment. It also offers background information to help you understand the logic behind the NetWorker backup schedule and index policy features.

For your convenience, NetWorker is shipped with preconfigured backup schedules, index policies, pools, label templates, directives, and notifications. A brief description of each preconfigured setting is presented in this appendix.

Guidelines for Choosing a Configuration

There are several factors that can help determine the NetWorker server configuration that best suits your backup and recover needs. The configuration consists of the hardware and software, including tape drives, jukeboxes, client systems, and network connections.

This section provides a few simple rules that you can use to guide your choices. Information focuses on backup, since backup requires far more server capacity than recovery. Keep in mind that these are *guidelines*; actual performances may vary.

The goal in selecting a configuration is to balance the different hardware and software limitations to achieve the overall data handling capabilities you require. Start by looking at the limits of the major NetWorker configuration components: tape drives, clients, network connection, jukeboxes, and the NetWorker server itself.

Tape Drives

Tape drives have a fixed maximum data transfer rate that they can handle. Since NetWorker automatically spans multiple tapes, the total tape capacity is not as important as the data rate. Refer to your hardware documentation to find out what the data transfer rate is for your drive.

NetWorker cannot back up faster than the data rate of your tape drive, but multiple tape drives can decrease backup time.

Clients

Different clients can generate data at different rates and, even within a single client, different types of files can generate different data rates. For example, symbolic links require as much processing as large data files, but produce no data. Consequently, the data rate produced by a backup of a single client can vary quite a bit. It is a good idea to run several clients simultaneously to help smooth out fluctuations in the data transfer rate for each client.

Network

Ethernet has an upper limit on bandwidth of about 1 megabyte (MB) per second, but in practice, most networks can handle only about 500 KB between a set of clients and a single server. Token ring has a lower maximum bandwidth (8 Mb per second) but a higher utilization ratio, so data transfer rates are approximately the same. FDDI is faster; see Table B-1.

Table B-1 Network Speed Comparison

Network	Rate
Ethernet	500 KB/s
Token Ring	500 KB/s
FDDI	5 MB/s

Server

The server must be able to handle the load of network packets, data movement, and tape drives in order to achieve the rates listed above. Most of the work on the server side is in data movement, context switching, and interrupt handling. The performance of all of these functions improves as the CPU speed increases. It takes approximately 20 MIPS to handle 500 KB/s of data, although this tapers off at high CPU speeds because bus bandwidth and other bottlenecks begin to affect the data movement. Approximately 16 MB of memory is required per 500 KB/s of data rate handled by the server.

Table B-2 shows the relationship of server power to backup throughput.

Table B-2 Server Throughput Comparison

Server CPU/Memory	Rate
20 MIPS/16 MB	500 KB/s
50 MIPS/32 MB	1000 KB/s
100 MIPS/64 MB	2000 KB/s

Jukeboxes

Jukeboxes (autochangers) provide automatic loading and unloading of tapes or optical disks. This assists the administrator in two different ways. During nightly backups, NetWorker uses the jukebox to automatically switch to new media as the data is backed up. During recovers, NetWorker uses the jukebox to load all of the media needed for the recovery without operator intervention. Refer to the jukebox documentation for data transfer rates and maximum capacity.

To determine the capacity requirements of a jukebox for a scheduled unattended backup, simply pick the jukebox with a capacity large enough to handle the largest possible amount of backup data. For example, a full backup of 60 GB requires an automatic tape loader(ATL) or Lago[®] jukebox.

To determine how much disk space the online indexes will require for quick recovers, do a rough calculation of the amount of data backed up in a single schedule period (for example, week, month, or quarter). Employ the guidelines in Table B-3 to determine how much data is backed up with different levels of backup.

Table B-3 Amount of Backup Data at Different Levels

Level	% of Data Backed Up
Full	100%
Level 1-9	25%
Incremental	10%

For example, a monthly schedule that has 1 full on the first Sunday, a level 5 on other Sundays, and incrementals every other day will look like Table B-4.

Table B-4 Sample Monthly Data Backup

Level	% of Data Backed Up
1 Full	100%
26 Incremental	260%
4 Level 5	100%
Total	460%

This illustrates that over the course of a month, 460% of the total amount of data will be backed up. For example, a total of 10 GB of client data backed up using this schedule would result in about 46 GB of data on tape per month.

In this example, different jukeboxes would provide the amount of online data indicated in Table B-5.

Table B-5 Capacity of Different Autochangers

Jukebox	Capacity	Months of Data
EXB-10e	50 GB	1
ATL/Lago	270 GB	5
EXB-120CHS	580 GB	12

Guidelines

Using the previous data transfer rate and capacity calculations, you may consider the following guidelines for configuring your network of servers and clients:

1. Assign 4 simultaneous clients per network, by setting the parallelism to 4 (if the clients are PCs, you may assign more).
2. Use one Exabyte 8500 or Sony 5200 per network, or two Exabyte 8200s.
3. Assign a NetWorker server with approximately 20-30 MIPS and 16 MB of memory per network.

Configuration Examples

This section includes two examples of NetWorker configuration to illustrate the reasoning behind selecting the components.

Example 1

Site A has approximately 30 GB of data on two networks of 50 clients and wants to schedule full backups for all of their data in one night (12 hours). The following equation calculates the required data transfer rate to achieve this goal:

$$30000 \text{ MB} \div 12 \text{ hours} = 2500 \text{ MB/hour} = 694 \text{ KB/second}$$

To back up the data with a single NetWorker server, the following configuration is suggested:

- NetWorker on a 30-40 MIPS CPU system with 24-32 MB of memory
- two Exabyte 8500 tape drives connected to the NetWorker server
- an ATL (or Lago) jukebox with two Exabyte 8500 tape drives connected to the NetWorker server
- two network interfaces
- one NetWorker Autochanger Software Module for the ATL (or Lago)

To back up the data with two NetWorker servers, the following configuration is suggested:

- NetWorker on two 20 MIPS CPU systems with 16 MB of memory
- an EXB-10e jukebox with an Exabyte 8500 tape drive connected to each server
- one network interface for each server
- two NetWorker Autochanger Software Modules for the EXB-10e

Example 2

Site B has 50 GB of data on a single network with 80 clients and they want to be able to schedule backups in a single night (12 hours). The full backups for the clients must be staggered due to the limit of 500 KB/sec data transfer rate per network. Calculate the backup capacity required to complete the backups in one night:

$$500 \text{ KB/sec} \times 12 \text{ hours} = 21.6 \text{ GB/night}$$

By staggering their full backups into three nights instead of one, using three different backup schedules, they reduce the load of the nightly backup data to about 20 GB each night, as shown in Table B-6.

Table B-6 Staggered Backup Schedules Reduce Load

Full:	$50\text{GB} \div 3 = 16.7\text{ GB}$
Incr:	$(50\text{ GB} - 16.7\text{ GB}) \times .1 = 3.3\text{ GB}$
Total:	$20\text{ GB} \div 12\text{ hrs} = 462\text{ KB/s}$

To back up the data with a single NetWorker server, the following configuration is suggested:

- NetWorker on one 20 MIPS CPU system with 16 MB of memory
- an EXB-10e jukebox with an Exabyte 8500 tape drive connected to the server
- one network interface
- one NetWorker Autochanger Software module for the EXB-10e

Measuring Performance

If you are interested in measuring the performance of your NetWorker environment, you must take into consideration the server system and the client system.

The factors to consider for the server system are the speed of the tape drive, the network speed, and the CPU speed. Factors to consider for the client system are filesystem traversing, generation of data, data on multiple disks, and CPU speed.

Server Performance

This section provides examples on how to measure the performance of the server.

Tape Drive Speed

Most tapes have step function in data rate. With most devices, NetWorker uses 32 KB per record.

To measure tape speed, follow these steps:

1. Create a large file (at least 20 MB) with non-zero data. For example:

```
# cat /vmunix /vmunix /vmunix /vmunix /vmunix > big
```

2. Use the `dd` command to write the large file to tape four times and time the results:

```
/bin/time dd if=big of=/dev/nrst8 bs=32k conv=sync
```

3. Divide the file's size by the average of the last three real times. The resulting number gives you a rate for the tape speed. For example:

```
20190 -rw-rw-r-- 1 root 20675420 Jan 7 11:04 big
95.2 real 13.0 user 11.9 sys
78.2 real 12.9 user 12.7 sys
78.0 real 12.8 user 12.5 sys
76.8 real 13.0 user 12.4 sys
```

Average: 77.67 (last three); Rate: 20190 KB ÷ 77.67 sec = 260 KB/sec

Network Speed

NetWorker for UNIX uses TCP and RPC/XDR as network communication protocols. To measure the network speed, follow these steps:

1. Create a large file (as in the tape speed measurement example) on a fast client.
2. Use the `rcp` command to copy the file from the client to the server and time the result:

```
# /bin/time rcp big server:/dev/null
```

3. To find the network speed, divide the number of bytes in the file by the real time. For example:

```
20190 -rw-rw-r-- 1 root 20675420 Jan 7 11:04 big
38.2 real 0.2 user 30.7 sys
```

Rate: 20190 KB ÷ 38.2 sec = 529 KB/sec

The most important factor affecting network speed is network errors. To determine the input error rate, the output error rate, and the collision rate, use the `netstat -i` command. If the input or output error rate is above 0.5%, or the collision rate is above 5%, network errors are slowing down the network speed.

CPU Speed

The CPU of a server limits the following:

- the total data throughput to tape
- the interrupts per second for network data
- context switches per second between processes

The best measure is the MIPS rating for the server (a larger MIPS or SPEC rating means a faster machine).

Memory

The memory on the server limits the amount of data buffered between the NetWorker **save** command, agent daemon, and media management daemon.

Client Performance

This section provides examples of measuring the performance of a NetWorker client.

Filesystem Traversing

To measure the filesystem traversing speed, follow these steps:

1. Time the **uasm** command with the **-bi** option. For example:

```
# /bin/time uasm -bi /usr
13931 records 2667396 header bytes 350849148 data bytes 124.9 real
10.8 user 34.0 sys
```

2. Divide the number of records by real time for rate per file. For example:

```
13931 records ÷ 124.9 sec = 111.5 files/sec
```

Data Generation Rate

To measure the rate at which a client generates data for a backup, follow these steps:

1. Time the **uasm** command with the **-si** option and redirect the output to */dev/null*. For example:

```
# /bin/time uasm -si /usr > /dev/null
```

2. Divide the number of bytes obtained (filesystem traversing) with the **uasm -bi** command by the real time generated by the **uasm -si** command. For example:
 $342626 \text{ KB} \div 1199 = 286 \text{ KB/sec}$

Data on Multiple Disks

NetWorker automatically backs up multiple disks in parallel.

To measure parallel disk speeds, follow these steps:

1. Use the **df** or **du** command to find two directories of approximately the same size located on different disks.
2. Run the same **uasm** speed tests for filesystem traversing and data generation rate as for one disk, but run the tests simultaneously on the two directories.
3. Add the data from each test (*files/sec* and *KB/sec*) to obtain a combined rate.

This rate reflects the performance of NetWorker backing up data on multiple disks.

CPU Speed

The CPU of a client limits the following:

- the total data throughput to the network interface
- interrupts per second for network data
- context switches per second between processes

The best measure is the MIPS rating for the client (a larger rating means a faster machine).

NetWorker Backup Schedules

The NetWorker server backs up each client system across your network according to a backup schedule. Create schedules in the Schedules window and assign individual clients to schedules in the Clients window. Schedules can be very simple or very sophisticated, depending on the needs of your environment. All clients can share the same schedule, or each client can have its own unique schedule.

This section discusses some of the considerations to keep in mind while determining which schedule best fits your situation.

Backup Levels

The schedule specifies the backup level at which NetWorker performs backups for a client on each day of a weekly or monthly period. NetWorker offers eleven different backup levels:

- Full—backs up all files, regardless of whether or not they have changed since the previous backup. A full backup is equal to a level 0 backup.
- Level 1 through level 9—back up files that have been modified since a previous full backup or a backup of a lower numbered level. For example, a level 3 backs up all the files that have changed since the previous level 2, level 1, or full backup.
- Incremental—backs up all files that have changed since the previous backup of any level. An incremental backup is equal to a level 10.

You can also skip a backup on a given day. You may want to schedule a “skip” backup on weekends or holidays when no one is available to load backup media. If you have a jukebox, this option becomes less important, since a jukebox automatically loads and unloads backup media.

NetWorker’s on-screen calendars present an easy method for setting up backups for each day of the month. You can designate a schedule and repeat it over a weekly or monthly period. For example, if you set up a full backup for one Friday, NetWorker automatically sets up a full backup for every Friday. Or, you can override the regularly scheduled backup level for a specific day.

There is no “correct” way to set up a backup schedule for a particular client or network of clients. The clients you need to back up probably vary considerably—some have a lot of critical data to back up, others may have a small amount of data that does not change very often. Consider the situation for each client, weigh the benefits of different backup schedules, and then select the best schedule for each client.

Full Backups Versus Incremental Backups

If your site has a small number of files, you may choose to perform a full backup every day, or perhaps once a week. This is a simple schedule to set up and execute, and it makes recovering from a disk crash easy—you simply need the last full backup volume.

The situations to consider are listed below:

- Full backups take more time to complete than incremental backups.
- If the full backup does not fit on a single piece of media, someone has to monitor the backup and change the media (unless you have a jukebox).
- Full backups cause the online indexes to grow more rapidly than incremental or level backups.

You may decide to schedule a full backup at the beginning of the period and schedule incremental backups the rest of the period. This schedule minimizes the amount of time that backups take, the size of the backups, and disk space required for the online indexes.

However, if you need to recover from a disk crash, you may need all the tapes used during the schedule, because the most current version of your files may be scattered across several different tapes. Although NetWorker asks for each tape that it needs for the recovery by name, loading and unloading them can be time-consuming (unless you have a jukebox, or all the incremental backups fit on one tape).

Using Level Backups

You can use level 1 through level 9 backups to moderate between the two extremes described above. Level 1 through level 9 backups allow you to set up a schedule for each client that balances your need for small, fast backups that do not take up too much index space and the need to recover quickly and easily from a disk crash.

A level backup serves as a checkpoint in your schedule since it collects into a single backup session all the files that have changed over many days or even weeks. Without a level backup, these files would be spread across tapes from many different backup sessions. As a result, a level backup can simplify and speed file recovery.

To illustrate the effect of level 1 to level 9 backups, consider two examples. In the first example, a full backup takes place on the first day, followed by a level 9, level 8, level 7, and so on down to a level 1 backup over time.

A full backup followed by level 9 to level 1 is illustrated in Figure B-1.

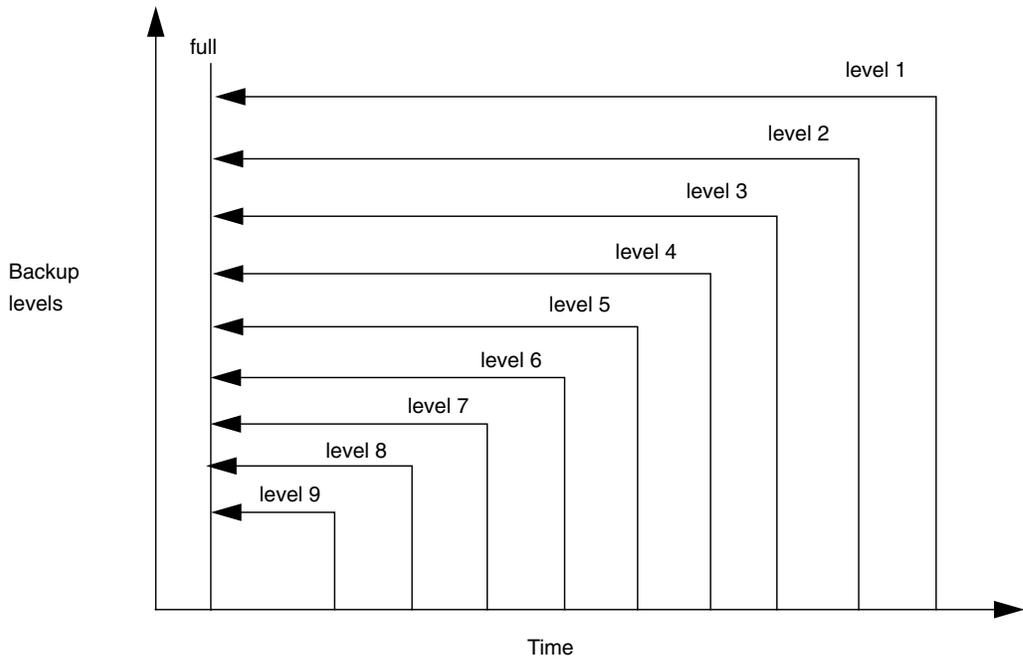


Figure B-1 Full Backup Followed by Levels 9 to 1

The advantage of this schedule is that to recover from a disk crash, you only need two tapes: the one with the full backup, and the one with the last level backup. The disadvantage is that with each day, there are more changed files to back up, so the backups take longer to complete.

In the second example, the backup schedule also starts out with a full, but the level backups that follow are in reverse order: starting with a level 1 on the first day following the full, on down to a level 9 backup. Each day, the backup backs up only the files that have changed on that day.

A full backup followed by level 1 to level 9 is illustrated in Figure B-2.

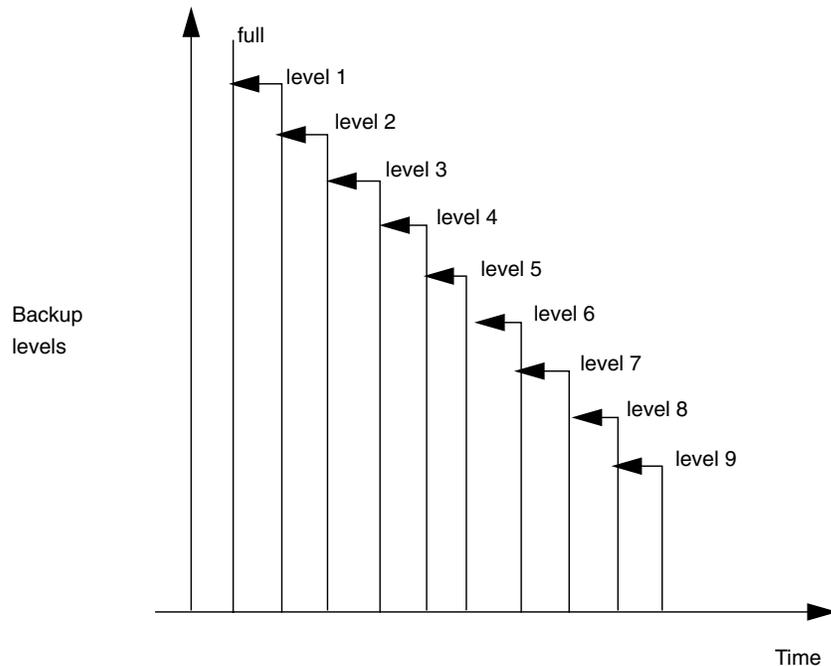


Figure B-2 Full Backup Followed by Levels 1 to 9

The advantage of this schedule is that each day's backup is small and completes quickly. The disadvantage is that recovering from a disk crash would require the full backup tape and all of the level backup tapes up until the day of the disk crash.

Neither of these backup schedules is practical. They simply illustrate how level backups work. The real power of level backups comes into play when you combine multiple levels along with fulls and incrementals.

