

**IRIX™ NetWorker®
Administrator's Guide**

Document Number 007-1458-030

CONTRIBUTORS

Written by Carolyn Curtis
Edited by Christina Cary
Production by Lorrie Williams
Engineering contributions by Mark Epstein and Robert Clark

© Copyright 1995, Silicon Graphics, Inc.— All Rights Reserved

This document contains proprietary and confidential information of Silicon Graphics, Inc. The contents of this document may not be disclosed to third parties, 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 Unpublished rights reserved under the Copyright Laws of the United States.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 94039-7311.

Silicon Graphics and IRIS are registered trademarks, and IRIX is a trademark of Silicon Graphics, Inc. Legato is a trademark and Legato NetWorker is a registered trademark of Legato Systems, Inc. Network License System and NetLS are trademarks of Apollo Computer, Inc., a subsidiary of Hewlett-Packard Company. EXABYTE is a trademark of EXABYTE Corporation. NFS is a registered trademark of Sun Microsystems, Inc. X Window System is a trademark of the Massachusetts Institute of Technology. Sun is a trademark of Sun Microsystems, Inc. IBM is a registered trademark of International business Machines Corporation. MIPS is a registered trademark of MIPS Technologies, Inc.

Contents

About This Guide xxiii

Audience xxiii

Contents of This Guide xxiv

Conventions xxv

Other Documentation xxvi

Product Support xxvii

- 1. Introduction to NetWorker** 1
 - Versions and Compatibility 2
 - NetWorker Basics 2
 - The NetWorker Client/Server Model 3
 - Backing Up Clients 4
 - Recovering Lost Files 5
 - Backup Devices and Media 5

- 2. Installing NetWorker** 7
 - Installing and Enabling NetWorker Software on Servers 9
 - Installing Server Software 9
 - Enabling Server Software 12
 - Enabling Server Software Except Autochangers 13
 - Enabling Autochangers 15
 - Configuring Clients to Recognize the Server 17
 - Using NIS to Configure Clients 17
 - Configuring Clients Without NIS or DNS 18
 - Installing NetWorker Software on Clients 19
 - Removing NetWorker 21

- 3. Using NetWorker Windows and Menus 23**
 - Starting the NetWorker Programs 23
 - Starting the NetWorker Administrator Program 24
 - Starting the Backup Program 27
 - Starting the Recover Program 28
 - Using NetWorker Menus 28
 - The File Menu 29
 - Clients Menu 29
 - Server Menu 30
 - Media Menu 30
 - Save Set Menu 31
 - Customize Menu 31
 - Options Menu 32
 - Changing the Polling Interval 33
 - Enabling or Disabling the Speedbar 33
 - Enabling or Disabling Tone 34
 - Help Menu 34
 - View Menu 35
 - Displaying a Tabular View 35
 - Displaying Details 38
 - Searching for Resources 39
 - Printing Window Contents 40
 - Saving Window Contents 41
 - Using NetWorker Generic Buttons 42
 - Create Multiple Button 42
 - Print and Save Buttons 43

- 4. NetWorker Test Drive 45**
 - Changing the NetWorker Server 46
 - Labeling and Mounting the Backup Volumes 47
 - Trying a Backup and Recover 48

Viewing the Backup Schedule	49
Viewing the Clients Window	50
Enabling the Scheduled Backup	52
Starting a Backup Now	54
Using an Autochanger	57
Using Concurrent Devices	57
Displaying NetWorker Licenses	58
The Next Step	59

5. NetWorker Configuration and Policy Issues	61
Determining Backup Schedules	61
NetWorker Backup Levels	63
Full Backups Versus Incremental Backups	65
Level Backups	66
Combining Full, Incremental, and Level Backups	68
Backup Time Requirements	69
Using the Compression Directive compressasm	71
Staggering the Backup Schedules	71
Balancing Convenience and Security	72
Using Save Sets	73
Using NetWorker Preconfigured Backup Schedules	73
Default	74
Full Every Friday	75
Full on First Friday of Month	76
Full on First of Month	77
Quarterly	78

	Determining Browse and Retention Policies (Index Policies)	80
	How Browse Policies Work	82
	Reclaiming Disk Space	84
	Recovering Files Removed From the Index	84
	How Media Retention Policies Work	85
	NetWorker Preconfigured Browse and Media Retention Policies	86
	Week Policy	86
	Month Policy	86
	Quarter Policy	87
	Year Policy	87
	Decade Policy	87
	Determining Autochanger Policies	88
	Jukebox Capacity for One Backup Cycle	88
	Jukebox Capacity for More Than One Backup Cycle	89
	Choosing a Jukebox	90
6.	Configuring a NetWorker Server	93
	Using NetWorker Preconfigured Settings	93
	Configuring Backup Groups	95
	Creating a Backup Group	96
	Changing a Backup Time	100
	Deleting a Backup Group	100
	Setting Up the Backup Schedules	101
	Creating a New Backup Schedule	104
	Example 1: A Weekly Schedule	105
	Example 2: A Monthly Schedule	106
	Creating Browse and Retention Policies	109
	Using Directives	112

Configuring a Server's Clients	115
Adding a New Client	121
Selecting Client Files to Be Backed Up	122
Giving Other Systems Access to a Client's Files	123
Scheduling Large Client Filesystems	124
Removing a Client	126
Fine-Tuning Server Settings	127
Changing Parallelism	130
Setting Concurrent Device Use	131
Adding or Changing Administrators	133
Changing the NetWorker Server	134

7. Configuring and Managing Volumes and Backup Devices 137

Labeling Backup Volumes	137
Using the Label Dialog Box	139
Labeling and Mounting a Backup Volume	141
Mounting and Unmounting Backup Volumes	142
Finding a Backup Volume Name	143
Automatically Managing Backup Volumes	144
Viewing the Volumes Scrolling List	146
Viewing the Save Sets Scrolling List	147
Manually Managing Backup Volumes	148
Changing the Status of a Save Set	149
Setting the Location for Your Backup Volumes	150
Changing the Mode of a Backup Volume	151
Determining Backup Media Capacity	153
Recycling Volumes	154
Removing a Backup Volume	155
Cloning Backup Volumes	156
Recovering from Cloned Volumes	158
Using the Clone Windows	158
Example: Cloning Backup Volumes	163