Typical Monthly Backup Schedule

Sites with even a few gigabytes of files to back up often choose a monthly schedule based on fulls, incrementals, and level backups. The example described in this section performs a full backup on the first day of each month, a level 5 backup on the 10th and 20th of the month, and incremental backups on all other days.

A monthly backup schedule minimizes the size of daily backups while also making it relatively easy to recover after a disk crash. This schedule offers several advantages.

First, the level 5 backups simplify recovery. Assume that a disaster strikes on the 24th of the month. All the files that you need to recover an entire client system are located on tapes from just five backup sessions:

- the incrementals from the 21st, 22nd, and 23rd
- the level 5 backup from the 20th
- the full backup at the beginning of the month

Second, the incremental backups are relatively small and quick to execute, even for large network environments. Several days of incrementals fit on a single tape. This further simplifies recovery and also avoids the need to have someone change tapes each day. Illustrated in Figure B-3 is level 5 and incremental backups after a full.

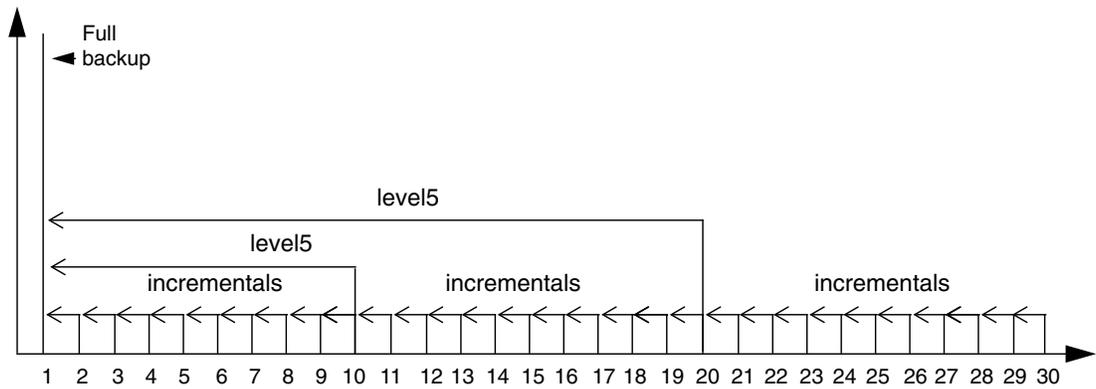


Figure B-3 Full Backup Plus Incrementals and Level 5

Backups Take Time

The amount of time you have to complete a backup on any given day also influences the schedule that you decide to use. Because of flextime and around-the-world operations, many networks must be up and running for users from early in the morning until very late in the evening. While NetWorker is able to back up live filesystems, most administrators want 100% of their network and systems capacity ready for users during work hours.

What number of files can NetWorker back up in, for example, a four-hour backup window? The answer is “it depends.”

- Select a backup server with enough CPU power, memory, and bus bandwidth so the backup server is not a bottleneck.
- Leave NetWorker’s parallelism feature turned on. This feature causes multiple client systems to send their files to the backup server in parallel. This keeps a stream of files ready for the tape drive, so that it does not start and stop. Find a parallelism that keeps the tape drive streaming without overloading the CPU.
- Experiment with compressing files on the client systems to reduce the size of the data sent across the network and written to tape. Compression may speed your backup as long as the client systems are still able to supply files to the backup server fast enough to keep the tape drive streaming.
- Take advantage of NetWorker’s ability to skip specified files during the backup. For example, you could choose to skip *core* files.
- Add a second backup device to your server. With NetWorker Server Edition and Network Edition, you can simultaneously back up to more than one device.

If your backup server can drive a single 8 mm tape drive at an average of 400 KB/second (its maximum speed is 500 KB/second and some time is invariably lost loading the tape or rewinding, for example), you will be able to back up a maximum of 5.76 GB in four hours. If you have more than this amount of data to back up, then full backups will be limited to weekends and holidays when users will not be affected.

Unless you have a jukebox and NetWorker’s optional Autochanger Software Module, you also have to schedule backups based on someone being available to load and unload media. Many administrators find that an incremental backup of their network fits onto a single 4 mm or 8 mm tape, but they must schedule multi-tape full and level backups for specific nights or weekends when an operator is on duty to load additional tapes. If an operator will not be available over a holiday weekend, you can set an override in the schedule to skip the backup on that day. You may also want to override the schedule just before a holiday with a full backup, for added peace of mind.

Using Compression During Backup

Software compression involves significant CPU usage. If you have a fast client, and performance monitors such as **gr_osview** indicate that CPU bandwidth remains, use **compressasm** to compress data before it goes over the wire to the server.

To use **compressasm**, open the Clients window and select Unix with compression directives in the Directive field. This is the default for IRIX NetWorker.

Software compression reduces network traffic, because **compressasm** typically achieves 2:1 compression (your mileage may vary). There is no harm in using **compressasm** in conjunction with a compressing tape drive, but the tape drive will probably not achieve further compression on the data.

If you are deciding between **compressasm** and a compressing drive solely to increase the amount of data on a tape, obtain the compressing drive—the hardware on the drive generally compresses faster than NetWorker and places *no* load on the client CPU for data compression.

Follow these guidelines when compressing data during backup:

- Use **compressasm** to minimize network bandwidth if you have the available CPU power.
- Use compressing drives to get more data on a tape.
- Any generic compressing algorithm typically achieves 2:1 compression, and tape drives are no exception. Sometimes you get more, sometimes less.
- Compressing already compressed data, such as GIF or JPEG images, has no effect (and might even expand the data).
- Do not use **compressasm** with a compressing drive and no networked clients.

Staggering the Backup Schedules

Networks with a large number of files can take a very long time, and require a lot of loading and unloading of tapes to complete a full backup. There may not even be time in a night or an entire weekend to complete a full backup of all the systems across a very large network. An easy way to handle this problem is to stagger the clients' backup schedules. Rather than have every client system perform a full backup on Monday and incrementals the rest of the week, for example, you can schedule some clients to perform a full backup on Tuesday and others on Wednesday.

NetWorker goes one step further to smooth the backup load for very large client systems. With NetWorker you can assign a separate backup schedule to each filesystem. Each filesystem, in essence, is treated as if it were a completely separate client.

Convenience Versus Security

You may leave the same backup volume mounted in the server's backup device throughout a week or month, and when it becomes full, replace it with a new labeled backup volume. NetWorker tracks all the backups, no matter what day of the week or month, or what part of the backup schedule cycle is in effect. The same backup volume may contain full, level [1-9], or incremental backups, and it makes no difference to NetWorker. For you, the benefits are fewer backup volumes to manage and the ability to recover from a disk crash with a minimum number of backup volumes.

Some sites prefer to segregate the full backups from the level [1-9] and incremental ones. The full backups protect the network from a catastrophic disk loss, and you want to guarantee their integrity. There is always a very small risk that if you leave the backup volume with the full backup sitting in the backup device, something could happen to it.

If a backup volume with incremental backups is ruined, users may lose one day of work. If the backup volume with the full backup is destroyed, users may lose all the work done since the last full backup. Therefore, some administrators prefer to remove the backup volume used for a full backup, put it in a safe place, and mount another backup volume for the following level [1-9] and incremental backups. The trade-off is that you may need a few more backup volumes to recover from a disk crash—the one with the last full, and the other volumes that contain the most recent level [1-9] and incremental backups.

NetWorker Browse and Retention Policies

NetWorker maintains online indexes of all the files backed up for each client and an index of the files stored on each piece of media. NetWorker lets you set policies that automatically control how long the information is retained in these online indexes. This section explains NetWorker's browse and retention policies and the trade-off between providing faster, easier recovery for your users or conserving disk space.

Browse Policy

One of NetWorker's popular features is the ability to browse many versions of a file that have been backed up over time and to choose which one to recover. However, each version of a file that NetWorker tracks takes up space in the client online index (about 220 bytes each). Since disk space is limited, you need to establish a policy of how far back in time you will keep information about backed-up files in the indexes.

The browse policy that you select specifies how long the entries for your files remain in the file indexes. A browse policy can be any number of days, weeks, months, or years. NetWorker automatically deletes entries older than the browse policy time and frees up disk space. The browse policy, like the backup schedule, can be different for each client.

How Browse Policies Work

To recover a complete directory or filesystem, you often need to recover some files from incremental and level backups as well as from a full. The incremental backup is dependent on the level backups and, in turn, on the full. NetWorker will not delete the entries from any backups on which other backups depend. As a result, you may find that entries are deleted later than you expect.

In Figure B-4 the browse policy is set to one week, which happens to equal one complete backup cycle.

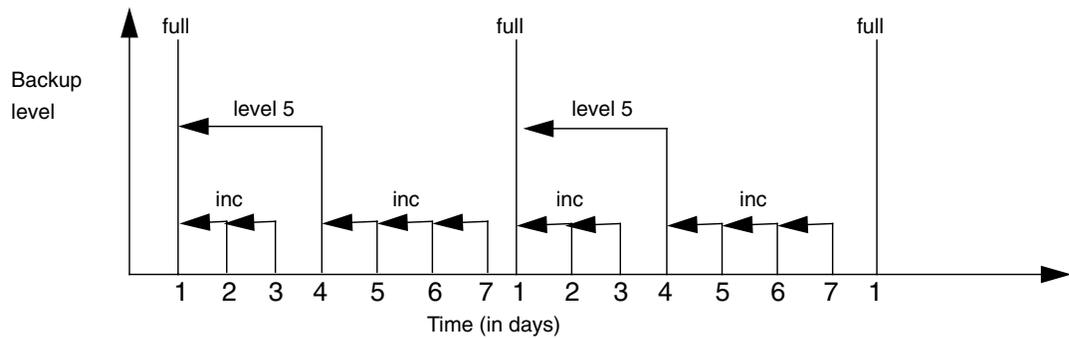


Figure B-4 Backups and the Browse Period

NetWorker will not remove the first full backup from the online file index until all the incremental and level-5 backups that depend on it have expired. As a result, the full backup actually stays in the online index for a period of time equal to the browse policy plus one full backup cycle.

The first full backup will not be removed from the online index in exactly one week, however, because there are incrementals and a level 5 backup, which have not yet expired, that depend on the full. Each incremental backup will be removed from the online index one week from the time it was completed. The level 5 backup will be

removed one week after the last incremental that depends on it is removed, and then the full backup will be removed at that same time.

The rule to remember is that a full backup actually remains in the online index for a period of time equal to the browse policy plus one complete backup cycle. A backup cycle is measured from one full backup to the next full backup. Also note that the browse policy is set for an entire client (or filesystem, if the filesystems are separately scheduled). Consequently, whatever policy you have for keeping full backups online and browsable in the file index you must also use for all incremental and level backups. With NetWorker you manage backup cycles (the period from one full backup to the next); you do not independently manage different levels of backups.

Reclaiming Disk Space

NetWorker automatically reclaims disk space that is freed up when entries are deleted from the online file indexes. However, the space is not returned immediately to your system. NetWorker takes some time, processing power, and swap space in order to reclaim this space, and having this constantly taking place on your backup server is inefficient. Instead, NetWorker first reuses this space to store information about new files that are backed up. When the file index for a client reaches a point where less than 50% of its space is being used by files that have not reached the end of their browse period, then NetWorker automatically invokes a process that returns the space to your system.

You may also reclaim disk space at any time by using the Reclaim space button in the Indexes window.

Recovering Files Removed From the Index

You can recover files whose entries have been removed from the online index because they have passed the Browse policy period as long as the files are still stored on a backup volume. However, the recover process is not as convenient as when the entries are still in the online index.

If you do not want to rebuild the index, the save sets you need are still in the media index. Since you know which save set contains the file you want, you can use the save set recover feature to recover the entire save set or selected directories and files. The save set recover feature is most useful for recovering from full backups and is limited to *root* and users belonging to the group *operator*.

If you want to rebuild the file index so that you can browse for the file you lost, here is the basic procedure to follow:

1. Use the Volume Management window to find the name of the backup volume that contains the save set.
2. Use the **mminfo** command to determine the save set ID. Use this syntax:

```
mminfo -v -s server -c client -N saveset volume_name
```
3. Rebuild the file index entries for the save set using the **scanner -i -s save_set_id#** command at the system prompt. Enter the save set ID number determined above for *save_set_id#*. Rebuilding the file index using the **scanner** command may take some time.
4. Use the NetWorker Recover window to identify the needed file(s) and initiate the recovery.

Recovery is considerably easier if the file information is still in the NetWorker online index. That is why you want to set a browse policy long enough to cover most recover requests.

Media Retention Policy

Your need to conserve disk space may lead you to establish a short browse period. NetWorker's media retention policies complement the browse policy by letting you specify a longer period of time during which files can still be recovered, although with more difficulty. NetWorker uses the retention policy to automatically recycle backup volumes.

NetWorker maintains a file index for each client system and a much smaller media index that tracks the save sets stored on each backup volume. When NetWorker removes entries that are older than the specified browse time from a file index, it leaves the corresponding save set information in the media index. The retention policy controls how long this information is kept and, as a result, how long a backup volume is kept before it can be overwritten with new backups.

As with the backup schedule and browse policy, you set the retention policy for each NetWorker client. Different clients can have different policies. The retention period can be any number of days, weeks, months, or years as long as the retention period is equal to or longer than the browse policy.

A NetWorker backup volume can contain save sets for many different clients over many days. As the retention period is reached for each save set, information about that save set is removed from the media index. When the retention period for every save set on a backup volume is reached, NetWorker marks the volume “recyclable.” This volume can then be reused for backups. At the time that the volume is actually reused, the old files are overwritten and can no longer be recovered.

The NetWorker browse and retention policies combine to give you a hierarchy of recovery capability while keeping the disk space needed for the online indexes to a minimum. Recovering a file is quick and easy using the NetWorker Recover window until the browse policy time is reached and the file information is removed from the file index. Then you can use save set recover or the more tedious process described to recover your files until the retention policy time is reached and the backup volume is recycled.

Setting Policies When Using a Jukebox

The NetWorker Autochanger Software Module automates your backup and recover activity. The capacity of the jukebox, the backup schedule you select, and the browse and retention policies you use determine whether you can walk away from backups for a week, a month, or even longer.

Jukebox Capacity

A jukebox is most useful if it has at least enough capacity to complete one entire backup cycle without intervention. This allows backups to run while you are out ill, on vacation, or busy with a user emergency, and helps minimize the time that you spend on backup (particularly if the backup server and jukebox are located some distance away). At the end of the cycle, you can move the used backup volumes offsite and load fresh tapes into the jukebox.

A jukebox with the capacity for one entire backup cycle also speeds file recovery. If a user accidentally deletes a file, there is at least one version (more if the user has recently edited the file) in the jukebox. With NetWorker, the user can quickly identify the lost file and initiate the recovery. The jukebox loads the needed tape and NetWorker completes the recovery without your help. Depending on the speed of the jukebox and the device used, the file should be recovered very quickly.

You need to design a schedule that fits the capacity of your jukebox. Start with your ideal schedule and then consider these suggestions to reduce the size of a complete backup cycle:

- Use more incremental backups and fewer level 1-9 backups.
- Back up systems with less critical files less often—perhaps only once a week.
- Use NetWorker’s directives to skip files during the backup, for example, *core* files.
- Shorten the length of the backup cycle.

Although your jukebox may have enough capacity for only one backup cycle, you can still set the browse and retention policies for a longer period. If a user tries to recover a file stored on a volume that is not in the jukebox, NetWorker will prompt you to load that volume. You can use the Location field in the Volumes window to keep track of volumes. Users can refer to this information when deciding which version of a file to recover and choose the one stored on a tape that is located in the jukebox.

Choosing the Jukebox Capacity

With just enough capacity in the jukebox for a single backup cycle, you must reload tapes at the end of each cycle. With more capacity you can set the schedule and the browse and recover policies so that the jukebox runs unattended for a long period of time. The jukebox automatically recycles tapes containing save sets that have passed their browse and retention times to continue backups virtually indefinitely.

Suppose you established a backup schedule for your network of systems that takes one week to complete (for example, you schedule a full backup once a week) and consumes a total of 12 GB of tape during the week. Assume that you are using a 50 GB Exabyte-10e jukebox. Each of the following combinations of browse and retention times will allow the jukebox to operate without intervention for an extended period of time:

- browse policy = 1 week, retention policy = 1 week
- browse policy = 1 week, retention policy = 2 weeks
- browse policy = 2 weeks, retention policy = 2 weeks

Each of these sets of policies has its advantages. With a browse and retention policy of just one week, your online indexes will be kept small. With a browse and retention policy of two weeks, your indexes will be larger but your users will have more versions to select from when they need to recover a file. A browse policy of one week and a retention policy of two weeks keeps your indexes small and allows you to recover older

files, although with a great deal more effort than if those files were still browsable in the index.

If you set the browse policy to four weeks, $4 \times 12 \text{ GB} = 48 \text{ GB}$ will fit in the jukebox. First, a full backup actually remains in the online index for a period of time equal to the browse policy plus one complete backup cycle. Thus with a browse policy of four weeks, essentially five weeks of backups would need to fit into the jukebox.

Second, since NetWorker cannot recycle a tape until all the save sets on that tape have expired, there is often some amount of “unavailable” tape in the jukebox.

Suppose that one year later the number of files that you have has grown so that the one-week backup cycle needs 18 GB of tape capacity. A browse policy of one week and a retention policy of one week still allow the jukebox to run continuously unattended.

If you want to keep files online in the jukebox longer, then you can use the methods listed earlier to reduce the size of the backup cycle. As an alternative, you can stretch out the backup cycle. For example, you can perform full backups every other week rather than every week. This should not greatly increase the size of a backup cycle and gives you more versions of files online in the jukebox.

Choosing a Jukebox

In the ideal situation, you first design the best backup schedule and set of policies for your environment and then determine the jukebox size that you need to purchase.

Assume that you have a network of systems with a total of 25 GB of files to back up and that you have selected a schedule that includes a full backup at the beginning of each month, a level 5 on the 10th and 20th of each month, and incrementals on all other days. The calculations in Table B-7 illustrate that one complete backup cycle should be about 64 GB in size.

Table B-7 Storage Requirements for One Backup Cycle

Level	Size	Frequency	Total
Full	25 GB	× 1 time /month	= 25 GB
Level 5	2.5 GB ^a	× 2 times/ month	= 5 GB

Table B-7 Storage Requirements for One Backup Cycle

Level	Size	Frequency	Total
Incremental	1.25 GB	× 27 times/month	= 34 GB
		Grand Total	= 64 GB

a. These size percentages are based on Legato experience over the past three years.

To determine the size jukebox that you need, start by estimating the size of a complete backup cycle. Now assume that you have decided on a browse policy of two months for all the client systems and a retention policy of six months. These policies let your users quickly recover any file and any version of a file that they had during the past two months, and with some effort, you can recover files that they owned at any time during the past six months. So you will need 6 months × 64 GB = 384 GB of capacity.

In practice you need a little extra jukebox capacity since there will be a small number of “unavailable” volumes; NetWorker must wait to recycle a tape until after all the save sets on that tape have expired.

Finally, remember to plan for growth in the number of your files. While sites differ in the rate at which their files are growing, a rule of thumb is to purchase a jukebox with about 50% more capacity than your current requirement.

Preconfigured Selections

NetWorker provides preconfigured settings for you to use so you can immediately start backing up your systems. This section offers an explanation of the different preconfigured settings.

Preconfigured Backup Schedules

For your convenience, NetWorker is shipped with several preconfigured backup schedules. If these schedules fit your backup requirements, you can use them “out of the box,” or you can create new ones to accommodate your site-specific needs.