Using Label Templates	164
Using the Preconfigured Label Templates	165
Choosing a Label Template	167
Creating a Label Template	169
Label Template Tips and Examples	171
Example 1: A Storage-Oriented Template	171
Example 2: A Sequential Label Template	174
Example 3: A Two-Sided Media Label Template	176
Using Volume Pools	178
Shortcut: Creating a Volume Pool	179
Using the Pools Window	180
Using Preconfigured NetWorker Volume Pools	185
Creating a New Pool	187
Volume Pool Examples	191
Example 1: Separate Departments	191
Example 2: Secure Device	191
Example 3: Targeted Backups	192
Example 4: Specific Filesystem	192
Adding and Deleting Backup Devices	193
Adding a Backup Device	196
Deleting a Backup Device	196
8. NetWorker Operations	197
Checking the NetWorker Server Status	197
Monitoring and Controlling Backups	199
Viewing Details of a Scheduled Backup	201
Previewing a Group Backup	202
Starting a Scheduled Backup Immediately	203
Stopping a Backup Group	204
Restarting a Backup Group	204
Handling Open Files	205
Monitoring Backup Devices	206

Monitoring and Managing Index Disk Space Usage	208
Viewing Index Information	208
Viewing Save Sets to Determine Resource Usage	210
Removing the Oldest Cycle	212
Reclaiming Index Space	213
Checking NetWorker Notices	215
Using the Notifications Window	216
Using NetWorker Preconfigured Notifications	218
Recovering Files Removed From the Index	219
9. Recovering and Cloning Save Sets	221
Shortcuts	221
Recovering Save Sets	221
Cloning Save Sets	222
Recovering Save Sets	223
Using the Save Set Recover Windows	225
Save Set Recover Details Window	228
Save Set Recover Status Window	229
Examples	232
Example 1: Recovering Individual Files	232
Example 2: Recovering a Disk Partition	234
Cloning Save Sets	236
Requirements for Cloning Save Sets	238
Using the Save Set Clone Windows	239
Using the Save Set Clone Details Window	242
Using the Save Set Clone Status Window	243
Viewing Additional Browsing Choices	244
Using Automatic Cloning	247
Examples	248
Example 1: Manually Cloning Save Sets	248
Example 2: Automatically Clone One Client	249
Example 3: Automatically Clone Several Clients	250

- 10. Configuring and Scheduling Archive Requests 257**
 - Backup, Archiving, and Hierarchical Storage Management 258
 - Features of the NetWorker Archive Product 259
 - Data Archiving 259
 - Data Verification 260
 - Data Retrieval 260
 - Configuring NetWorker for Archiving 261
 - Configuring Archive Services for a Client 263
 - Using Archive Save Sets 264
 - Creating New Volume Pools for Archiving 265
 - Scheduling and Monitoring Archive Requests 265
 - Scheduling an Archive Request 265
 - Starting or Disabling an Archive Request 268
 - Viewing the Progress of the Current Archive Request 269
 - Verifying Data on an Archive Volume 271
 - Cloning an Archive 271
 - Archive Example 272
 - Creating an Archive Client 272
 - Making an Archive Request 273
 - Checking the Archive Request 274
 - Rescheduling the Archive Request 275
 - Shortcut 275

- 11. Using NetWorker with Jukeboxes 277**
 - Testing the Jukebox Connection 278
 - Configuring a Jukebox 278
 - Using the Jukeboxes Window 279
 - Adding Devices to a Jukebox 282
 - Changing Available Slots 285

Labeling Volumes and Using Optional Bar Codes	286
Labeling Volumes in a Jukebox	286
Labeling With Bar Codes	290
Designing a Labeling Scheme	291
Labeling Volumes With Bar Code Labels	292
Using Bar Codes That Do Not Match Volume Labels	296
Viewing Jukebox Labeling Messages	296
Mounting Jukebox Backup Volumes Manually	297
Performing Routine Jukebox Operations	299
Viewing Jukebox Notices	299
Checking Jukebox Space	300
Recycling Versus Adding More Backup Volumes	301
Taking Inventory of the Jukebox Contents	302
Using Pools With a Jukebox	304
Cleaning the Devices	304
Calibrating the Devices	304

12. NetWorker Performance	305
Guidelines for Choosing a Configuration	305
Tape Drives	306
Clients	306
Network	307
Server	307
Jukeboxes	307
NetWorker Configuration: Example 1	309
NetWorker Configuration Example 2	310
Measuring Server Performance	311
Backup Device Speed	311
Network Speed	312
Server CPU Speed	312
Memory	313

Measuring Client Performance	313
Filesystem Traversing	313
Data Generation Rate	313
Data on Multiple Disks	314
Client CPU Speed	314
A. Troubleshooting	315
Maximum Filename Length	315
Checking the NetWorker Daemons	316
/nsr Not Configured	316
Solving Licensing Problems	317
Renaming a Client	317
Using Recover Access	318
Using Savegroup Completion Messages	319
Binding to Server Errors	319
Saving Remote Filesystems	320
File Changed During Backup	321
Cannot Print Bootstrap Information	322
Using Savegroup Completion Warning Messages	322
Errors in Executing External ASMs	323
Solving X Error Messages	323
Moving Indexes	324
Recovering Files From an Interrupted Backup	324
Determining the NetWorker Server	324
Using nsrexecd	325
B. Recovering From a Disk Crash	327
Shortcut	328
Preparing for a Crash	329
Filing the Bootstrap Information	330
Filing the Disk Information	331
Using Recover or Save Set Recover	333
Recovering a Secondary Disk	338

Recovering a Primary Disk on a Client	339
Recovering /nsr on a NetWorker Server	340
Replacing the /nsr/res Directory	343
Finding the Bootstrap Save Set ID	344
Recovering to a New Server	345
Disaster Recovery With Jukeboxes	346

C. Theory of Operations	349
NetWorker Media Format	351
The NetWorker System at Rest	351
Browsing and Modifying Resources	353
Backing Up Files	355
Recovering Files	357
Backing Up the Entire System	358
Automatic Network-Wide Backups	360
Event Processing	364
Media Management	365
Index Management	367
Disaster Recovery	368
NetWorker Security	370

D. Command Summary	373
NetWorker Reference Pages	376

Glossary	381
-----------------	-----

Index	389
--------------	-----

Figures

Figure 3-1	NetWorker Administrator Icon 24
Figure 3-2	NetWorker Administrator Window 25
Figure 3-3	NetWorker Backup Icon 27
Figure 3-4	NetWorker Recover Icon 28
Figure 3-5	File Pulldown Menu 29
Figure 3-6	File Menu in Other NetWorker Windows 29
Figure 3-7	Clients Menu 30
Figure 3-8	<i>Server Menu</i> 30
Figure 3-9	Media Menu 31
Figure 3-10	Save Set Menu 31
Figure 3-11	Customize Menu 32
Figure 3-12	Options Menu 32
Figure 3-13	Change Polling Dialog Box 33
Figure 3-14	Speedbar 33
Figure 3-15	Help Menu 34
Figure 3-16	View Menu 35
Figure 3-17	Clients Window 36
Figure 3-18	Clients Window: Tabular View 37
Figure 3-19	Groups Window: Displaying Details 38
Figure 3-20	Groups Window: Detail View 39
Figure 3-21	Search Dialog Box 40
Figure 3-22	Print Options Dialog Box 40
Figure 3-23	Save Options Dialog Box 41
Figure 3-24	Create Multiple Button 43
Figure 4-1	Change Server Window 46
Figure 4-2	Label Dialog Box 47
Figure 4-3	Schedules Window 49