This section explains the logic behind each schedule. After understanding how they work, you may want to use them as examples to set up your own schedules.

The most efficient way to protect the systems from file loss *and* maintain control over the number of backup volumes is to follow full backups with level [1-9] and incremental backups.

Each time you use the Schedules window to create a new weekly backup schedule, the preconfigured default schedule shown in Figure B-5 appears in the calendar as your starting point.

You are not allowed to change the name of an existing schedule. For example, if you want to *change* the schedule “Full Every Friday” to “Full Every Monday,” you must delete the “Full Every Friday” schedule and create a “Full Every Monday” schedule. You cannot change the existing schedule to complete full backups on Mondays instead of Fridays and then edit its name.

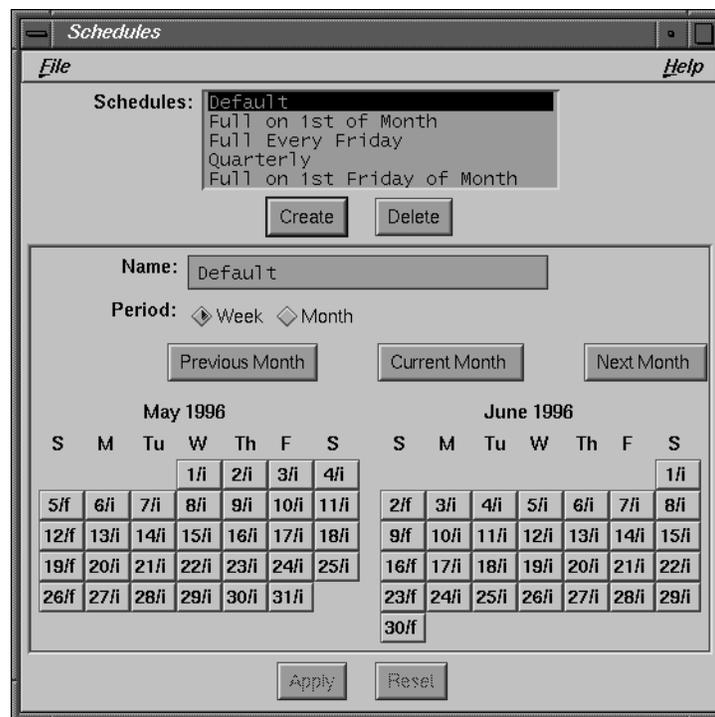


Figure B-5 Schedules Window

- **Default**—this is the only schedule you may not delete. It is a weekly schedule and completes a full backup every Sunday, followed by incremental backups all other days of the week.

This schedule is convenient if you want to premount the backup volume Friday night before you go home for the weekend. On Monday mornings, check your messages from NetWorker to make sure the backup completed. If you want to separate the full backups from the incrementals, remove the backup volume with the full backup and mount another one for the incremental backups.

- **Full Every Friday** — this weekly schedule completes a full backup every Friday, followed by incremental backups the other days of the week.

This schedule is identical to the Default schedule, except that instead of completing a full backup on Sundays, the full backup takes place on Fridays. Depending upon how much data changes on the network, the daily incremental backups might all fit onto one backup volume. In that case, if you had to recover from a disk crash, you would need only two backup volumes: the one with the last full backup, and the one with the incremental backups.

- **Full on 1st Friday of Month**—this monthly schedule completes a full backup on the first Friday of the month (not the first calendar day of the month. Incremental backups take place on all the other days.

The advantage of this schedule is that you complete a full backup only once a month. If you use this schedule, it would be a good idea to store the backup volume with the full backup in a safe place, and use other backup volumes for the incremental backups. It would also be a good idea to change backup volumes every few days for the incremental backups. If you allow all the incremental backups to be stored on one backup volume, and it is destroyed near the end of the month, you may not be able to fully recover from a disk crash.

Whenever you create a monthly schedule for a full backup on a *weekday* instead of a *calendar* day (like Friday, in this example), you must set the overrides in each month. (Notice the “f*” in the first Friday of each month.) This is because the first weekday (Monday through Friday) in a month may fall on any calendar day from 1 to 7.

Note: The Overrides you select for individual days do not carry over from one year to the next. Preconfigured schedules, however, do maintain the overrides for years into the future.

- **Full on 1st of Month** — this monthly schedule completes a full backup on the first calendar day of the month. On the other days of the month, an incremental backup takes place. This schedule has the same advantages and disadvantages as the “Full

on 1st Friday of Month” schedule. This schedule is easier to create because you do not have to set any overrides manually.

- Quarterly—the quarterly schedule completes a full backup on the first day of the quarter. A level 5 backup takes place on the first day of the other months in the quarter. Every seven days, a level 7 backup takes place. The other days of the month, an incremental backup takes place.

This schedule is convenient because a full backup takes place only once a quarter. On the first day of the month, a level 5 backs up everything that has changed since the first day of the quarter. Every seven days, the level 7 backup protects all the data that has changed since the first day of the month. The daily changes are protected by incremental backups.

If you use this schedule, it is a good idea to segregate the backup volumes and store them in a safe place. Use one volume for the full backup, one for the level 5 backups, one for the level 7 backups, and another one for the incremental backups. If you have a disk crash or a disaster, you risk losing only a few days’ work (the backup(s) on the mounted volume). If you change the backup volume every day for the incremental backups, you risk losing only one day’s work; however, you must use more tapes to recover from a disaster.

When you create a quarterly schedule like this one, use the Month period to set the level backups, then set each quarterly full backup on the calendar with an override.

To recover from a disk crash, you would need the backup volume with the full backup, the latest level 5, the latest level 7, and the incremental backups for the week.

Preconfigured Policies

NetWorker is shipped with five preconfigured policies: Decade, Month, Quarter, Week, and Year. Use these policies to choose the length of time to retain the entries in both the file index and media index. Remember, the *retention* policy you select affects the size of the *media* index and controls the length of time NetWorker tracks the backup volumes and the data on each volume.

The *browse* policy affects the size of the *file* index and the length of time NetWorker retains entries for every file backed up and visible in the Recover window. You must always choose a retention policy that is greater than or equal to the browse policy.

For example, if you choose Quarter for the retention policy for a client, and Month as the browse policy, the client will be able to browse all the file entries for backed-up files dating back a month. Each month the oldest entries for the client's files are automatically removed from the server's file index. However, the backup volumes that contain the files are still tracked by NetWorker in the media index.

The Policies window and the five preconfigured policies are shown in Figure B-6.

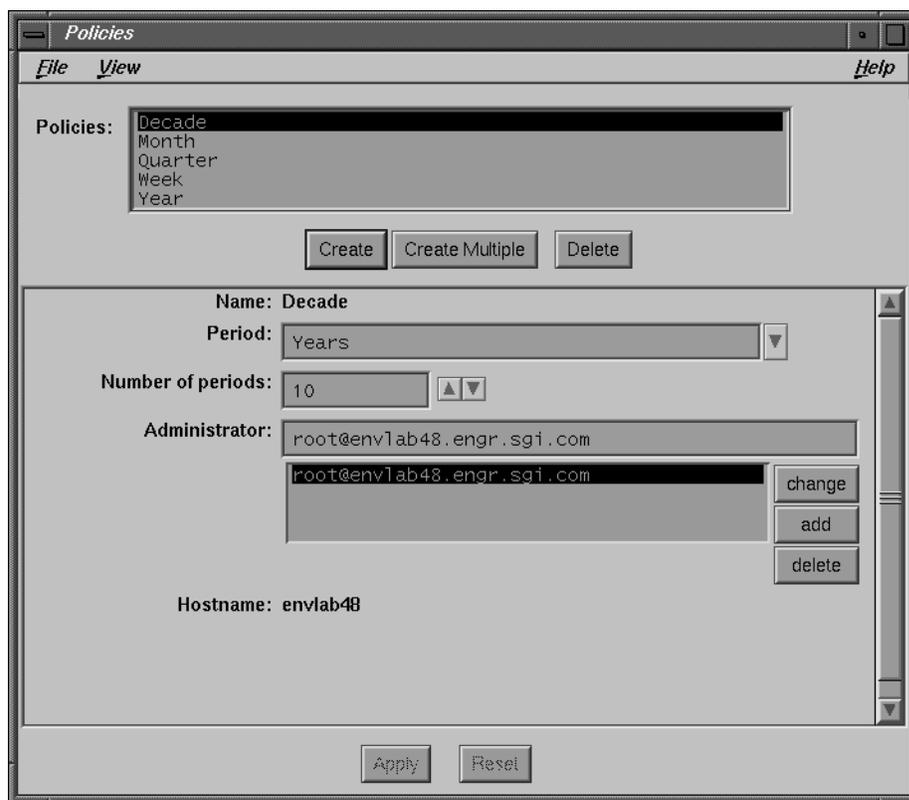


Figure B-6 Policies Window (Details)

See "Manually Managing the Online Indexes" on page 36 for an illustration of how browse and retention policies work.

- **Week**—this policy maintains the file index entries or the media index entries for one week after the last full backup. If you use this browse policy, the users will be able

to view and mark for recovery only those files that go back in time for a week. It is a useful browse policy when you have a limited amount of disk space and users do not expect to be able to recover versions of their data older than one week.

As a retention policy, *Week* means that your backup volumes will turn over quickly, and NetWorker will recycle through the tapes at a faster rate. Use this policy if you schedule weekly full backups and need to keep backup data for only one backup cycle plus a week.

- **Month**—This browse policy allows users to view and recover versions of files dating back at least a month. The *Recover* window displays versions of files backed up for one full month plus a number of weeks. As a retention policy, NetWorker maintains and tracks the backup volumes for one full backup cycle plus a month.
- **Quarter**—Use this policy if you need to keep backed-up data longer than a month. With this browse policy, the client can view and recover files for at least three months into the past. The retention policy tracks the backup volumes for at least three months plus one full backup cycle.
- **Year**—If you need to keep backed-up data online for several months, use the *Year* policy. For example, if your company requires ready access to information going back in time for at least three quarters, this is a good browse and retention policy. Realize, however, that NetWorker requires more disk space to maintain all the information online.
- **Decade** - This policy retains the entries in the server's indexes for ten years. It is useful for organizations that are required to retrieve individual files for very long periods.

Your NetWorker server will require *lots* of disk space for the online indexes if you choose *Decade* for your browse policy. Depending upon how much data you are backing up, ten years of file index entries could take up gigabytes of disk space.

It would make more sense to use *Decade* as the retention policy and use *Quarter* or *Year* as the browse policy. NetWorker can then track the backup volumes and the data on each one. You would always be able to retrieve data from an old backup volume using the *save set recover* feature if you needed to do so. NetWorker would still require disk space to maintain the media index, but it would be a much smaller amount of space using the *Quarter* or *Year* browse policies.

Preconfigured Pools

NetWorker is shipped with preconfigured pools and matching label templates. Each preconfigured volume pool has a set of unique preselected choices. If you do not choose a pool for your backups, they are automatically assigned to the preconfigured Default pool and are labeled using the Default label template.

The preconfigured pools have been included for your convenience and provide a variety of ways for organizing your data.

The preconfigured volume pools have matching label templates. The Two Sided label template is for labeling optical media and is the only template that does not have a matching volume pool.

You can use the Default, Default Clone, Archive, and Archive Clone pools without making any additional selections in the Pools window. To use the other preconfigured pools, you must first complete the selections and choose Yes from the Enabled choices. A pool must be enabled in order for NetWorker to sort data to that pool.

The preconfigured pools are described below:

- **Archive**—for archiving client data only. This pool cannot be modified or deleted. The preconfigured settings are Enabled—Yes, Label template—Archive, Pool type—Archive, Store Index entries—Yes. There are no selections for you to make for this pool.
- **Archive Clone**—for cloning archive data only. This pool cannot be modified or deleted. The preconfigured settings are Enabled—Yes, Label template—Archive Clone, Pool type—Archive Clone, Store Index entries—No. There are no selections for you to make for this pool.
- **Default**—automatically used if you do not choose a pool. If you decide not to use the pools feature, NetWorker automatically places all of your backup volumes in this pool. The Default pool cannot be deleted or modified. The preconfigured settings are Enabled—Yes, Label template—Default, Pool type—Default, Store Index entries—Yes. There are no selections for you to make for this pool.
- **Default Clone**—automatically used if you do not choose a pool for cloned data. If you decide not to use the pools feature, NetWorker automatically places all of your cloned backup volumes in this pool. The Default Clone pool cannot be deleted or modified. The preconfigured settings are Enabled—Yes, Label template—Default Clone, Pool type—Backup Clone, Store Index entries—No. There are no selections for you to make for this pool.

The Full, NonFull, and Offsite pools are intended for sorting data by levels.

- **Full**—use this pool for full backups only. This pool separates all of your full backups from the incremental and level backups. Using the Full pool allows you to easily track and separate your full backups from the incremental and level backups. Typically, you use this pool in conjunction with the NonFull pool. The preconfigured settings are Enabled—No, Label template—Full, Levels—full, Pool type—Backup, Store Index entries—Yes.
- **NonFull**—use for any backups other than full backups. This pool includes all incremental and level backups. Use the NonFull pool to easily keep your incremental and level backups separate from the fulls. Typically you use this pool in conjunction with the Full pool. The preconfigured settings are Enabled—No, Label template—NonFull, Levels—all level and incremental backups, Pool type—Backup, Store Index entries—Yes.
- **Offsite**—for volumes being stored offsite. The Offsite pool allows you to easily create a set of volumes to be stored offsite. If your onsite backup volumes are destroyed, you can still recover your valuable data with the volumes you have stored offsite. If you are also using the Full pool, you must disable it while you are sending data to the Offsite pool to ensure that all of the full backups will go only to the Offsite pool. The preconfigured settings are Enabled—No, Label template—Offsite, Levels—full, Pool type—Backup, Store Index entries—No.

Tip: Remember to enable pools you wish to have in effect during scheduled backups by selecting Yes from the Enabled choices.

Preconfigured Label Templates

The preconfigured label templates shipped with NetWorker are Archive, Archive Clone, Default, Default Clone, Full, NonFull, Offsite, and Two Sided. These are provided so that you can easily start labeling your backup volumes. There are also preconfigured volume pools with corresponding names (except Two Sided). The preconfigured volume pools automatically use the preconfigured label template with the same name.

The number range for all of the preconfigured label templates starts at 001 and ends with 999 to allow for expansion of the volume pools.

The Archive label template is used only for clients that need to archive data. It has three fields each separated with a period. The first field contains the NetWorker server name, the second field is “archive,” and the third field contains a number.

For example:

```
server.archive.number  
space.archive.001  
space.archive.099  
atlas.archive.325
```

The Archive Clone template has three fields separated by periods. NetWorker uses this template for backup volumes belonging to the Archive Clone pool. The first field contains the name of the NetWorker server and the letter “c.” The second field is “archive,” and the third field contains a number.

For example:

```
moon_c.archive.001
```

The Default template has two fields separated by a period. The first field contains the name of the NetWorker server and the second field contains a number.

For example:

```
server.number  
space.675  
space.800  
atlas.054
```

The Default Clone template has two fields separated by a period. NetWorker uses this template for backup volumes belonging to the Default Clone pool. The first field contains the server name and the letter “c.” The second field contains a number.

For example:

```
moon_c.002
```

The three preconfigured label templates Full, NonFull, and Offsite use the same labeling conventions. The name of the label template appears in the first field, and the second field contains a number.

For example:

```
label name.number  
Full.076  
NonFull.003  
Offsite.120
```

The Two Sided template is for use with two-sided media such as optical media. When labeling two-sided media, you need to be able to label both sides of the media. The first field contains the name of the server, the second a number, and the third either an “a” or “b” to differentiate between the two sides of the media.

```
server.number.side  
phoenix.001.a  
phoenix.001.b
```

Preconfigured Directives

NetWorker is shipped with preconfigured directives. Each directive covers a set of the most important and most useful backup instructions.

The preconfigured directives are listed below.

- Unix standard directives—use for most of your backups and when you do not need one of the other specialized directives. This selection
 - applies the directive “+skip: core” to the *root* directory (/), thus skipping the backup of all *core* files
 - contains a *swapasm* directive to back up the relevant information about all NFS based and local swap files, but not the data in them
 - contains a *mailasm* directive to ensure that your mail files are backed up, yet not marked as read, and *logasm* for directories containing log files
- Unix with compression directives—use when you want to compress your backup data. Compressing client files saves you media space and network bandwidth but takes more time and CPU cycles on the client. Overall, the entire network may back up faster if all the clients compress their files, and parallelism is set appropriately.
- DOS standard directives—use to back up your DOS clients.
- NetWare standard directives—use to back up your NetWare clients.
- NT standard directives—use to back up your NT[®] clients.
- NT with compression directives—use to compress data on an NT client during backup.
- Index directives—use to back up the online file index. This option is usually used only by the **savegrp** command.

Preconfigured Notifications

NetWorker is shipped with several notifications: Bootstrap, Cleaning cartridge expired, Cleaning cartridge required, Device cleaned, Device cleaning required, Index size, Log default, Migration attention, Migration completion, Registration, Savegroup completion, Tape mount request 1, Tape mount request 2, and Tape mount request 3.

Most of these notices alert you to important NetWorker events. For example, if a group of clients did not complete a nightly backup, NetWorker sends you a savegroup completion notice by electronic mail.

Registration

The registration notification sends a message to *root* notifying you that your NetWorker products are not properly registered. You will receive the registration notification once a day or each time you start NetWorker. The notification message includes related information about each of the NetWorker products that are not registered correctly. A default registration action is shown in Figure B-7.

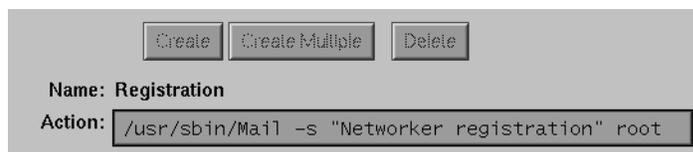
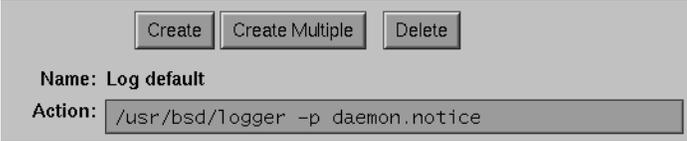


Figure B-7 Registration Notification

Log Default

The log default notification uses a UNIX facility called **syslog** to log and distribute notification about *all* NetWorker events. These events include requests for backup volume mounts, index size notices, and savegroup completion notices. How this information is distributed depends on how you have configured **syslog**. When NetWorker was installed, it created entries for logging and contacting operators. You can customize these entries. Refer to the `syslogd(1M)` reference page for information about configuring the distribution of log information. A default log action is shown in Figure B-8.

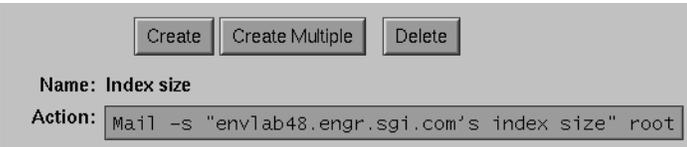


A screenshot of a configuration window. At the top, there are three buttons: 'Create', 'Create Multiple', and 'Delete'. Below the buttons, the text 'Name: Log default' is displayed. Underneath, the text 'Action: /usr/bsd/logger -p daemon.notice' is shown in a text input field.