Figure 4-4	Clients Window	51
Figure 4-5	Groups Window	53
Figure 4-6	Group Control Window	55
Figure 4-7	Group Control Details Window	56
Figure 4-8	Registration Window	58
Figure 5-1	Backup Levels	64
Figure 5-2	Full Backup Followed by Level 9 to Level 1	66
Figure 5-3	Full Backup Followed by Level 1 to Level 9	67
Figure 5-4	Full Backup Followed by Level 5 and Incremental Backups	68
Figure 5-5	Default Schedule	75
Figure 5-6	Full on First Friday of Month Schedule	76
Figure 5-7	Full on First of Month Schedule	78
Figure 5-8	Quarterly Schedule	79
Figure 5-9	Backups for a Browse Policy of One Week	83
Figure 6-1	Customize Pulldown Menu	96
Figure 6-2	Groups Window	97
Figure 6-3	Schedules Window	102
Figure 6-4	Selecting a Backup Level	105
Figure 6-5	Overrides Menu	108
Figure 6-6	Policies Window	110
Figure 6-7	Directives Window	114
Figure 6-8	Clients Menu	116
Figure 6-9	Clients Window	117
Figure 6-10	Clients Window, Lower Part	120
Figure 6-11	Example Pathname for Filesystem to Be Backed Up	123
Figure 6-12	Remote Access	123
Figure 6-13	Server Menu	127
Figure 6-14	Server Window	128
Figure 6-15	Parallelism, Active devices, and Sessions per device Fields	132
Figure 6-16	Changing Administrators	133
Figure 6-17	Opening the Change Server Dialog Box	134

Figure 6-18	Change Server Dialog Box	135
Figure 7-1	Opening the Label Dialog Box	139
Figure 7-2	Label Dialog Box	139
Figure 7-3	Mounting or Unmounting a Backup Volume	142
Figure 7-4	Opening the Volume Management Window	145
Figure 7-5	Volume Management Window	145
Figure 7-6	Changing the Status of a Save Set	149
Figure 7-7	Set Location Dialog Box	150
Figure 7-8	Changing Backup Mode	151
Figure 7-9	Changing Volume Mode: Confirmation Box	153
Figure 7-10	Recycling Volumes	154
Figure 7-11	Removing a Volume	155
Figure 7-12	Warning Message for Removing a Volume	156
Figure 7-13	Opening the Clone Volume Window	159
Figure 7-14	Clone Volume Window	160
Figure 7-15	Clone Volume Details Window	161
Figure 7-16	Clone Volume Status Window	162
Figure 7-17	Pools Window	168
Figure 7-18	Label Templates Window	169
Figure 7-19	Example Racks of Magnetic Tapes	172
Figure 7-20	Example Label Template	174
Figure 7-21	Example Two-Sided Media Label Template	177
Figure 7-22	Opening the Pools Window	180
Figure 7-23	Pools Window (Expanded View)	181
Figure 7-24	Example Label Templates Window	188
Figure 7-25	Example Pools Window	190
Figure 7-26	Opening the Devices Window	193
Figure 7-27	Devices Window	194
Figure 8-1	NetWorker Administrator Window	198
Figure 8-2	Opening the Group Control Window	199
Figure 8-3	Group Control Window	200
Figure 8-4	Group Control Details Window	201
Figure 8-5	Group Control Window	204

Figure 8-6	Media Menu 206
Figure 8-7	Jukebox Mounting Window 207
Figure 8-8	Opening the Indexes Window 208
Figure 8-9	Indexes Window 209
Figure 8-10	Instances Window 211
Figure 8-11	Removing a Cycle From an Index 213
Figure 8-12	Reclaim Space Confirmation Box 214
Figure 8-13	Reclaiming Space 214
Figure 8-14	Reclaiming Index Space 215
Figure 8-15	Opening the Notifications Window 216
Figure 8-16	Notifications Window 217
Figure 9-1	Opening the Save Set Window 225
Figure 9-2	Save Set Recover Window 226
Figure 9-3	Save Set Recover Details Window 228
Figure 9-4	Save Set Recover Status Window 229
Figure 9-5	Save Set Recover Volumes Required Window 231
Figure 9-6	Save Set Recover Options Window 231
Figure 9-7	Opening the Save Set Clone Window 239
Figure 9-8	Save Set Clone Window 240
Figure 9-9	Save Set Clone Details Window 242
Figure 9-10	Save Set Clone Status Window 244
Figure 9-11	Additional Fields in the Save Set Clone Window 245
Figure 9-12	Sort by Buttons 247
Figure 9-13	Automatic Cloning Scheme 252
Figure 9-14	Schedule for Acct Incr 253
Figure 9-15	Schedule for Acct Full+Clone 253
Figure 9-16	Acct Incr in the Groups Window 254
Figure 9-17	Acct Incr in the Groups Window 255
Figure 10-1	NetWorker Administrator Icon 261
Figure 10-2	NetWorker Administrator Window 262
Figure 10-3	Enabling Archive Services for a Client 263
Figure 10-4	Archive Request Window, Expanded View 266
Figure 10-5	Archive Request Control Window 268

Figure 10-6	Archive Request Details Window	270
Figure 10-7	Archive Request Schedule Dialog Box	275
Figure 11-1	Opening the Jukebox Window	279
Figure 11-2	Jukeboxes Window	280
Figure 11-3	Devices Display of Administrator Window	281
Figure 11-4	Devices Window	283
Figure 11-5	Entering the Jukebox Device Pathname	284
Figure 11-6	Listing Jukebox Device Pathnames	284
Figure 11-7	Changing Available Slots	285
Figure 11-8	Jukebox Labeling Dialog Box	287
Figure 11-9	Labeling a Single Volume	288
Figure 11-10	Labeling a Range of Volumes	288
Figure 11-11	Relabel Confirmation Dialog Box	289
Figure 11-12	Jukeboxes Window	292
Figure 11-13	Bar Code Label in Jukebox Labeling Dialog Box	293
Figure 11-14	Template Name in Jukebox Labeling Dialog Box	294
Figure 11-15	Volume Management Window with Complete Bar Code Selections	295
Figure 11-16	Jukebox Labeling Dialog Box: Status Message	297
Figure 11-17	Jukebox Mounting Window	298
Figure 11-18	Notifications Window With Three Tape Mount Request Notifications	300
Figure 11-19	Jukebox Inventory Dialog Box	303
Figure B-1	Backups Over Time	335
Figure B-2	Recover Save Set Window	336
Figure B-3	Damaged Secondary Disk	338
Figure B-4	Damaged Primary Disk	339
Figure C-1	Major Components of NetWorker	350
Figure C-2	NetWorker Monitors and Daemons	352
Figure C-3	NetWorker Resource Types	354
Figure C-4	Typical <i>save</i> Session	356
Figure C-5	Typical <i>recover</i> Session	358
Figure C-6	Saving Two Filesystems in Parallel	359

Figure C-7	Starting the Nightly Save	361
Figure C-8	Nightly Network Backup	363
Figure C-9	Notification Subsystem	365
Figure C-10	<i>nsrmm</i> d and Other Processes	366
Figure C-11	The File Index and Processes	367
Figure C-12	Index Information From Saving to Recovering	369

Tables

Table 2-1	Subsystems for NetWorker Servers	10
Table 2-2	Licensed NetWorker Products	12
Table 2-3	Subsystems for NetWorker Clients	20
Table 5-1	Choosing a Jukebox	90
Table 7-1	Changing Volume Mode	152
Table 7-2	Preconfigured Label Templates	166
Table 7-3	Label and Pool Names	171
Table 9-1	Acct Incr and Acct Full+Clone Schedules	251
Table 12-1	Tape Drives and Data Transfer Rates	306
Table 12-2	Client Data Transfer Rates	306
Table 12-3	Data Transfer Rates for Networks	307
Table 12-4	Jukebox Data Transfer Rates and Capacities	308
Table 12-5	Percent of Data Backed Up for Each Type of Backup	308
Table 12-6	Example of the Percent of Data Backed Up	308
Table C-1	Reference Pages for NetWorker at Rest	352
Table C-2	Reference Pages for Browsing and Modifying Resources	354
Table C-3	Reference Pages for Backing Up Files	356
Table C-4	Reference Pages for Backing Up the Whole System	359
Table C-5	Reference Pages for Automatic Nightly Backups	364
Table C-6	Reference Pages for Media Management	366
Table C-7	Reference Pages for Index Management	367
Table C-8	Reference Pages for Disaster Recovery	369
Table D-1	Commonly Used NetWorker Commands	373
Table D-2	NetWorker Maintenance Commands	374
Table D-3	Reference Pages for NetWorker Client Commands	376
Table D-4	Other NetWorker Client Reference Pages	377

Table D-5	Reference Pages for NetWorker Server Commands	378
------------------	---	-----

About This Guide

The *IRIXTMNetWorker[®] Administrator's Guide* documents the Silicon Graphics[®] implementation of Legato NetWorker from Legato Systems, Inc. IRIX NetWorker is adapted from the Legato 4.1 release of NetWorker. This guide was prepared in conjunction with Release 4.1.1 of NetWorker for IRIX Release 5.3 and 6.x.

The *IRIX NetWorker Administrator's Guide* contains information on how to install, configure, and manage NetWorker backup and recover software files. It also explains how to configure systems for NetWorker archiving and retrieval, if these optional features have been purchased. This manual provides enough information for the basic use and administration of NetWorker installed in a typical networked environment.

Audience

This manual is intended for system administrators who are responsible for installing software and maintaining the servers and clients on a network. It can also be used by operators who monitor the daily backups of the systems on a network and maintain the backup media and devices.

The *IRIX NetWorker Administrator's Guide* contains information on how to install NetWorker backup and recover software on Silicon Graphics servers and clients. If you have a heterogenous network and are backing up client systems of a platform different from the NetWorker server, you need to purchase a Legato ClientPak from Silicon Graphics, which includes documentation for other platforms.

Contents of This Guide

The *IRIX NetWorker Administrator's Guide* contains information on day-to-day NetWorker tasks such as managing the backup volumes and monitoring the backups, and the less frequent system administration tasks.

Chapter 1, "Introduction to NetWorker," introduces NetWorker features and terminology.

Chapter 2, "Installing NetWorker," explains the procedures for installing NetWorker software on NetWorker servers and clients and for activating software licensing on NetWorker servers (client software is unlicensed).

Chapter 3, "Using NetWorker Windows and Menus," introduces the main windows of the three NetWorker programs and the menus of the Administration program.

Chapter 4, "NetWorker Test Drive," is a tutorial designed to familiarize you with NetWorker features and tasks.

Chapter 5, "NetWorker Configuration and Policy Issues," explains backup levels, backup schedules, browse and retention policies, and jukebox policies.

Chapter 6, "Configuring a NetWorker Server," gives procedures for server configuration tasks such as using NetWorker preconfigured settings, configuring backup groups, creating backup schedules, creating browse and retention policies, using directives, configuring and adding clients, and fine-tuning server settings.

Chapter 7, "Configuring and Managing Volumes and Backup Devices," explains creating and using volume pools and other index and volume management tasks that are performed when needed.

Chapter 8, "NetWorker Operations," explains the day-to-day administration tasks of checking server status, monitoring and controlling backups, managing backup devices, monitoring and managing index space usage, and checking NetWorker notices.

Chapter 9, "Recovering and Cloning Save Sets," explains how to use the save set feature to retrieve data efficiently and how to clone save sets for

storing offsite, sending data to another location, or verifying backed-up data.

Chapter 10, “Configuring and Scheduling Archive Requests,” explains how to configure NetWorker for archiving, and how to schedule and monitor archive requests.

Chapter 11, “Using NetWorker with Jukeboxes,” explains how to configure and administer autochangers, including jukeboxes.

Chapter 12, “NetWorker Performance,” provides system, peripherals, and network performance information to help you choose a NetWorker server configuration, and explains how to measure server and client performance.

Appendix A, “Troubleshooting,” explains troubleshooting procedures and discusses error messages and their causes.

Appendix B, “Recovering From a Disk Crash,” explains how to save important information about your system so that it will be available to you in the event of a disk crash, and how to recover after a disk crash.

Appendix C, “Theory of Operations,” contains reference information about how NetWorker works.

Appendix D, “Command Summary,” describes the command line interface to NetWorker: commonly used commands, commands available on clients, and commands available on servers.

This guide concludes with a glossary of terms and an index.