Figure B-8 Log Default Notification

Index Size

NetWorker checks the size of its online indexes and sends a notification when it looks as if the indexes may run out of disk space. NetWorker automatically sends the electronic mail message to *root*. A default index size action is shown in Figure B-9.



A screenshot of a configuration window. At the top, there are three buttons: 'Create', 'Create Multiple', and 'Delete'. Below the buttons, the text 'Name: Index size' is displayed. Underneath, the text 'Action: Mail -s "envlab48.engr.sgi.com's index size" root' is shown in a text input field.

Figure B-9 Index Size Notification

The example above notifies you when the index for the client *atlas* is getting large. If you want the message to be mailed to someone other than *root*, you can edit Action and substitute *root* with a different user login name or mailing list.

Savegroup Completion

When NetWorker finishes backing up a group of clients, it sends a completion message via electronic mail to *root*. A savegroup completion action is shown in Figure B-10.



A screenshot of a configuration window. At the top, there are three buttons: 'Create', 'Create Multiple', and 'Delete'. Below the buttons, the text 'Name: Savegroup completion' is displayed. Underneath, the text 'Action: envlab48.engr.sgi.com's savegroup completion" root' is shown in a text input field.

Figure B-10 Savegroup Completion Notification

Backup Media Request Notices

When NetWorker needs backup media mounted for a backup, or a specific backup volume mounted to fill a recovery request, it displays a media request message in the NetWorker Administrator window. If no one fills the request, NetWorker sends another request after fifteen minutes. NetWorker sends a third request after another thirty-seven minutes, if no one fills the request.

The first mount request has a blank Action field, so the request appears only in the Pending display of the NetWorker Administrator window. The second mount request sends an alert to the logger, and the third request sends electronic mail messages to *root*. A tape mount request action is shown in Figure B-11.



Figure B-11 Tape Mount Request Notification

Summary

There are no rules for configuring NetWorker. The challenge is to understand how to best take advantage of the power and flexibility that NetWorker offers for your specific environment. Start using NetWorker with the preconfigured schedules and policies and then undertake small experiments. As your network of systems grows larger, as there are more and more files to back up, and as users see the advantages of NetWorker's fast file recovery, you will need to continue making adjustments.

Troubleshooting

This appendix contains troubleshooting information that addresses common questions concerning operating and configuring NetWorker.

Checking the NetWorker Daemons

If you have trouble starting NetWorker, the daemons may not be running properly. To check the daemons, enter the following command:

```
# ps -ef | grep nsr
```

The system displays output as shown below, showing these five daemons running.

```
111 ? IW    0:10 /usr/etc/nsrexecd -s localhost
116 ? S    176:15 /usr/etc/nsrd
158 ? IW    2:48 /usr/etc/nsrmmdbd
159 ? S    23:45 /usr/etc/nsrindexd
160 ? IW<  16:07 /usr/etc/nsrmmmd -n 1
```

If you discover that you need to start the NetWorker daemons, enter these commands:

```
# cd /
# nsrd
# nsrexecd
```

Displaying NetWorker

If you enter the **nwadmin** command and the NetWorker Administrator window does not appear, the DISPLAY variable on your system may not be set correctly.

To set the DISPLAY variable correctly, follow these steps.

1. Enter the following command at the system prompt for C shell or **tcsh**:

```
# setenv DISPLAY hostname:0.0
```

For a Korn shell or a Bourne shell, enter these commands:

```
# DISPLAY=hostname:0.0
# export DISPLAY
```

Replace *hostname* with the name of the machine where the user initially logged in.

2. Enter one of the following at the system prompt:

```
# xhost machine_name
```

or

```
# xhost +
```

Replace *machine_name* with the name of the machine where you are currently logged in, or the machine where you will log in.

3. Restart **nwadmin**.

Renaming a Client

NetWorker maintains an index for every client it backs up. If you change the name of the client, the index for that client is no longer associated with the client, and the client will not be able to recover any files backed up under its old name.

To change the name of a NetWorker client, you must first delete the old client name, then add the new client name, and rename the directory that contains the corresponding index. Follow these steps:

1. Create a new client, using the new name. Use the same configuration choices for the new client that you used originally for the old client.
2. Delete the old client.
3. As root on the NetWorker server, shut down the NetWorker daemons, using the **/etc/init.d/networker stop** command.
4. Change to the directory containing the client index directory. For example:

```
# cd /nsr/index
```
5. Delete the new client index directory (which is empty). For example:

```
# rmdir new_client_name
```

If you fail to remove the new client index directory, the old client index directory will be copied into the new client directory, as a subdirectory, with its old name.

6. As root on the NetWorker server, use the **mv** command to rename the client index directory. For example:

```
# mv old_client_name new_client_name
```
7. Restart the NetWorker daemons, using the **nsrd** and the **nsrexecd** commands. The daemon **nsrmmdbd** will rename all the instances of the save set under the old client name to the new client name.

Recover Access Issues

System administrators can control client recover access by configuring the client. The Recover access list in the Clients window displays the names of machines that can recover client files.

The following users have the ability to recover any files on any client:

- *root*
- *operator*
- a member of the *operator* group

Other users can recover only those files for which they have read permission, relative to the file mode and ownership at the time the file was backed up.

Files recovered by a user other than *root*, *operator*, or the *operator* group are created owned by that user.

Previewing a Backup

Every time you add a new client to NetWorker, it is a good idea to check if NetWorker can successfully back up the files for the new client. Use the Preview button in the Group Control window to see a “preview” of a group backup without actually backing up any files. You can also use the **savegrp -p** command at the system prompt to see a preview.

For example, this command shows you a backup preview of the clients assigned to the backup group *group_name*:

```
# savegrp -p group_name
```

If NetWorker cannot access a client in the backup group, check the following items:

- Make sure **nsrexecd** is running on the client machine and that it lists the hostname of the server in the command line. To make sure that **nsrexecd** is running, use the UNIX command **ps** on the client. See “Installing the Client Software” in the appropriate chapter for your platform for more information on **nsrexecd**.
- Using **nsrexecd** is the best method for backing up clients over the network. If you choose not to use **nsrexecd** and the clients cannot find the NetWorker binaries, add the location of the NetWorker binaries to the Executable path field in the Client window for each client. In other words, if the default PATH setup for root or Remote user does not include the appropriate path to the NetWorker binaries, add them to the client configuration. Display these attributes by choosing Details from the View menu.

Halting a Network Backup

To stop running a network-wide backup, click the Stop button in the Group Control window, shown in Figure C-1.

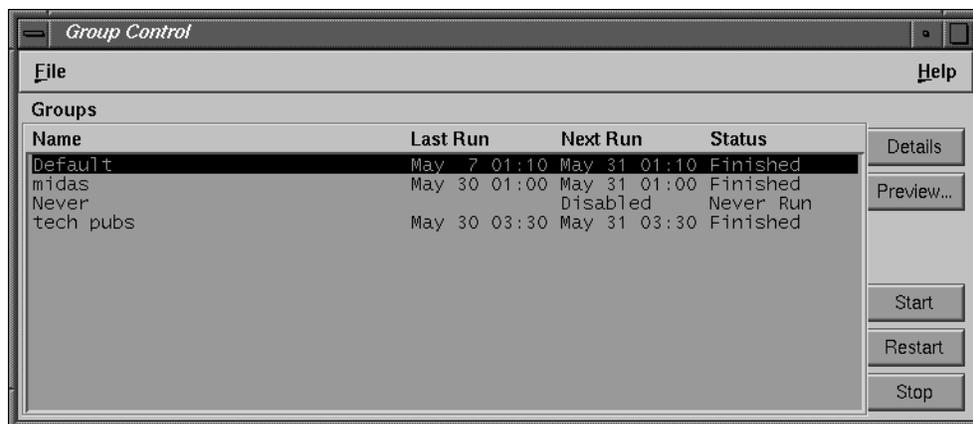


Figure C-1 Group Control Stop Button

The next network-wide backup will start as scheduled in the Start time field of the Groups window, or you may restart by clicking the Restart button in the Group Control window.

Backup Media Capacity

Occasionally you will find that NetWorker marks backup volumes as “full” when they are not really full. (The Volumes window and the output from the **mminfo -m** command display the details of the backup volumes.)

NetWorker marks magnetic tape “full” when it reaches the end of the tape or when there is a bad spot on the tape. For example, a backup tape that is reported as only “13% used” and marked as “full” has a bad spot on 13% of the length at the beginning of the tape. This tape can still be used for recoveries, but may not be used for any more backups.

If you see this “bad spot” behavior on many of the backup volumes, it may indicate the device needs to be cleaned or repaired.

Tapes are also marked “full” when they are recovered after being deleted from the media index.

Determining Jukebox Capacity

To find out how much space is available in the jukebox (autochanger), use either the Jukebox Mounting window or the **nsrjb** command. The Jukebox Mounting window displays all media in the jukebox and the percent used of each tape, as shown in Figure C-2.

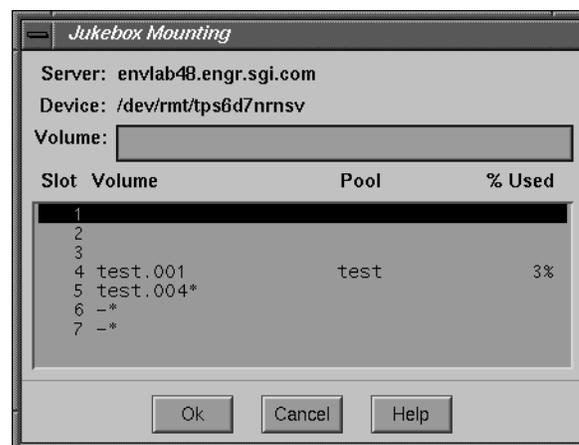


Figure C-2 Jukebox Mounting Window

If you prefer to use the **nsrjb** command, follow the steps below:

1. Switch user to root.
2. Enter the **nsrjb -v** command at the system prompt:

```
# nsrjb -v
```

NetWorker displays information about the backup volumes in the jukebox that looks similar to this:

```
Jukebox arc-db:
slot   volume      used    pool    mode
1:     moon.010
2:     moon.011    full   Default
3:     moon.012
4:     moon.013    full   Default

4 volumes, 2 less than 80% full.
2305 MB total capacity, 2200 MB remaining (5% full)
drive 1 (/dev/rmt/0hbn) slot 3: moon.012
```

Notice the information about the registered volumes, total capacity, and remaining capacity. This information tells you how much space is still available in the jukebox.

Savegroup Completion Messages

In the Notifications window, you configured NetWorker to mail the event notification about your savegroups. The Notifications window is preconfigured to mail the savegroup completion messages to *root*. This section contains descriptions of error messages that may appear in the savegroup completion mail. Possible solutions are included.

Binding to Server Errors

NetWorker is designed to follow the client/server model. In a client/server model, servers provide services to the client through the Remote Procedure Call (RPC). These services live inside of long-lived UNIX processes, known as daemons.

For clients to find these services, the services must be registered with a registration service. When daemons start up they register themselves with the registration service. In UNIX, the *portmapper* provides the registration service.

NetWorker servers provide a backup and recover service: they receive data from clients, store the data on backup media, and retrieve it on demand. If the NetWorker daemons are not running and a service is requested by **nwbackup**, **nwrecover**, or **mminfo**, for example, the following messages may appear in your savegroup completion mail:

```
Server not available
RPC error, remote program is not registered
```

These messages indicate the NetWorker daemons *nsrd*, *nsrindexd*, *nsrmmd*, *nsrmmdbd* might not be running.

To restart the NetWorker daemons, enter **nsrd** at the system prompt:

```
# nsrd
```

Saving Remote Filesystems

You may receive the following error message in your Savegroup completion notification when backing up a remote filesystem:

```
All: host hostname cannot request command execution
```

This means the **nsrexecd** on the client was not configured to allow the server hostname to back up its files.

You may also see this message:

```
All: sh: permission denied
```

This means **nsrexecd** is not running at all on the clients.

Make sure **nsrexecd** is running on the client machine and that it lists the server's *hostname* in the command line. To make sure that **nsrexecd** is running, use the UNIX command **ps** on the client. See "Installing the Client Software" in the appropriate chapter for more information on **nsrexecd**.

Using **nsrexecd** is the best method for backing up clients over the network. If you choose not to use **nsrexecd**, and the clients cannot find the NetWorker binaries, add the location of the NetWorker binaries to the Executable path field in the Client window for each client. In other words, if the default PATH setup for root or Remote user does not include the appropriate path to the NetWorker binaries, add the path to the client configuration. Display these hidden attributes by choosing Details from the View menu.

File Changed During Backup

NetWorker backs up the image that is in the filesystem at the time it comes across the file. NetWorker will notify you that the file was changed during the backup in the Backup Status window and the savegroup completion mail. You can back up the file manually after it has been closed, or wait until the next incremental backup.

Cannot Print Bootstrap Information

If your bootstraps are not being printed, you may need to enter the printer name as a hidden attribute using the following steps:

1. Open the Groups window and choose Details from the View menu.
2. Enter the name of the printer you are using to print the bootstrap in the Printer field.
3. Click Apply to save your changes.

Copy Violation

If you installed NetWorker on more than one server using the same NetWorker enabler code, you will receive the following messages in your savegroup completion mail:

```
--- Unsuccessful Save Sets ---
* quattro:/var save: error, copy violation - servers 'quattro' and
'spim' have the same software enabler code, '12345' (13)
* quattro:/var save: cannot start a save for /var with NSR server
'quattro'
* quattro:index save: error, copy violation - servers 'quattro' and
'spim' have the same software enabler code, '12345'
* quattro:index save: cannot start a save for /usr/nsr/index/quattro
with NSR server 'quattro'
* quattro:index save: cannot start a save for bootstrap with NSR server
'quattro'
* quattro:index /usr/etc/savegrp: bootstrap save of server's index and
volume databases failed
```

To complete a backup, you must kill the NetWorker daemons on both servers, de-install NetWorker from the extra server(s), and restart the NetWorker daemons on one server.

1. To kill the NetWorker daemons, log in to the NetWorker servers as *root* and enter the following command on all servers that have NetWorker installed:

```
spim# /etc/init.d/networker stop
...
quattro# /etc/init.d/networker stop
```

2. Use the following command to de-install NetWorker on the server(s) that you will not be using as a NetWorker server:

```
spim# versions remove networker4
```

3. Finally, restart the NetWorker daemons on one server with these commands:

```
quattro# nsrd &
quattro# nsrexecd &
```

Maximum Filename Length

NetWorker supports a maximum filename size of 1024 characters. This is the same as the UNIX *svid* limitation.

Savegroup Completion Warning Messages

Occasionally the savegroup completion message includes one or more messages. These messages contain information that help the administrator understand why NetWorker performs certain tasks.

Below is one of the messages you might see:

```
quattro:/usr no cycles found in media db; doing full save
```

In this example, the filesystem, */usr*, on the client *quattro* has no full saves listed in the media database. Therefore, despite the backup level pre-selected for that client's schedule, NetWorker will perform a full backup. This feature is important because it allows you to perform disaster recoveries for that client.

This message may also appear if the server and client clocks are not synchronized. To avoid this, make sure the NetWorker server and client

- are in the same time zone
- have their clocks synchronized

The following savegroup message may also appear:

```
NetWorker_server:index Saving server index because server is not in an
active group
```

If your server belongs to a group that is not enabled, NetWorker will, to avoid a long recovery process, save the server bootstrap information along with this group. As soon as possible, enable the group to which your NetWorker server belongs.

X11 Errors

The following error message may appear when the **nwadmin &** command is executed:

```
Xlib: connection to "client:0.0" refused by server
Xlib: Client is not authorized to connect to Server
X error: Cannot open display on window server: client:0.0 (Server pkg)
```

This indicates that the client is not authorized to display NetWorker.

To correct this situation do the following at the client machine:

```
client% xhost NetWorkerServer
```

Remotely log in to the NetWorker server and run the following command at the server prompt:

```
% setenv DISPLAY client:0.0
```

For the Korn shell or the Bourne shell, use the following commands:

```
# DISPLAY=client:0.0
# export DISPLAY
```

Moving Indexes

Because the index databases are holey files, **cp** creates a file that consumes more disk space than the original file. To move indexes, execute the following command in the */nsr/index* directory:

```
# uasmm -si clientIndexDirectoryName | (cd target_dir; uasmm -rv)
```

Recovering Files From an Interrupted Backup

You cannot recover files from a backup terminated by killing the NetWorker daemons because the media index was not updated before the daemons exited. Consequently, NetWorker does not know on which volume the requested file is located.

Determining the NetWorker Server

If you start NetWorker from a remotely mounted directory, you may receive the following message:

```
Using server server_name as server for client_name.
```

NetWorker looks for the system that is the fileserver of a remotely mounted directory and uses the NetWorker server assigned to that system as the backup server. To bypass this message, start NetWorker from a local filesystem.

Using nsrexecd

The **nsrexecd** daemon runs on NetWorker client machines. This daemon provides a secure and restrictive way for NetWorker to start automatic backups on clients. The **nsrexecd** daemon allows you to restrict access to a select set of NetWorker servers. When you install NetWorker on a client, **chkconfig** automatically turns NetWorker on, so **nsrexecd** will be started each time the client reboots. Security is increased by the use of a challenge/response scheme to ensure that only the NetWorker server is initiating connections, and not another program.

The file modified for each client type is shown in the table below. If you ever need to reconfigure **nsrexecd**, for example, to allow a different NetWorker server to back up the client, edit the appropriate file on the client, make the changes to the **nsrexecd** startup command (see the `nsrexecd(1M)` reference page for a description of the command-line configuration options), and restart **nsrexecd**.

Make sure you enter the **nsrexecd** command in exactly the same way as it is listed in the boot-time file, complete with all command-line options. Alternatively, on non-IRIX systems you can use **nsr_ize -c -u** to deinstall the client software, entering **no** whenever it asks you questions. Then, use **nsr_ize -c -i** and follow the instructions as if you were performing an NFS client install.

Table C-1 shows the location of boot-time files on different operating systems.

Table C-1 Where to Start nsrexecd

Operating System Type	Boot-Time File
AIX	<i>/etc/rc.nsr</i>
HP-UX	<i>/etc/rc</i> or <i>/sbin/init.d/networker</i>
IRIX	with <i>chkconfig(1M)</i>
SCO	<i>/etc/rc2.d/S95networker</i>
Solaris 2.x	<i>/etc/rc2.d/S95networker</i>
SunOS 4.1.x	<i>/etc/rc.local</i>
Ultrix	<i>/etc/rc.local</i>
others	<i>/etc/rc2.d/S95networker</i>

Command Summary

This appendix contains a list of the most commonly used NetWorker commands entered at the system prompt and a table listing the NetWorker maintenance commands.

NetWorker System Commands

If you are not using the X Window System, or are using NetWorker from an ASCII terminal, the following commands are the ones you use most often:

Use:	To:
mminfo	Display information about the backup volumes and save sets.
nsradmin	Configure networker resources.
nsrclone -c -N	Perform a super-full backup.
nsrck -F	Compact the index size after purging or deleting backup volumes.
nsrexecd	NetWorker client execution daemon.
nsrhmck	Checks that stubs and filenames are consistent.
nsrinfo <i>client</i>	Lists contents of client file index.
nsrjb -l <i>volume_name</i>	Load and mount a backup volume in the jukebox.
nsrls	Display information about the server online indexes.
nsrmig	Migrate files meeting defined criteria or rules.
nsrmm -d <i>volume_name</i>	Delete the backup volume named <i>volume_name</i> .
nsrmm -d -P <i>volume_name</i>	Purge the file index entries from both the file index and the media index for the backup volume named <i>volume_name</i> .