Conventions

There are some conventions used in this manual to make information easy to access and understand. The conventions are as follows:

- Command names are given in *italic* typeface. For example:
To invoke NetWorker, enter the *networker* command at the prompt in a shell window.

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

```
NetWorker successfully installed on 'mars'.
```

- Examples of what you type are displayed in bold fixed-width typeface. For example:

```
mars% /usr/etc/networker &
```

- Names of buttons are displayed in *italics*. For example:

Press the *Cancel* button.

- Important pieces of information and cautionary notes written to prevent you from making a mistake are marked with the word “Note.” For example:

Note: If you do not highlight any save sets in the Save sets scrolling list, then NetWorker removes all the oldest cycles of all the save sets displayed in the scrolling list for the client highlighted in the Clients scrolling list.

- Tips or suggestions that you do not necessarily have to follow, but may give you hints as to how to set up NetWorker at your site are marked with the word “Tip.” For example:

Tip: You may want to store a copy of the weekly or monthly full backup offsite.

- Entries that you type and are substituted with a variable are displayed as italic. For example:

```
atlas% man name
```

Other Documentation

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.

Both the *IRIX NetWorker Administrator's Guide* and *IRIX NetWorker User's Guide* are available online, as InSight manuals. To view the manuals, use *insight(1)* and install the subsystems *networker4.books.NetWorker_AG* and *networker4.books.NetWorker_UG*.

Product Support

Silicon Graphics offers a comprehensive product support and maintenance program for IRIS products. For information about using support services for this product, refer to the *Release Notes* that accompany it.

Introduction to NetWorker

Everyone knows the importance of backing up computer systems to protect from the loss of valuable data. In a networked environment, where the number of files grows as systems are added to the network, the need to protect against loss of data becomes crucial.

IRIX NetWorker is a software product that reliably protects files against loss across an entire network of systems. NetWorker saves valuable administrator time by speeding and simplifying daily backup operations. An easy-to-use graphical user interface 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. Furthermore, as the network and number of files expand, NetWorker has the capacity and performance to handle the load.

Whether you need to recover an entire filesystem or a single file, NetWorker tracks the location of the files on the backup volumes and requests the ones you need by name. You do not have to think about which backups are on which backup volumes. All you need to do is mount the requested backup volume in a NetWorker server device to recover the desired files. Loading and mounting volumes is automatic if you are using an *autochanger*. An autochanger is a device that can move media among various components located in the device, including slots, media drives, media access ports, and transports; examples are jukeboxes and tape libraries.

This chapter introduces

- versions and compatibility
- NetWorker basics
- the NetWorker client/server model
- backing up clients

- recovering lost files
- backup devices and media

Versions and Compatibility

NetWorker 4.1.1 server software runs on IRIX 5.3 and 6.x. NetWorker requires the base release of the software and enough disk space for the installation and the index files. NetWorker uses about 1 MB of index space for each gigabyte of backup.

The NetWorker 4.1.1 client software runs on IRIX 4.x, 5.3, and 6.x; see the *IRIX NetWorker Release Notes* for more information. NetWorker 4.1.1 is fully compatible with all index and resource files from NetWorker version 4.0.2 and 4.0.2A.

Note: For information on licenses and NetWorker products, see “Enabling Server Software” in Chapter 2, “Installing NetWorker.”

Caution: Once you have upgraded from NetWorker version 1.x to 4.1.1, you cannot go back because of the changes made to the index files during the upgrade. Index files from 4.0.2 and 4.0.2A are compatible with 4.1.1.

NetWorker Basics

A set of files backed up by a NetWorker client is called a *save set*. *Backup volumes* are the media to which NetWorker backs up data and are recognized and managed by NetWorker. *Volume pools* are used by NetWorker to sort specific data to preselected backup volumes.

The backup volume names and client save sets are maintained in two indexes on the server: the *file index* and the *media index*. NetWorker uses these two indexes to track file backups to specific backup volumes. The *media manager* works with the two indexes to track where the backups are located on the backup volumes, making it easy to recover lost files.

When backing up files, NetWorker may request that you mount a writable backup volume in the server device. When recovering files, NetWorker

requests a specific backup volume by its name. All you need to do is mount the requested backup volume or volumes in the device(s) attached to the server. If you have a jukebox or other autochanger, the NetWorker TurboPak option, and a NetWorker Autochanger Software Module, NetWorker automatically loads the requested volume if it is stored in the autochanger.

The NetWorker Client/Server Model

A client/server architecture is the key to the ability to support network-wide, heterogeneous backup and recover. NetWorker *servers* provide a backup and recover service: they receive files from NetWorker *clients*, store the files on backup volumes, and retrieve the files on demand.

A NetWorker server is a system equipped with one or more backup devices. The server responsibilities are to

- write files from many client systems to media
- maintain an index of information about all the files backed up and where they are stored
- allow the clients to browse the file index
- fill requests from clients to back up or recover files
- initiate automatic network-wide backups according to a specific schedule
- give notices about NetWorker events

The NetWorker clients are all the other systems on the network that use the backup and recover service provided by the server. Clients can be desktop workstations, PCs with small disk drives, or large systems with gigabytes of data.

The clients on the network have software installed that

- gives the client access to the server for backing up and recovering files
- allows the client to browse the online index on the server to select files for recovery

A single NetWorker server can provide the backup and recover services to many client systems.

Backing Up Clients

NetWorker performs network-wide system backups. Such backups are more efficient and easier to maintain than local backup because of these NetWorker features:

- NetWorker has preconfigured backup schedule suitable for most small to medium-sized networks. You can easily edit this schedule using the NetWorker X™ Window System interface. You may create multilevel backups spanning any time period: weekly, monthly, bimonthly, or quarterly.
- A sophisticated media manager prompts you for all of its media requirements. When a volume is needed for backup or file recovery, NetWorker sends a message via its window-based status monitor or by electronic mail.
- NetWorker enables you to back up to premounted multiple devices or to a jukebox for unattended backups.
- The filesystems or individual files being backed up can span multiple backup volumes, or numerous filesystems and files can be saved to the same backup volume.
- The online backup feature allows you to perform backups while server and client systems remain in operation.
- The backups from many clients take place in parallel. This feature keeps a steady stream of data supplied to the backup device or devices so they can operate at full speed. The parallel backup option also ensures that no single client can monopolize the server and prevent other clients from accessing the server. With IRIX NetWorker and the TurboPak, NetWorker can back up to several backup devices at once, thus optimizing server performance.
- You can separate the media containing your full backups from media containing non-full backups.
- You can sort data during backup to preselected volume pools.

Recovering Lost Files

The NetWorker online index and media manager work together to expedite the recovery of lost files. The basic procedure for recovering files is as follows:

- Users browse the online index of previous file backups to identify the lost files and mark the files they want to recover, using the graphical user interface.
- The NetWorker media manager determines which backup volumes are needed to recover the files and requests only the ones it needs for the recovery.
- Once the volume is loaded in the server device, the files are recovered to the client.
- The optional NetWorker Support for Jukeboxes automates tape loading operation for fast recoveries.

Backup Devices and Media

NetWorker supports a number of backup devices. For the current list, see the *IRIX NetWorker Release Notes*.

Caution: A no-rewind tape backup device is required (except for optical 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 work with IRIX NetWorker, digital linear tape (DLT) drives in all peripherals must have Silicon Graphics personalities. If the DLT is not bought from Silicon Graphics, you must request the correct version of the firmware from the vendor.

Note: DLT2700 stackers that are not from Silicon Graphics do not work with IRIX NetWorker.

Installing NetWorker

This chapter describes the procedure for installing and licensing NetWorker 4.1.1 for IRIX on NetWorker servers and clients. The NetWorker software distribution CD-ROM contains all the software you need for Silicon Graphics servers and clients, including

- IRIX NetWorker, which supports 10 or fewer clients and one or two tape drives
- options for additional client connections in additive units of 5, 25, or 100
- the NetWorker TurboPak option, which supports archive services, save set cloning, concurrent device support, autochanger support, and parallel backup and recovery
- ClientPak Software option
- Autochanger Software Module options

Licenses are required to enable each of these items except the ClientPak Software option.

NetWorker server software is license-protected with NetLS™ node-locked licenses, so the NetWorker software distribution is shipped with a temporary license that enables you to use all NetWorker features to back up a maximum of 10 clients for a short period until you obtain permanent licenses from Silicon Graphics. Additional client connections are separate licenses; each server and each autochanger require a license.

This chapter explains how to install NetWorker software on Silicon Graphics servers and clients, and how to install and activate the licenses that enable the server software. The basic steps are as follows:

For the server:

1. Use *inst(1M)* to install NetWorker software on servers. This procedure is described in the section “Installing Server Software.”
2. Enable NetWorker server software by following the directions in the section “Enabling Server Software.”
3. Follow the directions included with the NetWorker package to obtain permanent licenses from Silicon Graphics, if you did not install permanent licenses earlier.
4. Install permanent licenses to replace temporary licenses on the appropriate servers by again using the procedure in the section “Enabling Server Software.”

For the client:

1. Use *inst(1M)* to install NetWorker software on clients. This procedure is described in the section “Installing NetWorker Software on Clients.”
2. Configure each client to recognize its NetWorker server by following the directions in the section “Configuring Clients to Recognize the Server.” No licenses are required on clients.

A Silicon Graphics NetWorker server can back up other vendors’ systems as well as Silicon Graphics clients. NetWorker software for those clients and installation information for those clients is included in NetWorker ClientPak products, which are available from Silicon Graphics.

Note: The NetWorker ClientPak products may include client software for Silicon Graphics systems running IRIX 4.0.x. On systems running IRIX 4.0.x, you should install client software from the IRIX NetWorker CD-ROM instead of from the ClientPak media.

Installing and Enabling NetWorker Software on Servers

To install and enable NetWorker software on servers, you must perform these three tasks:

- Install NetWorker software on servers, as described in the section “Installing Server Software.”
- Install and activate licenses and perform other tasks to enable the NetWorker products, as described in the section “Enabling Server Software.”
- Configure network information service (NIS) master hosts or other systems so that clients can recognize the NetWorker servers, as described in the section “Configuring Clients to Recognize the Server.”
- If you have purchased the optional Archive product (in TurboPak) and you are using an autochanger (jukebox or tape library), it is a good idea to premount an archive volume in one of the slots.

Installing Server Software

To install NetWorker on a server, you need:

- IRIX Release 5.3 or 6.0.x running on the server
- A filesystem on the server with enough free disk space to store the NetWorker indexes. For each client (including the server itself), the size of the indexes on the server is approximately 2% of the disk space in use on the client, presuming average-size files and the default backup schedule.

Caution: If you are upgrading from NetWorker 4.0.x, you cannot go back to the earlier version once you have upgraded to 4.1.1, because upgrading changes the index files.

Follow this procedure to install NetWorker server software:

1. Load the software distribution CD-ROM (p/n 812-0202-003).

2. On the server, become root and enter the *inst(1M)* command:

```
% /bin/su -
Password:
# inst
```

3. Specify the location of the NetWorker software distribution and get a list of the subsystems, for example:

```
Inst> from /CDROM/dist
Inst> list
```

For additional information on *Inst*, see the *Software Installation Administrator's Guide*.

4. Make subsystem selections in *Inst*. Use Table 2-1 to guide you in selecting subsystems. Subsystems not listed in Table 2-1 should not be installed on servers. For example, to install the required software, the reference pages, the release notes, and the InSight versions of the NetWorker manuals, enter these commands:

```
Inst> keep *
Inst> install networker4.*.server networker4.*.client_5
Inst> install networker4.man.relnotes
```

For additional information on these subsystems, see the *IRIX NetWorker Release Notes*.

Table 2-1 Subsystems for NetWorker Servers

Subsystem	Description	Required/Optional
<i>networker4.sw.server</i>	NetWorker server software	Required
<i>networker4.sw.client_5</i>	NetWorker client software for IRIX 5.3 and 6.x	Required for IRIX 5.3 and 6.x systems
<i>networker4.sw.client_4</i>	NetWorker client software for IRIX 4.x	Required for IRIX 4.x systems
<i>networker4.man.server</i>	Reference pages for NetWorker server software	Optional
<i>networker4.man.client_5</i>	Reference pages for NetWorker client software	Optional
<i>networker4.man.relnotes</i>	Release notes	Optional

Table 2-1 Subsystems for NetWorker Servers

Subsystem	Description	Required/Optional
<i>networker4.books.NetWorker_UG</i>	InSight version of the <i>IRIX NetWorker User's Guide</i>	Optional; requires IRIX 5.0 or later
<i>networker4.books.NetWorker_AG</i>	InSight version of this guide	Optional; requires IRIX 5.0 or later
<i>networker4.sw.archive.server_5</i>	Optional archive and retrieve software for server	Requires purchase of TurboPak option

5. Perform the installation and exit *inst*:

```
Inst> go
Inst> quit
```

6. Choose a filesystem for the NetWorker index files. As a rule of thumb, the amount of disk space used by the index files is about 2% of the amount of disk space used by the backed up files, presuming average-size files and the default backup schedule. You can use part of a filesystem, or dedicate a filesystem. Dedicated filesystems are typically called */nsr*.