- nsrmm -l** *volume_name*
Label a backup volume with the name *volume_name*.
- nsrmm -m** Mount the backup volume.
- nsrmm -m -l -R** Recycle and mount backup volume.
- nsrmm -u** Unmount the backup volume.
- nsrpmig** Pre-migrates files meeting defined criteria or rules.
- /etc/init.d/networker stop**
Shut down the NetWorker daemons and processes dependent on those daemons.
- nsrwatch** Display the character-based NetWorker status monitor.
- savegrp** *group_name*
Start the backup of the clients in the group *group_name*.
- savegrp -p** *group_name*
See a preview of the backup of the group named *group_name*.

NetWorker Maintenance Commands

Table D-1 lists the NetWorker maintenance commands.

Table D-1 Maintenance Commands and Reference Pages

Command	Reference Page	Description
jb_config	jb_config(1M)	jukebox resource configuration program
jbexercise	jbexercise(1M)	jukebox diagnostic program
mmrecov	mmrecov(1M)	command to recover the online indexes of a NetWorker server
nsrmm	nsrmm(1M)	NetWorker media interface command
mminfo	mminfo(1M)	NetWorker media and save set information reporting command
mmlocate	mmlocate(1M)	accesses and manages the backup volume location information contained in the media database
mmpool	mmpool(1M)	NetWorker media pool reporting command

Table D-1 Maintenance Commands and Reference Pages

Command	Reference Page	Description
nwadmin	nwadmin(1M)	window-based display of NetWorker server status
(info only)	nsr(1M)	guide to using NetWorker
nsradmin	nsradmin(1M)	NetWorker character-based program for system administration
nsrck	nsrck(1M)	NetWorker check and repair program for the server indexes
nsr_crash	nsr_crash(1M)	how to use NetWorker to recover from crashes
nsrclone	nsrclone(1M)	save set cloning command
nsrd	nsrd(1M)	NetWorker server daemon
nsrexecd	nsrexecd(1M)	NetWorker client execution daemon
nsrshmck	nsrshmck(1M)	checks and corrects inconsistencies in HSM migrated files
nsrim	nsrim(1M)	NetWorker index management program, usually invoked from the savegrp command
nsrindexasm	nsrindexasm(1M)	module for saving and recovering NetWorker indexes
nsrindexd	nsrindexd(1M)	NetWorker file index daemon
nsrinfo	nsrinfo(1M)	file index reporting command
nsrjb	nsrjb(1M)	jukebox control command
nsrls	nsrls(1M)	lists the statistics of NetWorker index files
nsrmig	nsrmig(1M)	migrates files for long term storage
nsrmmmd	nsrmmmd(1M)	NetWorker media management daemon
nsrmmdbasm	nsrmmdbasm(1M)	module for saving and recovering NetWorker media databases
nsrmmdbd	nsrmmdbd(1M)	NetWorker media index daemon
nsrpmig	nsrpmig(1M)	pre-migrates files for long term storage
nsrretrieve	nsrretrieve(1M)	retrieves archived save sets

Table D-1 Maintenance Commands and Reference Pages

Command	Reference Page	Description
nsrwatch	nsrwatch(1M)	displays the NetWorker server status from an ASCII terminal
recover	recover(1M)	command to browse the online indexes and recover files from the NetWorker server
save	save(1M)	command to save files to the NetWorker server
savefs	savefs(1M)	command to save filesystems to the NetWorker server (see also savegrp)
savegrp	savegrp(1M)	command to save the files of a group of NetWorker clients
scanner	scanner(1M)	command to read the contents of a volume to recover from NetWorker server crashes
tapeexercise	tapeexercise(1M)	command to exercise a tape drive in order to uncover problems
uasm	uasm(1M)	NetWorker module for saving and recovering generic UNIX files

Frequently Asked Questions

This chapter covers some common problems and error messages reported by NetWorker customers. Questions and answers were developed by Pat Darisme and Doug Murphy, with help from Sandeep Cariapa, Luc Chouinard, and Steve Modica.

Licensing and Installation

The most common problem reported during NetWorker installations is that people forget that, in addition to installing NetWorker licenses with **LicenseManager**, you must also run **nsrcap** to initialize licenses. Without **nsrcap**, the following error message appears in the **nwadmin** window:

```
Server is disabled (Install base enabler)
```

Run the *nsrcap* command to reread the license configuration file and enable the server.

```
# /usr/etc/nsrcap -v
```

Q: The NetWorker server does not work, and the error message “No server available” keeps appearing. What can I do?

A: Check that all required subsystems are installed. See the *IRIX NetWorker Installation Guide* for a list of required server subsystems, and for a list of required client subsystems. After installing any new subsystems, stop and restart the NetWorker daemons, and reinitialize licensing information with **nsrcap**:

```
# /etc/init.d/networker stop  
# /etc/init.d/networker start  
# /usr/etc/nsrcap -v
```

Q: How can I tell which licenses I have enabled on the machine?

A: In the **nwadmin** window, choose Registration from the Server menu. The Registration window appears, containing a list of enabled licenses.

Q: When I run **nsrkap -v**, the following error message appears:

```
nsrkap: RPC error, Program not registered
```

A: You must launch the NetWorker daemons. The following command shows if the daemons are present:

```
# ps -ef | grep nsr
root  704  527  0  Sep 14 ?        0:00 /usr/etc/nsrmmdbd
root  527    1  0  Sep 14 ?        0:46 /usr/etc/nsrd
root  543    1  0  Sep 14 ?        0:01 /usr/etc/nsrexecd
```

If the NetWorker daemons are not present, you can start them with this command:

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

Finally, ensure that NetWorker daemons will be launched at boot time by doing:

```
# chkconfig networker on
```

Q: When trying to license an autochanger (jukebox), this error message appears:

```
The Jukebox "Jukebox Name" is not properly licensed
Make sure the jukebox is properly enabled and authorized
```

A: Using the **LicenseManager** command, install the appropriate autochanger license, and then rerun the **nsrkap -v** command.

Q: What is the difference between NetWorker server edition and network edition?

A: The server edition is intended for backing up a single machine, perhaps a database server. No client connections are allowed, meaning that server edition cannot back up other machines over the network. Server edition can be upgraded to network edition, which allows up to 10 client connections, meaning that NetWorker can back up 10 client machines over the network, or 9 plus the server itself.

Q: What is the difference between a client connection and ClientPak?

A: Client connections are licensed objects, whereas ClientPak is a software product distributed on CD-ROM. Out of the box, NetWorker network edition allows 10 client connections, allowing you to back up machines over the network, plus the server itself. Additional client connections can be purchased at any time, in units of 5, 25, and 100.

ClientPak products provide NetWorker clients for heterogeneous networks of machines, including most major hardware platforms and operating systems.

Q: Whenever I try to modify NetWorker settings I get this message:

```
Modify failed: user name@host needs to be on administrator list
```

A: The root user (or other user on the administrator list for that NetWorker server) must add you to the list of privileged users. They must choose **Server Setup** from the **Server** menu, then modify the **Administrator** field.

Q: What licenses are available for NetWorker?

A: See the *IRIX NetWorker Installation Guide* for a list of NetWorker license options.

Q: I always see the following error messages when I launch NetWorker. Why?

```
Warning: Could not find Scheme file: SgiSpec  
Warning: Unmatched quotation marks in string "", remaining fonts in list unparsed
```

These messages are normal on systems without product image *desktop_eoe.sw* installed, and can be ignored.

Tapes and Devices

Q: What device do I specify for backups?

A: The selection of an *nrnsv* (non-rewinding non-byte-swapping) device is critical. The **nwadmin** window's **Devices** panel often shows a *nrnsv* device selected, or if hardware compression is needed, a *nrnsv* device. Thus the panel reads something like this:

```
/dev/rmt/tps1d6nrnsv      8mm 5GB (mounted)
```

Use of a rewinding tape device can result in severe problems restoring backups, such as being able to restore only 3 MB of a 10 GB save set. If this accident occurs, your backup tapes are useless. Change the backup device by choosing **Devices** from the **Media** menu and clicking the **Create** button to designate an *nrnsv* device. Then relabel all tapes and perform a full backup as soon as possible.

Q: Although I know the tape is not full, when launching a backup I get this message:

```
No Space Left On Device
```

A: NetWorker sometimes issues this message if it encounters write errors. This forces you to change tapes in order to ensure the safety of backups.

The probable cause of write errors is that the tape drive needs cleaning. A general rule of thumb is that a tape drive should be cleaned after every 30 to 50 hours of use.

Also bear in mind that the tape might actually be full. There is some overhead associated with NetWorker backups; in some cases it can be as much as 30%. This can be caused by empty tape space being used to keep the tape running at constant speed during backup, especially when the network is slow or for other reasons when data blocks arrive empty.

Q: Even though the tape is not write protected, when mounting a tape I get this message:
Tape write protected

A: After failing to write data to a write-protected tape, NetWorker marks the tape in the media index as full. Even if the tape is nowhere near full, NetWorker says it is, as a way to deal with write protection. When you later make the tape physically read/write, you can no longer append to this tape, since the media index claims that it is full.

Fortunately the data on the tape is still readable and recoverable. Before proceeding, it would be wise to recover all data from this tape. Alternatively, you could retain the tape until all the save sets contained on it have expired.

The only way to use this tape again with NetWorker is to “recycle” it as a new tape. One way of doing this is to remove the index references to the tape so that NetWorker no longer thinks it is full. This can be done with the following command:

```
# nsrmm -dV volume
```

This deletes the volume indexes, thus removing the volume from the media index and making the tape writable again. Warning: the command above makes it impossible to recover any existing save sets from that volume.

Q: Do tapes have to be relabeled before NetWorker starts rewriting to the new tapes?

A: Yes, relabeling a tape is an effective way of rewriting old tapes. However you can also recycle old tapes using the Volume Management window. Choose Volumes from the Media menu, then select a volume in the Volume Management window and choose Recycle from the Volume menu. This recycles a tape, allowing you to write to it again.

Q: In what format does NetWorker write data to the tape?

A: NetWorker writes data in XDR format (external data representation), with the media and file indexes before and after a save set. Thus tapes are readable only by NetWorker.

Autochangers (Jukeboxes)

Q: What kind of autochanger can I use with NetWorker? Will my existing jukebox work with NetWorker?

A: Consult the NetWorker release notes for the currently supported list of jukeboxes, since this list is dynamic and depends on the version of NetWorker and IRIX that a system is running. Another way to obtain a list of supported jukeboxes is to run the **jbm_enabler** command, which outputs a list of jukeboxes that work, usually without any need to add special drivers into the kernel. The last entry is called "other" and may be used to configure custom support for hitherto unsupported jukebox.

If you have Web access, you can obtain the list of supported jukeboxes at this address: <http://www.sgi.com/Products/hardware/challenge/Networker/NetWorker.html> (look at the answers to frequent questions). Silicon Graphics recommends that only these supported devices be used with NetWorker. If you buy a jukebox from a third-party vendor, make sure that the firmware version matches the version recommended. Silicon Graphics supports jukeboxes from third-party vendors, if the firmware matches.

Q: When I run **jbm_enabler**, the following error message appears:

```
jbm_enabler: RPC error, Program not registered
```

A: This usually means one of two things: either the NetWorker daemons are not running, or the customer has not installed (or has improperly installed) permanent licenses for the Autochanger option. Each jukebox requires a separate autochanger license. Autochanger licenses come in many sizes; see the marketing and product information for details. To check the NetWorker daemons, try the following test:

```
# ps -ef | grep nsr
root  704   527  0   Sep 14 ?           0:00 /usr/etc/nsrmmdbd
root  527     1  0   Sep 14 ?           0:46 /usr/etc/nsrd
root  543     1  0   Sep 14 ?           0:01 /usr/etc/nsrexecd
```

If the NetWorker daemons are not present, you can start them with this command:

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

If the daemons have already started, then check the */var/FLEXlm/license.dat* file and verify that the licenses are installed and that they have not expired. NetWorker ships with temporary licenses for extra-cost options such as Jukebox and Archive, to assist you in evaluating the product.

A licensing problem typically results in this error message inside the **nwadmin** window:

```
The jukebox 'model_name' is not properly licensed.  
Make sure the jukebox is properly enabled and authorized.
```

To help pinpoint a licensing problem, there are several things you can try. First run the **jbexercise** command to test jukebox connection and activity.

In the **nwadmin** window, choose Jukeboxes from the Media menu; then in the Jukeboxes window, choose Details from the View menu and verify that the Enabler code is present. Click the Label button and attempt jukebox labeling; helpful error messages may result. Try deleting and (re)creating the jukebox, by choosing Devices from the Media menu.

Once the jukebox is licensed, reinventory with the **nsrjb -IE** command.

Q: Where is the configuration file for my jukebox?

A: The jukebox configuration file */nsr/res/nsrjb.res* contains licensing, scheduling, and current use information. This file is the roadmap for jukebox operation, and should not be modified, although sometimes it must be removed to force reinstallation of a jukebox (be sure to save a copy before removal).

Q: My jukebox is confused. It seems to expect some tapes to be mounted but cannot find them, or it continuously asks that I change tapes in the slots.

A: A jukebox's status can become confused for various reasons throughout operation. There are two ways around this. The first is to empty the jukebox and inform NetWorker that it is empty, then fill it up again. The second is to look at the *nsrjb.res* file, figure out what NetWorker expects in the jukebox, and put in what it wants.

The status of the jukebox is kept in the */nsr/res/nsrjb.res* file, where two lines are critical to the state of the jukebox:

```
loaded slots: 5;  
loaded volumes: pipedream.005;
```

These lines indicate that slot 5 is loaded and contains the volume *pipedream.005*. If this is not the actual state of the jukebox, then NetWorker gets confused. To alleviate this situation, stop the NetWorker daemons:

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

Edit the two lines so that they read:

```
loaded slots: ;
loaded volumes: "";
```

Remove tapes from the jukebox device, and restart the NetWorker daemons:

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

Finally, run **nsrjb -IE** to reinventory the jukebox and return it to a stable state.

To avoid having to do this repeatedly, you might want to make a copy of the *nsrjb.res* file, for example *nsrjb.res.orig*, when the jukebox is in a stable state. The next time the jukebox gets confused, replace the seemingly corrupted *nsrjb.res* file with the good copy, after stopping the NetWorker daemons:

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

If you are familiar with Legato NetWorker on other platforms, note that the **nsrjb -H** option (hardware reset) is not supported in IRIX NetWorker. However, by opening the door of a jukebox or powering it down, you can get a similar effect.

Q: How do I change the cleaning slot on my jukebox?

A: Choose Jukeboxes from the Media menu. In the Jukeboxes window, change the slot number in the Cleaning slots field, and adjust the number of Available slots accordingly.

The cleaning slot entry corresponds to the loaded volumes entry in the *nsrjb.res* file. Here we see an error:

```
volumes: -, -, -, -, -, -, Cleaning Tape (5 uses left);
...
cleaning slots: 1;
```

In this case, either the 7th volume should be moved to first place, or the cleaning slot should say 7, depending on which hardware slot actually contains the cleaning tape.

Q: I started my backup last Friday. When I came in Monday, a message in the **nwadmin** window says NetWorker is waiting to mount a writable tape. What happened?

A: The advantage of a jukebox is that the administrator can set up a variety of backup schedules that do not require anybody to insert new tapes into a backup device. It is still possible to create schedules that require more tapes than a jukebox can hold. For answers on setting up schedules, see the section "Backup Schedules" on page 284.

Q: I had a jukebox license before, and I had to update it. Now when I run **nsr`cap` -v**, I get an error message saying the machine cannot find a valid license. What happened?

A: If you look at the `/nsr/res/nsrjb.res` file, you'll notice that it contains current license information for the jukebox. If the information in `nsrjb.res` does not match the license information in `/var/flexlm/license.dat`, then `nsrjb.res` wins. You might need to edit these to files by hand to bring them in sync. Then restart NetWorker and rerun **nsr`cap`**:

```
# /etc/init.d/networker start
# /usr/etc/nsrcap -v
```

Q: My jukebox cannot seem to mount volumes for backup or recovery. What's wrong?

A: If backup and recovery has worked in the past (you eliminated configuration errors), look at the `/nsr/res/nsrjb.res` file. Make sure entries in the file indicate all instructions the jukebox is getting from the NetWorker GUI tools or from the command line. Look for discrepancies between entries in this file and the perceived setup. This might give you some hints about what went wrong.

Q: I have 100 GB of data—why is a 7-slot DLT jukebox unable to hold this much data?

A: We talk more about compression directives later, but be aware that compression with NetWorker is optimized by choosing either software or hardware compression. If you try compressing files that have already been compressed, you will see a degradation in performance of your backup device. Also, make sure what you are backing up is in fact what you think you are backing up. Having directories NFS-mounted or CDRom drives on clients could possibly affect what NetWorker tries to back up.

Q: When should I inventory my jukebox? Should I ever reinventory my jukebox?

A: Choose Inventory from the Media menu whenever you have changed the set of tapes in the jukebox, or when NetWorker has trouble finding the correct volume to mount, which sometimes occurs if labels or indexes become corrupted. It never harms NetWorker to reinventory the jukebox.

Q: When performing a full backup, the jukebox loads tape 1, then tape 2, but then jumps to tape 7. What gives?

A: This is normal. NetWorker arbitrarily picks the order for writing tapes in a jukebox.

Q: What is the proper syntax for running the **jbexercise** command?

A: This command requires 2 arguments: the SCSI port of the jukebox controller, and the model number of the jukebox. To obtain a list of supported jukebox models, run the **jbexercise** command without arguments. Here is an example with proper syntax:

```
jbexercise -c /dev/scsi/sc6d610 -m "Silicon Graphics DLT2700"
```

Q: When adding a jukebox, how do I know what control port the jukebox device is on? What is the "pathname of the control port" for a jukebox device?

A: Run the **hinv** command, which displays something like this:

```
Jukebox: unit 6, lun 1 on SCSI controller 5
```

This indicates that the jukebox device name would be */dev/scsi/sc5d610* (sc = SCSI controller, d = device unit, l = logical unit).

Q: How do I know the "pathname of media drive 1" to select a tape drive in the jukebox?

A: Looking at the */dev/rmt* directory entries for the SCSI controller and device of the jukebox (with **hinv** output) you can pick the appropriate tape device. As per the previous example: *tps5d6nrnsv* or *tps5d6nrnsvc* for hardware compression.

Q: Why did NetWorker just disable my jukebox device?

A: Sometimes NetWorker disables a device after encountering an error. For instance, trying to move a tape into a full drive results in a disabled jukebox. To reset in **nwadmin**, choose Devices from the Media menu, and click the Enabled field's Yes button.

Q: In what format does NetWorker write data to the tape?

A: NetWorker writes data in XDR format (external data representation), with the media and file indexes before and after a save set. Thus tapes are readable only by NetWorker.

Q: Where can I get more information on DLT Jukeboxes?

A: Look at the Web site

http://www.sgi.com/Products/Mass_storage.html

Q: Where can I get more information on Exabyte Jukeboxes?

A: Look at the Web site

http://www.sgi.com/Products/appsdirectory.dir/DeveloperIXExabyte_Corporation.html

Backup Schedules

Q: What's the difference between a level 5 and a level 9 backup?