7. If you've chosen a filesystem other than a dedicated filesystem named */nsr*, use this command to tell NetWorker where to put the index files:

```
mars# ln -s directory /nsr
```

directory is the name of a directory in the filesystem you've chosen for the index files, for example, */bigfs/nsr*.

8. Start NetWorker daemons with this command:

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

Because the licenses haven't been enabled yet with the *nsrkap* command (see the next section, "Enabling Server Software"), these messages, which can be ignored, appear in */var/adm/SYSLOG* and in mail sent to root:

```
NetWorker Registration: (notice) server disabled; No
valid Base or Advanced licenses found,
nsrd: either install new licenses or recover a valid
resource file. (15017)
NetWorker registration: (info) Server is disabled
(Install base enabler).
```

If this is the first time that NetWorker has been installed on this system, you may also see this message, which you can also ignore:

```
Verification could not mount '/nsr/mm/mmvolume'
```

After you install software on the server, you must enable it by following the instructions in the next section, and you must ensure that clients can recognize the hostname alias *nsrhost* by following the instructions in “Configuring Clients to Recognize the Server,” later in this chapter.

Enabling Server Software

When you purchase any of the NetWorker products, you must enable them on servers by installing NetLS licenses and activating the licenses with the *nsrcap(1M)* command.

All NetWorker products are included on the NetWorker software distribution CD. Table 2-2 lists the NetWorker products, their features, and licensing information.

Table 2-2 Licensed NetWorker Products

Product	Description	Licensing Information
IRIX NetWorker	Supports 10 or fewer clients and one or two tape drives; can be used with TurboPak option	IRIX NetWorker license is required on each NetWorker server.
NetWorker TurboPak option, including the archive services, Concurrent Device Support, autochanger (disk jukeboxes and tape minilibraries) support, and parallelism	Supports 10 clients and up to 64 concurrently active tape drives; allows addition of jukeboxes	IRIX NetWorker license is required on each NetWorker server. In addition, a separate TurboPak license is required for TurboPak. For archive capability,
5 Client Connection	Supports 5 additional clients	Client licenses are additive up to a maximum of 1000 clients per server.
25 Client Connection	Supports 25 additional clients	Client licenses are additive up to a maximum of 1000 clients per server.
100 Client Connection	Supports 100 additional clients	Client licenses are additive up to a maximum of 1000 clients per server.

Table 2-2 Licensed NetWorker Products

Product	Description	Licensing Information
ClientPak Software Module	Enables client to perform backups or administration tasks; available for UNIX, PC desktops, NetWare, and Windows NT clients	
Autochanger Software Module	Enables the use of disk jukeboxes and tape minilibraries; available for 1-6 slots, 1-16 slots, 1-64 slots, and more than 64 slots	One autochanger license is required for each connected jukebox or tape minilibrary

The software distribution CD-ROM is usually shipped with a temporary NetWorker IRIX NetLS license, a temporary NetLS license for an autochanger with 1-16 slots, and a temporary license for the NetWorker TurboPak option. These licenses allow you to use these products on any Silicon Graphics system for a short period.

To continue to use NetWorker products after the temporary licenses expire, you must obtain a permanent, node-locked license for each NetWorker product you purchase. If you do not already have the appropriate licenses for your NetWorker product(s), see *Network License System Administration Guide* and the *IRIX NetWorker Release Notes* for information on obtaining licenses. Additional information on installing NetLS licenses is available in the *Network License System Administration Guide*.

The next two sections describe how to enable NetWorker server software by installing and activating NetLS licenses. The section “Enabling Server Software Except Autochangers” explains how to enable all server software except the NetWorker Autochanger Software Module. The section “Enabling Autochangers” explains how to enable IRIX NetWorker Jukebox options.

Enabling Server Software Except Autochangers

Follow the steps below to enable all NetWorker server software except jukebox support. This procedure can be used when the software is initially installed and when upgrades are added, for example when a license for the NetWorker 25 Client Connection is added to the system.

1. Find the NetWorker NetLS license information stickers or sheets in your package. (Purchase the licenses from the Silicon Graphics licensing organization.)
2. Become root on the NetWorker server.

```
mars% /bin/su -  
Password:
```

3. If the subsystem *netls_eoe.sw.netls_eoe* is not already installed, use *inst* to install it from the IRIX CD that you used when installing the system software on the server.
4. To begin entering license information, edit */var/netls/nodelock*. For example:

```
## "Silicon Graphics" "NetWorker TurboPak" "A" "1995/05/29"  
546fb4684914.02.c0.1a.3d.52.00.00.00 svav6qttzsd78c5ja75ukmia  
## "Silicon Graphics" "NetWorker Autochanger/16" "A" "1995/05/29"  
546fb4684914.02.c0.1a.3d.52.00.00.00 zi3xheugriecc5ja75ukmia  
## "Silicon Graphics" "Networker 4.1.1 Archive Option" "4.1.1"  
"1995/04/06"  
546fb4684914.02.c0.1a.3d.52.00.00.00 xjzihimf6varac5ja75ukx2a  
  
546fb4684914.02.c0.1a.3d.52.00.00.00 925gjbqjmxkqgc9999998mia # Desks A  
1995/05/15
```

Using the information from the license information sheet, fill in the Product Name, Version, Expiration Date, and License Password fields.

5. Start NetWorker:

```
# /usr/etc/nwadmin &
```
6. Activate the licenses by running the *nsrkap* command as root:

```
mars# /usr/etc/nsrkap -v  
Found a NetLS license for installation:  
name: NetWorker /10  
enabler code: 2b2d6b-a3b02d-6a9c80  
license type: B10  
expires: 990 days (Nov 18, 1996)
```

The output of *nsrkap* lists each license you installed. These licenses are now activated.

Note: If you have the Registration window open at the time you activate NetWorker licenses with *nsrkap*, close it to refresh the window.

Enabling Autochangers

The procedure for enabling a NetWorker Autochanger Software Module differs from the procedure for enabling other NetWorker software. Use the procedure in this section to enable each Autochanger Software Module you purchase. You can use this procedure to enable when you install NetWorker initially and to enable autochangers that you add later as upgrades.