A: A level 5 backs up more stuff, potentially. A level 1 backs up everything since the last full backup, even if level 1-9 backups have occurred in the meantime. A level 5 backs up everything to the last level 1-4 or full backup. Backup schedules using various levels are a good way to reduce the number of tapes needed for disaster recovery, while decreasing the frequency of time-consuming full backups. For more information about scheduling backups of different levels, see "Using Level Backups" on page 233.

Q: How many tapes do I need for a complete recovery of my client if it crashes?

A: That depends on the backup schedule you set for the client. The default schedule is for full backup on Sunday, and incrementals the remainder of the week. Under this schedule, the number of tapes needed for recovery increases each day, from 1 on Sunday to 7 on Saturday. (Incrementals could be contained on the same tape, although this provides poor protection against fire). If you set a schedule with full backup on Sunday, and level 9 backups the remainder of the week, the number of tapes needed for recovery should never exceed 2. (Of course level 9 backups would take longer than incrementals, especially at the end of the week.)

Q: How long should it take for NetWorker to backup my clients?

A: This depends on the device you use, how your schedules are staggered, the network speed, parallelism, and other factors. If you assume a rate of 1MB/second for a DLT drive, then it should take close to 3 hours to backup 10GB of data. If you find that your backups are taking too long, then investigate factors such as network speed, CPU speed, tape drive speed, NetWorker parallelism, and concurrent devices.

Browse and Retention Policies

Q: What is the difference between Indexes and Save Sets?

A: An index is a browsable directory tree indicating which files are available for recovery. Although they are copied to tape, browsable indexes are stored on server disk, and thus do not need to be retrieved from tape. They contain information about files backed up for each NetWorker client. A save set is a bunch of files backed up on tape.

Since indexes consume lots of disk space, they are typically browsable (stored on disk) for much less time than save sets are retained on tape. The browse policy governs the longevity of client indexes, whereas the retention policy controls the minimum age of save sets when a tape gets marked as recyclable.

Q: How do I define browse and save set retention policies?

A: Both the browse policy and the retention policy may be set by choosing Client Setup from the Clients menu. The default browse period is one month; the default retention period is one year. Although there is seldom a need for doing so, policy periods may be set by choosing Policies from the Customize menu.

Q: How can I keep a backup longer than the retention policy? What is archiving?

A: You can make a clone of a backup volume by choosing Clone Volume from the Media menu. You can make a clone of a save set by choosing Clone from the Save Set menu. Cloned volumes and save sets can be stored forever, and recovered by means of the **scanner** command.

Archiving involves the long-term storage of data on media called archive volumes. Archive volumes are like backup volumes, except their retention period is unlimited, and archives are not tracked in the file index. Archiving requires an Archive license, and works with either server or network edition.

Backing Up Data

Q: Where do I define which directories to back up? How can I control the contents of a save set?

A: Each client's save sets may be specified by choosing Client Setup from the Clients menu. In the Save set field, indicate which directories and/or files you wish to have backed up on that client.

Q: What's the difference between the media index and the file index?

A: The media index contains information about tapes, whereas the file index contains information about backed-up files and directories. As long as files remain in the online file index, they are browsable and recoverable from the **nwrecover** window.

The media index works in tandem with the file index to tell NetWorker which volume contains a particular saveset. As long as volume entries remain in the media index, save sets on that volume are recoverable by choosing Recover from the **nwadmin** Save Set menu. Both the media index and the file index exist online, and are also copied to tape (twice) during each scheduled backup.

Q: What does labeling a tape do?

A: The tape label is how NetWorker (and the operator) identify a tape and determine which save sets reside on which tape. The physical label and the magnetic label should match. When you relabel a tape, you remove all entries for that volume from the media index. This makes it impossible to recover data from that tape using the Save Set Recover window.

When a tape is marked recyclable, you must relabel it before overwriting it.

Q: How can I check the status of a magnetic tape label? How do I verify a label?

A: To verify a tape label, or to see if it is corrupted, try the following command:

```
# nsrmm -v -p -f /dev/rmt/tpsXdYnrnsv
```

Replace X with the SCSI controller number, and Y with the device number of the drive. If the label is good, **nsrmm** should display output similar to the following:

```
/dev/rmt/tps0d3nrnsv: verifying label, moving backward 2 files
/dev/rmt/tps0d3nrnsv: verified label of vol_test2
Verified 4mm tape vol_test2 on /dev/rmt/tps0d3nrnsv
expiration time is Tue Jun 16 16:43:52 1998
```

Q: I am unable to launch a group backup from the **nwadmin** window. Why?

A: Try launching this backup from the command line by typing:

```
# savegroup group_name
```

This might provide debugging information that can help you isolate the problem.

Q: How do I force an immediate launch of a backup without waiting for the scheduled time to arrive?

A: In the **nwadmin** window, click the Group Control button (or choose Group Control from the Server menu), select a group from the Groups list, and click the Start button.

Q: I keep getting the following error message. Why?

```
Permission denied... backup failed.
```

A: The client needs to be informed that the NetWorker server needs to read its files. On the client system, run **chkconfig networker on**, which starts the **nsrexecd** daemon to perform backups. It is possible to modify to modify */.rhosts* to grant NetWorker server access, but this method is less secure than **nsrexecd**.

Q: I keep getting the following error message. Why?

```
root is not an authorized user.
```

In the Remote access field, you also need to add the username of the agent (probably root) who runs the backup remotely.

Recovering Data

Q: When trying to recover files, I get the following error message. Why?

```
hostname is not a registered client.
```

A: Automatic server detection has located a server on which the current system is not a registered client. If a different server is appropriate, run the command again with the **-s** option to specify server. If the server is correct, register that client on the server. In **nwadmin**, choose Client Setup from the Clients menu. In the Clients window, click Create, and type the client name in the Name field.

Q: I keep getting the following error message. Why?

```
user root on machine server is not on server.FQDN's remote access list
```

This is usually the result of not having both a bare hostname and a fully qualified domain name (FQDN) for the NetWorker server. In **nwadmin**, choose Client Setup from the Clients menu, select the server in the Clients window, scroll down to the Aliases field, and add the bare hostname as an alias for the server. If the aliasing information is correct, check the Remote access field for accuracy.

Q: Where do I find the browsable indexes?

A: If indexes are still browsable, you can see them by choosing Indexes from the Clients menu in the **nwadmin** window, and users can browse them in the **nwrecover** window.

Q: How do I recover files if the file index gets erased or their browse period expires?

A: The media index, which is not affected by the browse period, contains information about which save sets are on each volume. You can browse the media index by choosing Volumes from the Media menu; in the Volume Management window, locate the save set you want. Then choose Recover from the Save Set menu; in the Save Set Recover window, select the correct version and click the Recover button.

Q: How can I recover files without using **nwrecover**?

A: To avoid the **nwrecover** window, you can use the **mminfo** and **recover** commands in combination. First run **mminfo -v** to print the Save Set IDs (ssid) list, and choose the one you wish to recover. Then run **recover -S ssid** to interactively recover files.

```
# mminfo -v
volume      client      date      time      size      ssid fl  lvl name
test.006    envlab48   9/23/96   20:16:58  31 MB    12821 cb full /n/nw-test-dir
# recover -s servername -S 12821
```

You can recover specific directories from specific savesets without going through the **recover** interface by using the **scanner** command:

```
# scanner -s ssid /dev/rmt/tpsXdXnrnsv -x uasm -vr /full/path/dirname
```

You can rename the recovery directory by specifying **-m /full=/fullnew** (for example). You can also use **scanner** to recover individual files, or the Save Set Recover feature of **nwadmin**. Before recovering any client files, run **nsrck** first. If there are no problems with the index, it should exit quickly:

```
# nsrck clientname
```

Q: The **mminfo -v** command displays nine fields: volume, client, date, time, size, ssid, fl, lvl, and name. What does the “fl” field indicate, and what do its values represent?

A: The “fl” (flags) field indicates whether a save set spans multiple volumes.

```
cr  complete, recoverable
hr  just the head, recoverable
tr  just the tail, recoverable
```

Q: How can I recover a single file using **nwadmin** (instead of **nwrecover**)?

A: In **nwadmin**, choose Recover from the Save Set menu.

In the Save Set Recover window, click the Recover button (at the bottom right). In the Save Set Recover Status window, type the full path name of the file that you wish to recover in the Paths to recover field, then press [Enter] to register your selection. If you wish to relocate or overwrite the file, click the Options button. Finally, click the Start button to initiate recovery.

Q: When a file is recovered, does its modification time get changed?

A: No, a recovered file retains the modification time it had when it was backed up.

Q: How can I determine which save sets are on a tape?

A: To list the save sets on a tape, run the **scanner** command:

```
# scanner /dev/rmt/tpsXdXnrnsv
```

Q: How could I list the files in a particular save set, such as #8738783?

A: Run the **scanner -s** command, piped to the **uasm** command with **-rvn** option:

```
# scanner -s 8738783 /dev/rmt/tpsXdXnrnsv | uasm -rvn
```

Q: How do I rebuild the online file index for a save set, such as #8738783?

A: Run the **scanner** command with **-i** and **-s** options:

```
# scanner -i -s 8738783 /dev/rmt/tpsXdXnrnsv
```

Q: How can I recover files if the media index has been corrupted or removed?

A: First get the bootstrap ID of the most recent full backup (use the **mminfo -v** command) then rebuild the media index (use the **scanner -i** command, which rebuilds both online file and media indexes). Of course recent changes to files will be lost.

Q: Can I recover files after their retention period has expired?

A: Perhaps, depending on how you set up schedules. Normally NetWorker does not reuse volumes until one backup cycle has completed. So before giving up, check your most recent backups volumes and you might find some tape labels that have not yet been rewritten. Of course, once a tape has been relabeled, not just marked for recycling, old data is no longer recoverable. Browse files in the save set to determine if you wish to recover them. Use the **recover -a** command to force an automatic (non-interactive) recover of the entire save set without prompting for each file.

Directives

Q: Where is the documentation for using directives?

A: The `nsr(4)` reference page contains complete information about directives.

Q: What is a directive?

A: A directive is an improved method for backing up special files. Directives may be stored in the `.nsr` file of any directory, and are parsed by NetWorker commands, called ASMs (application specific modules), to determine how to save that directory. If there are no directives, everything in the directory gets saved.

Q: What is the correct syntax for a `.nsr` directive file?

A: There are three forms, as described on the `nsr(4)` reference page. The first is a particular ASM specification. The second is a “save environment” directive. The third is a `<<dir>>` directive. A `#` at the beginning of a line indicates a comment.

```
[+] ASM [args ...] : pattern ...  
save environment  
<< dir >>
```

In the first line the colon (`:`) is an internal field separator between the ASM specification and its arguments. As usual, patterns requiring special characters must be enclosed in double quotes (`"`). The following patterns are illegal: `".."` and `"/"`. This is because all patterns must match the current directory. The optional `[+]` key means the directive file is propagated to subdirectories. The save environment and `<<dir>>` directives are both adequately explained by the `nsr(4)` reference page.

Q: What ASMs (application specific modules) are available with NetWorker?

A: The principal NetWorker ASM is **uasm**, the module for saving and recovering UNIX filesystem data. You can use **uasm** manually during save and recover, for instance where you extract a particular file from piped output. For example, this command extracts a save set for client `mars` and relocates it from root to the `/a` filesystem:

```
# scanner -c mars /dev/rmt/tps3d4nrnsv | uasrm -rv -m /=/a
```

See the `uasrm(1M)` reference page for more information about **uasm**, and for descriptions of the following ASMs: `skip`, `null`, `nullasm`, `holey`, `always`, `logasm`, `mailasm`, `atimeasm`, `tarasm`, `swapasm`, `xlateasm`, `compressasm`, `nsrmmdbasm`, `nsrindexasm`.

All these have the same basic options, so please refer to the reference pages when determining which ones to use for complicated configurations.

Compression

Q: Where is the compression being done on my backups?

A: There are only two ways of compressing backups: in the software using compression directives (**compressasm**) defined in a *.nsr* file, or in the hardware using the tape drive's firmware implementation of compression.

Q: How do I enable hardware (H/W) compression?

A: For hardware compression you must specifically use the "c" (compression) device. For instance, specify */dev/rmt/tpsXdYnrnsvc* instead of */dev/rmt/tpsXdYnrnsv*.

There are some NetWorker configuration changes to support hardware compression. You must inform NetWorker that tape data capacity has been enlarged. In the **nwadmin** window, choose Devices from the Media menu. Then in the Devices window, choose Details from the View menu.

For a DLT device, select the hardware compression drive in the Devices field, and enter "20GB" in the Volume default capacity field. Note that NetWorker does not use this value until it labels a tape, so it does not take effect immediately.

Some words about compression: software compression is done on the client during backup, whereas hardware compression is done on the server. This has implications for network throughput, drive throughput, CPU usage, and so forth.

Q: What is the compression rate?

A: The Legato **compressasm** uses their implementation of Lempel-Ziv coding to compress file pages (one page is always 8 KB) while saving them. This can result in less data being transmitted over the network and less data being stored. Since network and media subsystems are normally major performance bottlenecks, compression provides significant performance improvement, especially if the client processor is relatively fast. Compression of 50-60% is typical, but results vary with file content. Already-compressed files, such as GIF or JPEG images, do not benefit further from either software or hardware compression.

Logging and Notification

Q: How can I see a log of NetWorker activities?

A: People often remark that information in the Messages panel of the **nwadmin** window does not go back far enough. For NetWorker server messages going back to when you started the networker daemons, see the standard log file */nsr/logs/daemon.log*.

If you want to save daemon log files from previous NetWorker starts, you can change the */etc/init.d/networker* file to save the old log file before starting **nsrd**.

Q: I no longer receive mail notification for savegroup completion. Why?

A: Verify that mail for root is working. Mail failure is the likely cause of this problem.

Glossary

1-9

Intermediate backup levels. Each number represents a backup level. Lower levels back up more files.

4mm, 8mm

Choice in the *Devices* window; represent four- and eight-millimeter magnetic tapes.

archive

The process by which NetWorker backs up directories or files to an archive volume and then optionally deletes them to release disk space.

ASM

Application specific module. 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.

attribute

A piece of information describing a NetWorker resource, with a name and a list of values.

autochanger

A device that can move media among various components located in the device, including slots and media drives. Autochangers automate the loading, labeling, and mounting of media, during both backups and recovers.

browse policy

The browse policy determines how long entries for your files will remain in the online file index.

client

A machine that accesses the NetWorker server to back up or recover files. Clients may be workstations, PCs, or file servers with gigabytes of data.

cloning

The process by which NetWorker makes an exact copy of saved data (save sets). NetWorker can clone individual save sets or entire backup volumes.

command line

The shell prompt, where you enter commands.

compressasm

The NetWorker directive used for compressing and decompressing files.

daemon

A long-lived program that implements a service. For example, *nsrd* is a daemon that implements the NetWorker backup and recover service.

device

The backup device connected to the NetWorker server; used for backing up and recovering client files.

directive

Instruction to maximize the efficiency of a backup and to handle special files.

file index

A database of information maintained by NetWorker which tracks every file or filesystem backed up.

fileserver

A machine with disks that provides services to other machines on the network.

filesystem

1. A subtree of a UNIX file tree that is on a specific disk partition or other mount point.
2. The entire set of all UNIX files.
3. A method of storing files.

full (f)

A backup level in which all files are backed up, regardless of when they last changed.

group

A client or group of clients that starts backing up files at a designated time.

head (h)

Represents the beginning of a save set that spans two backup volumes.

heterogeneous

Heterogeneous networks are networks with systems of different platforms that interact meaningfully across the network.

hint

Choice of in the Devices window; represents half-inch magnetic tape.

holey

A directive used to back up files that do not have all of their data blocks allocated.

HSM

Hierarchical storage management. Provides a way to automatically move data between a local disk and other storage locations to conserve network storage resources.

incremental (i)

A backup level in which only files that have changed since the last backup are backed up.

interoperability

The ability of software and hardware on multiple machines from multiple vendors to communicate meaningfully.

job pack

A tray or carousel that holds several backup volumes.

jukebox

See autochanger.

level [1-9]

A backup level that backs up files that have changed since the last lower-level backup.

machine

Any computer, including file or compute servers, disk-based workstations, or diskless workstations.

mailasm

A directive that adheres to spool mail file-locking conventions and resets a file's access time back to its pre-saved values, so users can tell if new mail arrived after NetWorker backed up their mail.

media

Magnetic tape or optical disks used to back up files.

media index

A database of information maintained by NetWorker to track every backup volume.

media manager

The NetWorker component that tracks save sets to backup volumes.

media pool

The collection of backup volumes recognized and managed by NetWorker.

migrate

The process of moving files meeting administrator-defined criteria from local disk to slower, less-expensive storage devices.

NetWorker

The network-based software product used to back up and recover filesystems.

NetWorker client

A machine that can access the backup and recover services from a NetWorker server.

NetWorker daemons

Daemons specific to the NetWorker environment.

NetWorker server

The machine on a network running the NetWorker software, containing the online indexes, and providing the backup and recover services to the clients on a network.

NetWorker resources

Components of NetWorker software configuration information, described by a list of attributes and values.

namespace

An application associated with an entry in the online file index.

notice

A response to a NetWorker event.

nsrhost

The logical *hostname* of the machine that is the NetWorker server.

online indexes

The databases located on the server that contain all the information pertaining to the client backups and backup volumes.

operator

The person who monitors the server status, loads backup volumes into the server devices, and otherwise executes day-to-day tasks using NetWorker.

override

A backup level that takes place instead of the scheduled one.

pools

See *volume pools*.

preconfigured

Existing selections or configurations for different NetWorker features.

pre-migration

Command used for making a copy of a file in preparation for migration.

preview

A look at what a NetWorker command will do without actually executing the command.

print

Send data to a printer.

qic

A choice in the Devices window; represents quarter-inch cartridge tape.

recall

The process of moving a migrated file from a storage location back to the original location on local disk.

recover

The NetWorker command used to browse the server index and recover files from a backup volume to a client's disk.

retention policy

The retention policy determines how long entries will be retained in the media index and thus be recoverable.

recycle

A volume whose data has passed both its browse and retention policies and is available for relabeling.

resources

See NetWorker resources.

save

The NetWorker command that backs up client files to backup volumes and makes data entries in the server index.

save set

A set of files or a filesystem backed up onto backup media using NetWorker.

save set ID

An internal identification number assigned to a save set by NetWorker.

save set recover

A process by which NetWorker recovers an entire save set instead of individual files in the save set.

scanner

The NetWorker command used to read a backup volume.

server

The machine on a network running the NetWorker software, containing the online index,

and providing backup and recover services to the clients on a network.

shell prompt

The command line to which you enter UNIX commands.

skip

The directive to skip files during a backup. Useful for avoiding the copying of files that do not require backup.

skip (s)

A backup level in which files are skipped and not backed up.

SNMP

Simple network management protocol.

stub

A UNIX symbolic link that serves as a place holder for a migrated file and as a pointer to the new location of a migrated file.

super-full

A clone of the most recent full backup of a save set and of all migration save sets for the designated client filesystem.

superuser

A UNIX user with *root* privileges.

system administrator

The person normally responsible for installing, configuring, and maintaining NetWorker.

tail (t)

Represents the end of a save set that spans two backup volumes.

user

Any person who employs NetWorker from his or her workstation to back up and recover files.

volume

Backup media, such as magnetic tape or optical disk.

volume ID

The internal identification assigned to a backup volume by NetWorker.

volume name

The name you assign to a backup volume when it is labeled.

volume pools

A feature that allows you to sort your backup data to selected backup volumes. A volume pool contains a collection of backup volumes that have specific data sorted during the backup process.

Index

Numbers

1-9 backup levels, 232

A

aborted save set status, 103, 120

abort status, 72

access

remote, 27

restricting clients, 31, 35

access control lists, 111

Action field, 165

Active devices field, 49, 51

adding

administrators, 52

backup devices, 96

clients, 30

directives, 149

index entries, 39

jukebox devices, 171

Administrator Guide

conventions, xxiv

Administrator program, NetWorker, 2

administrators, adding or changing, 52

Administrator scrolling list, 49

Aliases field, 27, 31

alleviating network traffic, 159

Always prompt toggle button, 107

Annotation field, 194

appen mode, 70

Application Specific Modules, 148

archive

contrasted with backup, 204

defined, 191, 204

grooming, 202

verification, 202

Archive Client field, 194

Archive clone pool list, 194

Archive completion command field, 194

Archive list, 195

Archive pool list, 194

archive request, making, 197

Archive Request Control window, 195

Archive Request Details window, 196

Archive Request Schedule window, 199

Archive Requests scrolling list, 194

Archive Requests window, 193

Archive services field, 25

Archive Start time setting, 194

Archive Status choices, 194

Archive users field, 27, 192

Archive Verify choices, 194

autochanger

automatic device cleaning, 170

cleaning slots, 170

cleaning the devices, 174

exclusive to NetWorker, 175

managing media automatically, 175

messages, 188

- modifying, 167
- notifications, 187
- requirements, 167
- unlabeled volumes, 170
 - See also* jukebox.
- Auto clean choices, 170
- Auto command, 76
- automatic
 - cloning, 129
 - cloning of save sets, 122
 - index management, 146
 - labeling, 64
 - managing of volumes, 67
 - media management, 175
 - mounting, 64
- Auto media management
 - choices, 170
 - enabling, 98
- auto media management, 99, 175
- Auto media verify choices, 89
- Autostart choices, 160
- Available slots field, 170

B

- backup
 - client, preventing manual, 49
 - concurrently to devices, 51
 - customizing, 131
 - failed, 57
 - files changed during, 266
 - Last Run, 55
 - monitoring and controlling, 54
 - Next Run, 55
 - preview, 261
 - previewing, 57
 - simulated, 57
 - speeding up, 51
 - stopping, 262
- Backup command field, 27
- backup contrasted with archive, 204
- backup device speed, 228
- backup group
 - changing time, 163
 - configuring, 159
 - creating, 161
 - defined*, 159
 - deleting, 163
 - details, 56
 - restarting, 59
 - selecting for client, 162
 - status, 55
 - stopping, 58
- backup level
 - 1-9, 232
 - defined*, 232
 - full, 232
 - full versus incremental, 233
 - incremental, 232
 - setting up, 232
- backup levels
 - defined*, 141
 - overriding, 133
 - selecting, 141
 - trade offs, 141
 - using with directives, 142
- backup media capacity, 263
- Backup program, NetWorker, 5
- backup schedule
 - convenience or security, 239
 - Default, 247
 - defined*, 231
 - Full Every Friday, 248
 - Full on 1st Friday of Month, 248
 - Full on 1st of Month, 248
 - preconfigured, 246
 - Quarterly, 248
 - staggering, 238
 - time considerations, 236

- typical, 235
- backup schedules, 139
 - calendars, 133
 - changing, 131
 - creating, 134
 - defined*, 131, 140
 - deleting, 131
 - monthly example, 136
 - overriding, 139
 - preconfigured, 131, 132
 - setting up, 133, 140
 - skipping, 139
 - weekly example, 134
- backup status, 55
- Backup Status window, 266
- backup volume
 - changing locations, 74
 - changing mode, 74
 - cloning, 78, 130
 - cloning example, 82
 - defined*, xxv
 - destination, 113
 - displaying information, 68
 - finding the name, 67
 - internal labels, 65
 - labeling, 61
 - labeling in a jukebox, 178
 - labeling with bar codes, 180
 - label templates, 64
 - location, 73
 - managing automatically, 67
 - managing manually, 72
 - mode, 70, 81
 - mounting and unmounting, 65
 - needed for recovery, 67
 - overriding retention policy, 76
 - premounting, 66
 - recycling manually, 77, 178
 - recycling versus adding, 188
 - removing, 77
 - requested for recovery, 67
 - sharing, 100
 - source, 113
 - unlabeled, 100
- backup volume, full, 263
- bar code labels
 - defined*, 180, 182
 - matching volume label, 181
 - using identical labels, 183
 - using on existing volumes, 183
- Bar code reader choices, 169
- bootstrap, 206
 - cannot print information, 266
 - event notification, 164
 - save set ID, 219
- browsable save set status, 103, 120
- browse policy, 146, 239, 249
- Browse policy field, 25
- browsing index entries, 39
- brows status, 71
- button
 - Clone, 81, 116
 - Create Multiple, 20
 - Current Month, 133
 - Details, 55, 80, 103, 116, 195, 196
 - generic, 20
 - Instances, 45
 - More, 116
 - Next Month, 133
 - Options, 106
 - Preview, 55
 - Previous Month, 133
 - Print, 21
 - Query, 116
 - Reclaim Space, 42
 - Recover, 103
 - Remove oldest cycle, 42
 - Restart, 55, 59
 - Save, 21

- Schedule, 195
- Speedbar, 3
- Start, 55, 106
- Stop, 55, 58
- Update Server List, 54
- Volumes, 106

C

- calendars, backup schedules, 133
- calibrating devices, 190
- Change command, 53
- Change Mode command, 74
- Change Server dialog box, 53
- Change Server speedbar button, 53
- Change Status command, 73
- changing
 - administrators, 52
 - backup group time, 163
 - backup volume locations, 74
 - directives, 149
 - index policies, 146
 - NetWorker server, 53
 - status of save sets, 73
 - volume modes, 75
- check outcome of archive request, 198
- choices
 - Auto clean, 170
 - Auto media management, 170
 - Auto media verify, 89
 - Autostart, 160
 - Bar code reader, 169
 - Clones, 161
 - Devices, 89
 - Disabled, 160
 - Discard recovered file, 107
 - Duplicate file resolution, 107
 - Enabled, 160
 - Group, 26, 88
 - Label template, 88
 - Manual saves, 49
 - Match bar code labels, 169
 - Maximum, 120
 - Model, 169
 - Overwrite existing file, 107
 - Period, 132, 145
 - Pools, 63
 - Pool type, 88
 - Rename recovered file, 107
 - Separator, 153
 - Status, 120
 - Store file index entries, 89
- choosing server configurations, 223
- cleaning cartridge notification, 164
- cleaning devices
 - cleaning cartridge slots, 170
 - for autochangers, 170, 174
 - notifications, 174
- Cleaning slot field, 170
- client
 - CPU speed, 231
 - data generation rate, 230
 - data transfer rate, 224
 - index, 260
 - performance, 230
 - recover access, 261
 - renaming, 260
- Client field, 102, 116
- Client retries
 - field, 161
 - sliding bar, 162
- Clients
 - pulldown menu, 9
 - scrolling list, 24
 - window, 23
- clients
 - access, restricting, 35

- adding, 30
 - applying schedules, 143
 - automatically clone multiple example, 123
 - cloning automatically example, 122
 - configuring, 29
 - creating an archive, 197
 - defined*, 29
 - failed backup, retrying, 162
 - manually cloning save sets example, 121
 - permissions, 35
 - preventing manual backups, 49
 - removing, 34
 - scheduling large filesystems, 32
 - security, 35
 - selecting backup group, 162
- Client Setup command, 23
- Clients field, 89
- Clients scrolling list, 41
- Clients window, 192
- Clone button, 81, 116
- Clone choices, 194
- Clone command, 114
- cloned data
- changing volume modes, 76
 - index entries, 84, 129
 - recovering, 85
 - removing entries from index, 77
 - retention policies, 129
 - storing off-site, 128
 - tracking save sets, 113
- Clone Pool field, 82, 118, 161
- clone pools
- preconfigured, 130
 - selecting, 118
 - using, 130
- Clones choices, 161
- Clone Volume command, 79
- Clone Volume Details window, 80
- Clone Volume Status window, 81
- Clone Volume window, 79
- cloning an archive volume, 200
- cloning data
- automatically, 129
 - automatically clone multiple clients example, 123
 - benefits, 83
 - defined*, 78, 113
 - destination volume, 113, 128
 - example of automatically, 122
 - example of manually, 121
 - integrity of data, 129
 - minimizing network traffic, 80, 129
 - requirements, 114
 - save sets with schedules, 130
 - selecting save sets, 115
 - selecting the volume, 82
 - source volume, 113, 128
 - time needed for, 129
 - viewing details, 117
 - volume capacity, 128
 - volume example, 82
- command
- Auto, 76
 - Change, 53
 - Change Mode, 74
 - Change Status, 73
 - Client Setup, 23
 - Clone, 114
 - Clone Volume, 79
 - Devices, 96
 - df**, 208
 - Directives, 148
 - dkinfo**, 208
 - Group Control, 54
 - Groups, 159
 - Indexes, 40
 - Inventory, 184
 - jb_config**, 167, 272
 - jbexercise**, 272
 - Jukeboxes, 168

- Label, 62
- Label Templates, 152
- Manual, 76
- mminfo**, 39, 271, 272
- mmlocate**, 272
- mmpool**, 272
- mmrecov**, 215, 272
- Mount, 65
- mv**, 261
- networker stop**, 272
- Normal, 73
- Notifications, 164
- nsr_crash**, 273
- nsradmin**, 271, 273
- nsrck**, 273
- nsrck -F**, 271
- nsrclone**, 273
- nsrclone -c -N**, 271
- nsrd**, 273
- nsrexecd**, 36, 269, 271, 273
- nsrhmck**, 271, 273
- nsrim**, 273
- nsrindexd**, 273
- nsrinfo**, 39, 271, 273
- nsrjb**, 264, 271, 273
- nsrls**, 271, 273
- nsrmig**, 271, 273
- nsrmm**, 271, 272
- nsrmm -d**, 271
- nsrmm d**, 273
- nsrmm dbd**, 273
- nsrmm -d -P**, 271
- nsrmm -l**, 272
- nsrmm -m**, 272
- nsrmm -m -l -R**, 272
- nsrmm -u**, 272
- nsrpmig**, 272, 273
- nsrretrieve**, 273
- nsrwatch**, 272, 274
- nwadmin**, 2, 273
- nwbackup**, 5
- nwrecover**, 6
 - On Help, 14
 - On Version, 14
 - On Window, 13
 - Policies, 143
 - Polling Interval, 12
 - Pools, 86
 - Print, 7
 - Recover, 102
 - recover**, 35, 274
 - Remove, 78
 - Save, 8
 - save**, 35, 274
 - savefs**, 35, 274
 - savegrp**, 36, 207, 272, 274
 - savegrp -p**, 261, 272
 - scanner**, 72, 78, 83, 219, 274
 - Schedules, 132
 - Search, 19
 - Server Setup, 48
 - Set Location, 73
 - Suspect, 73
 - Tabular, 8
 - tapeexercise**, 274
 - Topics, 14
 - Unmount, 65
 - Volume, 68
- Command field, 8, 21
- Complete Save Sets scrolling list, 56
- components of label templates, 153, 159
- compressasm directive, 149
- compressing data, 149
- compression, 237
- concurrent devices
 - controlling backup traffic, 51
 - defined*, 51
 - distributing backups, 52
- configuration examples, 227
- configuring

- backup groups, 159
 - clients, 29
 - jukeboxes, 169
 - confirmation box
 - Reclaim Space, 44
 - controlling backups, 54
 - Control port field, 169
 - conventions for *Administrator's Guide*, xxiv
 - Copies field, 8, 21
 - copies of data, 128
 - copy violation, 266
 - CPU speed, 230, 231
 - Create Multiple button, 20
 - creating
 - backup groups, 161
 - backup schedules, 134
 - clients, 30
 - copies of files, 113
 - directives, 148
 - index policies, 142, 145
 - label templates, 153
 - multiple configurations, 20
 - notifications, 165
 - volume pools, 85, 94
 - Current Month button, 133
 - customer support, xxv
 - Customize pulldown menu, 11, 132
 - customizing backups, 131
- D**
- daemons, NetWorker, 259
 - data
 - cloning, 128
 - compressing, 149
 - integrity, 129
 - recovering large amounts, 112
 - recovering unbrowsable, 112
 - reducing for backup, 149
 - retrieval, 203
 - scanning, 112
 - selecting for pools, 89
 - sorting with pools, 95
 - verification, 202
 - databases, NetWorker, 142
 - data transfer rate
 - clients, 224
 - jukeboxes, 225
 - network, 224
 - NetWorker server, 224
 - Date field, 116
 - Date last cleaned field, 98
 - Decade policy, 251
 - Default cleanings field, 170
 - Default schedule, 247
 - definition of archiving, 204
 - deleting
 - administrators, 53
 - backup devices, 96
 - backup groups, 163
 - clients, 34
 - devices, 99
 - index policies, 145
 - save sets, 33
 - volume pools, 88, 95
 - destination volume, 113
 - details
 - backup group, 56
 - displaying window, 17
 - Details button, 17, 55, 80, 103, 116, 195, 196
 - device cleaning notification, 164
 - Devices
 - choices, 89
 - command, 96
 - display, 4
 - field, 62, 169

- scrolling list, 97
- window, 96
- devices
 - adding and deleting, 99
 - calibrating, 190
 - cleaning for autochangers, 174
 - handling unlabeled volumes, 98
 - pathname for device, 99
 - reserved for recovery, 49
 - supported by NetWorker, 96
 - using, 96
 - using concurrently, 51
 - using multiple, 96
- df** command, 207, 208
- dialog box
 - Instances, 45
 - Label, 62
 - Print Options, 8, 21
 - Resource, 19
 - Save Options, 8, 21
 - Save Set Recover, 109
 - Save Set Recover Details, 103, 108
 - Save Set Recover Options, 106, 107
 - Save Set Recover Status, 108
 - Server, 53
 - Set Location, 73
- directive
 - nsrindexasm, 273
 - nsrmmdbasm, 273
 - uasm, 274
- Directive field, 25, 149
- Directive list, 194
- Directives
 - scrolling list, 149
 - window, 148
- directives
 - changing, 149
 - defined*, 148
- Directives command, 148
- directories
 - recovering complete, 38
 - recovering save sets, 105
- Disable, 195
- Disabled choice, 160
- disabling
 - a network backup, 160
 - automatic device cleaning, 170
 - element status, 186
- discarding recovered files, 107
- Discard recovered file choice, 107
- disk
 - reclaiming space, 41, 44
 - recovering, 112
 - recovering partition example, 109
 - reducing amount used, 112, 143
 - space required for index, 39
- disk, replacement, 213
- disk crash
 - preparing for, 205
 - recovering from, 205
- disk space
 - determining, 225
 - reclaiming, 241
- display
 - Devices, 4
 - Messages, 4
 - Pending, 4
 - Server status, 4
 - Sessions, 4
- displaying
 - backup volume information, 68
 - tabular view, 14
 - window details, 17
- displaying NetWorker, problems, 259
- DISPLAY variable, setting, 259
- Duplicate file resolution choices, 107
- duplicating volume pools, 94

E

Enabled choice, 160
 enabling
 scheduled backup, 160
 event notifications, 164
 examples
 automatic cloning, 122
 backup schedules, 134, 136
 clone multiple clients automatically, 123
 creating a volume pool, 90
 manually cloning save sets, 121
 sequential label template, 156
 storage oriented label template, 154
 two-sided label template, 157
 excluding data from recovery, 105
 exclusive NetWorker autochangers, 175
 expired cleaning cartridges, 170
 external labels, 151

F

failed backup, 57
 Failed Save Sets scrolling list, 56
 field
 Action, 165
 Active devices, 49, 51
 Aliases, 27, 31
 Annotation, 194
 Archive services, 25
 Archive users, 27
 Available slots, 170
 Backup command, 27
 Browse policy, 25
 Cleaning slot, 170
 Client, 102, 116
 Client retries, 161
 Clients, 89

Clone Pool, 82, 118, 161
 Command, 8, 21
 Control port, 169
 Copies, 8, 21
 Date, 116
 Date last cleaned, 98
 Default cleanings, 170
 Devices, 62, 169
 Directive, 25, 149
 Fields, 153
 Files to recover, 105
 First slot, 178
 Jukebox options, 186
 Last slot, 178
 Migration clone pool, 161
 Name, 25, 49, 169
 Next, 153
 Number of periods, 145
 Parallelism, 49
 Password, 27, 31
 Period, 144
 Pool, 119
 Relocate recovered data to, 107
 Remote access, 27, 31
 Remote user, 27, 31
 Retention policy, 25
 Save Set, 116
 Save set, 26, 30
 Schedule, 25
 Search for resource, 19
 Server, 25
 Sessions per device, 49
 Sort by, 120
 Starting with, 178
 Start time, 161
 Version, 49
 Volume, 116
 Volume name, 63
 Fields field, 153
 file
 changed during backup, 266

file index
 defined, 37, 142
 index actions, 39
 reclaiming disk space, 44
 removing entries, 146
 removing oldest entries, 42
 report, generating, 39
 viewing information, 40
filename, maximum length, 267
File pulldown menu, 7
files
 creating copies, 113
 open, 59
 recovering entire save sets, 105
 recovering individual example, 108
 recovering unbrowsable, 112
Files to recover field, 105
filesystems, scheduling, 32
finding a volume name, 67
Finished status, 55
First slot field, 178
full backup level, 141, 232
full backup volumes, 263

G

generic buttons, 20
grooming, *defined*, 202
Grooming choices, 194
Group choices, 26, 88
Group Control
 command, 54
 speedbar button, 54
 window, 55
Groups command, 159
Groups scrolling list, 160
Groups window, 159

H

Help pulldown menu, 13
Hierarchical Storage Management (HSM), 204

I

icon
 NW Admin, 2
 NW Backup, 5
 NW Recover, 6
immediately starting backups, 161
incremental backup level, 141, 232
index, client, 260
index actions, 39
index entries
 cloned data, 84
 cloned save sets, 129
 including, 89
 reclaiming filesystem space, 39
 recovering, 111
 removing, 39, 129, 146
indexes
 automatically managing, 146
 conserving space, 37
 file, 37, 142
 manually managing, 36
 media, 37, 142
 moving, 268
 reducing size, 41, 44
 size, 37
Indexes command, 40
Indexes speedbar button, 40
Indexes window, 40
index location, 206
index policies
 browse policy, 146
 changing, 146

- creating, 145
- defined*, 142
- preconfigured, 143
- retention policy, 146
- index size notification, 164, 256
- in-progress save set status, 103, 120
- inpro status, 72
- inserting index entries, 39
- InSight manuals, xxv
- Instances button, 45
- Instances dialog box, 45
- Instances scrolling list, 103
- integrity of cloned save sets, 129
- internal volume label, 65
- Inventory command, 184

J

- jb_config** command, 167, 272
- jbexercise** command, 272
- jukebox
 - accidentally relabeling media, 179
 - adding devices, 171
 - available slots, 170
 - calibrating devices, 190
 - capacity, 243
 - changing available slots, 173
 - choosing, 245
 - configurations, 169
 - daily operations, 176
 - data transfer rate, 225
 - defined*, xxv
 - determining capacity, 263
 - disaster recovery, 220
 - labeling volumes, 178
 - making changes, 168
 - mounting volumes manually, 176
 - pathnames, order of, 172

- policies, 243
- recycling volumes manually, 178
- using pools, 190
- See also* autochanger.

- Jukeboxes
 - command, 168
 - scrolling list, 169
 - window, 168
- Jukebox Mounting window, 176, 263
- Jukebox options field, 186

K

- killing the NetWorker daemons, 267

L

- Label
 - command, 62
 - dialog box, 62
 - speedbar button, 62
- labeling
 - accidentally relabeling, 178
 - backup volumes, 61
 - backup volumes in jukeboxes, 178
 - before backing up, 64
 - consistently with templates, 179
 - identical bar code labels, 183
 - matching bar code and volume names, 181
 - relabeling, 64
 - using unique names, 65
 - with bar codes, 180
 - with label templates, 64
- labels
 - external, 151
 - unrecognized, 100, 170, 175
- Label template choices, 88
- Label Templates

- command, 152
- scrolling list, 152
- window, 152

label templates

- choosing, 151
- components, 153, 159
- creating, 153
- defined*, 150
- matching to pool, 154
- naming, 150
- preconfigured, 150
- sequential example, 156
- storage-oriented example, 154
- two-sided example, 157
- using with pools, 95

label templates, preconfigured, 253

Last Run backup, 55

Last slot field, 178

location, backup volume, 73

log default notification, 256

M

managing

- backup volumes, manually, 72
- indexes manually, 36
- media automatically, 99, 175

man recyc

- mode, 70

Manual command, 76

manually

- cloning save sets, 121
- managing backup volumes, 72
- mounting jukebox volumes, 176
- recycling volumes, 77

Manual recycle toggle button, 63

Manual saves choices, 49

Match bar code labels choices, 169

matching label names, 179

Maximum level choices, 120

media

- accidentally relabeling, 179
- defined*, xxv
- index, 37
- recycling versus adding, 188
- relabeling, 64
- verification, 202

media index

- changing volume mode, 74
- defined*, 37, 142
- index actions, 39
- recycling volumes, 76
- removing volume entries, 77
- setting volume location, 73
- viewing, 68

Media pulldown menu, 10, 62

media request notification, 257

media set, *defined*, xxv

Media type scrolling list, 97

Menu bar, 3

messages, savegroup completion, 264, 267

Messages display, 4

Migration clone pool field, 161

migration notification, 164

Migration Setup, 36

mminfo command, 39, 271, 272

mmlocate command, 272

mmpool command, 272

mmrecov command, 215

mmrecov command, 272

mode

- appen, 70
- changing backup volume, 74
- defined*, 81
- man recyc, 70
- Read Only, 75

- recyc, 70
 - Model choices, 169
 - monitoring backups, 54
 - monthly backup schedule, 136
 - Month policy, 251
 - More button, 116
 - Mount after labeling toggle button, 63
 - Mount command, 65
 - mounting volumes, 65
 - Mount speedbar button, 65, 68
 - moving indexes, 268
 - multiple configurations, 20
 - multiple devices, using, 96
 - multiple disks, backing up, 231
- N**
- Name field, 25, 49, 169
 - naming, 139
 - backup schedules, 139
 - index policies, 145
 - label templates, 150
 - NetWare version of NetWorker, xxvi
 - network data transfer rate, 224
 - NetWorker
 - Administrator program, 2
 - Backup program, 5
 - browse policy, 239
 - daemons, 259, 266
 - database modules, xxvi
 - databases, 142
 - measuring performance, 228
 - NetWare version, xxvi
 - notifications, 164
 - other documentation, xxv
 - preconfigured policies, 249
 - problems starting, 259
 - Recover program, 6
 - retention policy, 242
 - setting up environment, 223
 - starting for the first time, 52
 - supported devices, 96
 - trouble displaying, 259
 - troubleshooting, 259
 - NetWorker Administrator window, 2
 - NetWorker programs, 1
 - NetWorker server
 - changing, 53
 - choosing configurations, 223
 - client of itself, 24
 - data transfer rate, 224
 - errors, 264
 - parallelism, 50
 - shutting down and restarting, 218
 - networker stop** command, 218, 272
 - network speed, 229
 - network traffic
 - alleviating, 159
 - cloning volumes, 129
 - Never Run status, 55
 - Next field, 153
 - Next Month button, 133
 - Next Run backup, 55
 - no-rewind device name, 99
 - Normal command, 73
 - Not Finished status, 55
 - notification
 - about NetWorker events, 164
 - autochanger, 187
 - backup media request, 257
 - cleaning devices, 174
 - creating, 165
 - index size, 256
 - log default, 256
 - preconfigured, 164
 - registration, 256

- savegroup completion, 257
- Notifications command, 164
- Notifications scrolling list, 165
- Notifications window, 164, 264
- nsr_crash** command, 273
- nsradmin** command, 271, 273
- nsrck** command, 273
- nsrck -F** command, 271
- nsrclone -c -N** command, 271
- nsrclone** command, 273
- nsrd** command, 273
- /nsr directory*, 206
- nsrexecd**, 262
- nsrexecd** command, 36, 269, 271, 273
- nsrshmck** command, 271, 273
- nsrim** command, 273
- nsrindexasm directive, 273
- nsrindexd** command, 273
- nsrinfo** command, 39, 271, 273
- nsrjb** command, 264, 271, 273
- nsrjls** command, 271, 273
- nsrmig** command, 271, 273
- nsrmm** command, 271, 272
- nsrmmdbasm directive, 273
- nsrmmdbd** command, 273
- nsrmm -d** command, 271
- nsrmm d** command, 273
- nsrmm -d -P** command, 271
- nsrmm -l** command, 272
- nsrmm -m** command, 272
- nsrmm -m -l -R** command, 272
- nsrmm -u** command, 272
- nsrpmig** command, 272, 273
- nsrretrieve** command, 273
- nsrwatch** command, 216, 272, 274

- Number of periods field, 145
- nwadmin** command, 2, 191, 273
- NW Admin icon, 2
- nwbackup** command, 5
- NW Backup icon, 5
- nwrecover** command, 6
- NW Recover icon, 6

O

- On Help command, 14
- online indexes, 37
- On Version command, 14
- On Window command, 13
- open files, handling, 59
- optical device *pathnames*, 99
- Options button, 106
- Options pulldown menu, 12
- overriding
 - backup level, 133
 - backup schedules, 139
 - volume retention policy, 76
- Overwrite existing file choice, 107
- overwriting recovered files, 107

P

- paragraph, xxiv
- parallelism, 237
 - clients, 50
 - turning off, 50
- Parallelism field, 49
- Password field, 27, 31
- pathnames*
 - for devices, 99
 - optical, 99

- order for jukeboxes, 172
- Pending display, 4
- Pending Save Sets scrolling list, 56
- performance
 - client, 230
 - measuring, 228
 - server, 228
- Period choices, 132, 145
- Period field, 144
- permissions
 - client, 35
 - root@server_name*, 52
 - user, 53
 - user group, 53
- Policies
 - command, 143
- policies, jukebox, 243
- policies, recovering files, 140
- Policies scrolling list, 144
- Policies window, 143
- policy
 - Decade, 251
 - Month, 251
 - preconfigured, 249
 - Quarter, 251
 - Week, 250
 - Year, 251
- Polling Interval command, 12
- pool creation, 200
- Pool field, 119
- pools, preconfigured, 251
- Pools choices, 63
- Pools command, 86
- Pools scrolling list, 88
- Pools window, 87
- Pool type choices, 88
- posting NetWorker notifications, 165
- post-processing, Alternate Backup Command, 27
- preconfigured
 - backup schedules, 131, 246
 - clone pools, 130
 - directives, 148, 255
 - index policies, 143, 249
 - label templates, 150, 253
 - notifications, 164, 255
 - pools, 251
 - schedules, 132
 - settings, 24
- premounting backup volumes, 66
- pre-processing, Alternate Backup Command, 27
- Preview button, 55
- previewing backups, 57
- Preview Run status, 55
- Previous Month button, 133
- Print
 - button, 21
 - command, 7
- printing window contents, 7, 21
- Print Options dialog box, 8, 21
- privileges, *root*, 35
- program
 - NetWorker Administrator, 2
 - NetWorker Backup, 5
 - NetWorker Recover, 6
- prvtoc** command, 207, 208
- pull-down menu
 - Clients, 9
 - Customize, 11, 132
 - File, 7, 10
 - Help, 13
 - Media, 10, 62
 - Options, 12
 - Save, 11
 - Save Set, 73, 114
 - Server, 48
 - View, 14

Q

- Quarterly backup schedule, 248
- Quarter policy, 251
- Query button, 116

R

- Read Only mode, 75
- reclaiming
 - disk space, 44
 - filesystem space, 39
- reclaiming disk space, 241
- Reclaim Space button, 42
- Reclaim Space confirmation box, 44
- Recover
 - command, 102
 - window, 37
- recover, 209
- recoverable save set status, 103, 120
- Recover access scrolling list, 31
- Recover button, 103
- recover** command, 35, 274
- recovering
 - access control lists, 112
 - a primary disk, 214
 - client access to, 261
 - directories, complete, 38
 - disk, 112
 - disk partition example, 109
 - file policies, 140
 - files, xxv
 - files, from index, 241
 - files from other clients, 27
 - from a disk crash, 205
 - from a secondary disk crash, 212
 - from reserved device, 49
 - indexes, 215

- individual files example, 108
- large amounts of data, 112
- /nsr* on a NetWorker server, 215
- save sets, 101
 - to a new server, 219
 - unbrowsable data, 112
 - volumes needed for, 67
 - with jukeboxes, 220
- Recover program, NetWorker, 6
- recovery and retrieval, 203
- recov status, 71
- recyc
 - mode, 70
 - status, 72
- recyclable save set status, 103, 120
- reducing
 - backup data, 149
 - index size, 39, 41, 44
 - used disk space, 112
- reference pages, xxv
- registration notification, 164, 256
- relabeling
 - media, accidentally, 179
 - volumes, 64
- relocate recovered data, 107
- Relocate recovered data to field, 107
- Remote access field, 27, 31
- Remote user field, 27, 31
- Remove command, 78
- Remove oldest cycle button, 42
- removing
 - backup volumes, 77
 - clients, 34
 - index entries, 39, 129, 146
 - oldest save sets, 42
- Rename recovered file choice, 107
- renaming a client, 260
- renaming recovered files, 107

- replacement disk, 213
 - requirements for cloning save sets, 114
 - rescheduling an archive request, 199
 - Resource dialog box, 19
 - Restart button, 55, 59
 - restarting a backup group, 59
 - restricting client access, 31, 35
 - retention policies for cloned volumes, 129
 - retention policy, 146, 242, 249
 - Retention policy field, 25
 - retrieval, *defined*, 203
 - retrying failed client backup, 162
 - root* privileges, 35
 - Running status, 55
- S**
- Save button, 21
 - Save command, 8
 - save** command, 35, 274
 - savefs** command, 35, 274
 - savegroup
 - completion messages, 264, 267
 - error messages, 265
 - savegroup completion notification, 164, 257
 - savegrp** command, 36, 207, 272, 274
 - savegrp -p** command, 261, 272
 - Save Options dialog box, 8, 21
 - save session, 51
 - Save Set Clone Details window, 117
 - Save Set Clone Status window, 118
 - Save Set Clone window, 114
 - Save Set field, 116
 - Save set field, 26, 30
 - Save Set pulldown menu, 11, 73, 114
 - save set recover, 209
 - feature, 209
 - using, 209
 - Save Set Recover Details dialog box, 103, 108
 - Save Set Recover dialog box, 109
 - Save Set Recover Options dialog box, 106, 107
 - Save Set Recover Status dialog box, 108
 - Save Set Recover window, 102
 - save sets, 113
 - automatically cloning, 122
 - changing status, 73
 - choosing for recovery, 103
 - cloned, index entries, 129
 - cloning, 113
 - cloning with schedules, 130
 - date created, 116
 - defined*, xxv, 37
 - details, viewing, 45
 - integrity of cloned, 129
 - manually cloning, 121
 - recovering, 101
 - recovering individual files example, 108
 - removing oldest, 42
 - selecting for browsing, 115
 - save sets, archive versus backup, 201
 - Save Set scrolling list, 102
 - Save sets field, 89
 - Save Sets scrolling list, 42, 71, 105, 116
 - Save sets scrolling list, 26, 118
 - save set status
 - aborted, 103, 120
 - browsable, 103, 120
 - in-progress, 103, 120
 - normal, 73
 - recoverable, 103, 120
 - recyclable, 103, 120
 - scanned-in, 103, 120
 - suspect, 103, 120
 - saving remote filesystems, 265

- saving window contents, 8, 21
- scanned-in save set status, 103, 120
- scanner** command, 78, 83, 216, 219, 274
- scanning data, 112
- scann status, 72
- schedule
 - Full Every Friday, 248
 - Full on 1st Friday of Month, 248
 - Full on 1st of Month, 248
 - Quarterly, 248
- Schedule button, 195
- scheduled backups, 131, 140
- Schedule field, 25
- Schedules command, 132
- Schedules window, 132
- scrolling list
 - Administrator, 49
 - Clients, 24, 41
 - Complete Save Sets, 56
 - Devices, 97
 - Directives, 149
 - Failed Save Sets, 56
 - Groups, 160
 - in Archive Request window, 194
 - in Clients window, 192
 - Instances, 103
 - Jukeboxes, 169
 - Label Templates, 152
 - Media type, 97
 - Notifications, 165
 - Pending Save Sets, 56
 - Policies, 144
 - Pools, 88
 - Recover access, 31
 - Save Set, 102
 - Save Sets, 42, 71, 105, 116, 118
 - Save sets, 26
 - Server, 54
 - Status, 82, 106, 118
 - Volumes, 69, 80, 81
- Search command, 19
- Search for resource field, 19
- securing clients, 35
- selecting
 - backup schedules, 141
 - backup volume for cloning, 82
 - clone pool, 118
 - save sets for browsing, 115
 - specific data for pools, 89
- Separator choices, 153
- Server field, 25
- server performance, 228
- Server pulldown menu, 10, 48
- Servers command, 48
- Server scrolling list, 54
- Server status display, 4
- Server window, 48
- Sessions display, 4
- Sessions per device field, 49
- Set Location command, 73
- Set Location dialog box, 73
- setting up backup schedules, 133, 140
- setting up NetWorker environment, 223
- sharing backup volumes, 100
- simulated backups, 57
- size of index, 37
- skip backup level, 141
- skipping backup schedules, 139
- sliding bar, Client retries, 162
- slots
 - available, 170
 - changing for jukeboxes, 173
 - cleaning cartridges, 174
- Sort by field, 120
- sorting data to pools

- volume pools
 - sorting data, 86
 - source volume, 113
 - speed
 - backup device, 228
 - CPU, 230
 - network, 229
 - Speedbar
 - defined*, 3
 - turning on and off, 13
 - speedbar button
 - Change Server, 53
 - defined*, 3
 - Group Control, 54
 - Indexes, 40
 - Label, 62
 - Mount, 65, 68
 - Unmount, 65
 - Speedbar toggle button, 13
 - staggering backup levels, 238
 - Start button, 55, 106
 - starting
 - backup from Groups window, 58
 - Backup program, 5
 - backups immediately, 57
 - NetWorker for the first time, 52
 - Recover program, 6
 - starting NetWorker, problems, 259
 - starting NetWorker daemons, 259
 - starting time for archive, 198
 - Starting with field, 178
 - Start time field, 161
 - status
 - backup, 55
 - backup group, 55
 - save set, 71
 - Status choices, 120
 - Status scrolling list, 82, 106, 118
 - Stop, 195
 - Stop button, 55, 58, 262
 - stopping a backup, 262
 - stopping a backup group, 58
 - Store file index entries choices, 89
 - summary of archive request, 201
 - Suspect command, 73
 - suspect save set status, 103, 120
 - susp status, 72
- T**
- Tabular command, 8
 - Tabular toggle buttons, 15
 - tabular view, displaying, 14
 - tape drives, transfer rates, 223
 - tapeexercise** command, 274
 - tape mount request notification, 164
 - technical support, xxv
 - time needed for cloning data, 129
 - toggle buttons
 - Always prompt, 107
 - Details, 17
 - Manual recycle, 63
 - Mount after labeling, 63
 - Speedbar, 13
 - Tabular, 15
 - Tone, 13
 - Tone, turning on and off, 13
 - Tone toggle button, 13
 - Topics command, 14
 - tracking cloned save sets, 113
 - transfer rate, tape drives, 223
 - troubleshooting NetWorker, 259
 - typical backup schedule, 235

U

uasm directive, 274
unlabeled volumes, 100, 175
Unmount command, 65
unmounting volumes, 65
Unmount speedbar button, 25
Update Server List button, 54
using directives, 148

V

verification, *defined*, 202
verify existing archive, 200
Version field, 49
View pulldown menu, 14
Volume field, 116, 176
volume Location, 70
Volume name field, 63
volume pools
 changing, 93
 clone, 118, 130
 creating, 85, 94, 200
 defined, 85
 deleting, 88, 95
 duplicates, 94
 example, 90
 including index entries, 89
 selecting specific data, 89
 sorting data, 87, 95
 using label templates, 95
 using with jukeboxes, 190
volume pools, preconfigured, 251
Volumes
 button, 106
 scrolling list, 80, 81
volumes

changing modes, 74
cloned, 130
cloning capacity, 128
defined, xxv
destination, 113, 128
finding a name, 67
internal label, 65
labeling, 61
label templates, 64
location, 73
managing, 67
managing manually, 72
mounting, 61
mounting and unmounting, 65
needed for recovery, 67
relabeling, 64
removing, 77
source, 113, 128
Volumes command, 68
Volumes scrolling list, 69
Volumes window, 68

W

weekly backup schedule, 134
Week policy, 250
window
 Backup Status, 266
 Clients, 23
 Clone Volume, 79
 Clone Volume Details, 80
 Clone Volume Status, 81
 Devices, 96
 Directives, 148
 Group Control, 55
 Group Control Details, 56
 Groups, 159
 Indexes, 40
 Jukeboxes, 168

Jukebox Mounting, 176, 263
Label Templates, 152
NetWorker Administrator, 2
Notifications, 164, 264
Policies, 143
Pools, 87
printing contents, 7, 21
Recover, 37
Save Set Clone, 114
Save Set Clone Details, 117
Save Set Clone Status, 118
Save Set Recover, 102
saving contents, 8, 21
Schedules, 132
Server, 48
Tabular, 14
Volumes, 68

X

`xlv_mgr` command, 207

Y

Year policy, 251

Tell Us About This Manual

As a user of Silicon Graphics products, you can help us to better understand your needs and to improve the quality of our documentation.

Any information that you provide will be useful. Here is a list of suggested topics:

- General impression of the document
- Omission of material that you expected to find
- Technical errors
- Relevance of the material to the job you had to do
- Quality of the printing and binding

Please send the title and part number of the document with your comments. The part number for this document is 007-1458-050.

Thank you!

Three Ways to Reach Us

- To send your comments by **electronic mail**, use either of these addresses:
 - On the Internet: techpubs@sgi.com
 - For UUCP mail (through any backbone site): *[your_site]!sgi!techpubs*
- To **fax** your comments (or annotated copies of manual pages), use this fax number: 415-965-0964
- To send your comments by **traditional mail**, use this address:

Technical Publications
Silicon Graphics, Inc.
2011 North Shoreline Boulevard, M/S 535
Mountain View, California 94043-1389