Follow this procedure for each autochanger:

1. Run the `jbm_enabler(1M)` command as root to tell NetWorker about the autochanger (jukebox or tape library). For example:

Note: The list of supported jukebox types is subject to change.

```
mars# /usr/etc/jbm_enabler
Enter the letter corresponding to the type of jukebox
you are installing:
a) SGI DLT2700
b) Odetics 2640
c) ADIC-1200c
d) ADIC-VLS
e) ARC-DiamondBack
f) DLI Libra Series
g) Digital DLT
h) EXB-10e
i) EXB-10i
j) EXB-60
k) EXB-120
l) EXB-210
m) EXB-400 Series
n) HP-Optical 630MB/1.3GB
o) HP-C1553A
p) Lago 340
q) Lago (Serial) Datawheel
r) Metrum (SCSI)
s) STK-9709/Odetics
t) STK-9708/Lago 380 (SCSI) Datawheel
u) ATL (Serial)
v) other
Choice? a
Installing a SGI DLT2700 jukebox.
Name you would like to assign to the jukebox device? DLT
Pathname of the control port for the jukebox
device? [/dev/scsi/sc1d510] /dev/scsi/sc4d610
```

```
Enter pathname of media drive 1
(eg: /dev/rmt/tps1d6nrnsv): ? /dev/rmt/tps4d6nrnsv
Enter the slots you wish to use for backup? [1-7] 1-7
```

Jukebox has been added successfully

The numbers that you enter for the slots represent a range of slots, not the total number of slots available. Usually, for a jukebox, you do not enter a single slot number. Entering a single slot number means that only a single slot would be used for backup and recovery with NetWorker.

2. To begin entering license information, edit `/var/netls/nodelock`, as explained in “Enabling Server Software Except Autochangers,” earlier in this chapter.
3. Start NetWorker:

```
# /usr/etc/nwadmin &
```

4. Activate the license by running the `nsrkap` command as `root`:

```
mars# /usr/etc/nsrkap -v
Found a NetLS license for installation:
  name:          JSM-10
  enabler code: 2b2d6b-a3b02d-6a9c80
  license type: J10
  expires:      990 days (Nov 18, 1996)
```

Note: If you have the Registration window open in the NetWorker Administration program (`nwadmin`) when you enable a NetWorker license with `nsrkap`, close it to refresh the window.

Configuring Clients to Recognize the Server

The clients on the network can be configured so that when they use the *nwbackup* and *nwrecover* commands (and, if TurboPak and an archive license have been purchased, the *nwarchive* and *nwretrieve* commands), it automatically points to their server. (If clients are not configured this way, users must use the “Change Server” command or speed button to specify their server.) Configure the clients by defining a hostname alias called *nsrhost* that can be resolved by each client to be the name of their NetWorker server.

Configuring clients to recognize the NetWorker server can be done in these ways:

- If your network uses the network information service NIS, use the section “Using NIS to Configure Clients.”
- If your network uses the Domain Network Service (DNS), refer to the documentation on the BIND Name Server in the *IRIS Advanced Site and Server Administration Guide*.
- If your network uses neither NIS nor DNS, use the section “Configuring Clients Without NIS or DNS.”

For more information about how clients bind to NetWorker servers, see the *nsr(1M)* reference page.

Using NIS to Configure Clients

If you have NIS on the network, follow these steps to configure clients to recognize the NetWorker server:

1. Identify the NIS master for each client by using the *ypwhich(1)* command on each client:

```
jupiter% ypwhich -m hosts
venus
```

In this example, venus is the master NIS host for the client jupiter.

2. As root on each of the master NIS hosts, edit the */etc/hosts* file, adding *nsrhost* as an alias next to the name of the NetWorker server. For example:

```
137.69.1.3 mars nsrhost
```

Note: Only one host in */etc/hosts* should have the *nsrhost* alias. If several NetWorker servers back up files on clients served by a single NIS master host, some users may have to specify their server using the “Change Server” command on the Server pulldown menu (or the *Change Server* button) of *nwadmin*.

3. On each of the master NIS hosts, change directory to */usr/etc/yp* and use the *ypmake(1)* command:

```
jupiter# cd /usr/etc/yp
jupiter# /usr/etc/yp/ypmake
```

4. To verify that each client can recognize the NetWorker server, you can enter this command on each client after waiting a suitable period for NIS databases to propagate:

```
jupiter# ypmatch nsrhost hosts
137.69.1.3 mars nsrhost
```

The hostname listed in the output should be the name of the NetWorker server.

Configuring Clients Without NIS or DNS

This section describes the procedure for configuring a client to recognize the NetWorker server when NIS and DNS (BIND) are not used on the network. This procedure can also be used if NIS or DNS is in use on the network.

1. Log on to the client as root:

```
jupiter#
```

2. Edit the file */etc/hosts*, adding *nsrhost* as an alias next to the NetWorker server name:

```
137.69.1.3 mars nsrhost
```

In this example, the client named *jupiter* will be backed up by the NetWorker server named *mars*.

Note: If you have purchased NetWorker for more than one server, only one host should have the *nsrhost* alias in each client’s */etc/hosts* file.

3. If NIS and/or DNS are in use on the network, verify that *nsrhost* is correctly aliased:

```
jupiter% /usr/etc/ping -c1 nsrhost
PING mars.engr.xyz.com (192.48.155.20): 56 data bytes
...
```

In the command, the `-c` is followed by a numeral 1. The first line of output should contain the name of the NetWorker server. If the output doesn't contain the name of the NetWorker server, modify the client's `/usr/etc/resolv.conf` file so that the local `/etc/hosts` file is checked before NIS or BIND. See the *resolver(4)* reference page for more information.

Installing NetWorker Software on Clients

Follow the procedure in this section to install NetWorker client software. Because licenses for NetWorker clients are not installed on the clients, no enabling of the client software is required.

1. On the client, become root and enter the *inst(1M)* command:

```
% /bin/su -
Password:
# inst
```

2. Specify the location of the NetWorker software distribution and get a list of the subsystems, for example:

```
Inst> from /CDROM/dist
Inst> list
```

For additional information on *inst*, see the *Software Installation Administrator's Guide*.

3. Make subsystem selections in *Inst*. Use Table 2-3 to guide you in selecting subsystems. Subsystems not listed in Table 2-3 should not be installed on clients. For example, to install the required software, the references pages, the release notes, and the InSight versions of the NetWorker manuals on an IRIX 5.2 client, enter these commands:

```
Inst> keep *
Inst> install networker4.*.client_5
Inst> install networker4.man.relnotes
Inst> install networker4.books.*
```

For additional information on these subsystems, see the *IRIX NetWorker Release Notes*.

Table 2-3 Subsystems for NetWorker Clients

Subsystem	Description	Required/Optional
<i>networker4.sw.client_4</i>	NetWorker client software	Required (install on systems running IRIX 4.0.x only)
<i>networker4.sw.client_5</i>	NetWorker client software	Required (install on systems running IRIX 5.3 only)
<i>networker4.man.client_4</i>	Reference pages for NetWorker client software	Optional (install on systems running IRIX 4.0.x only)
<i>networker4.man.client_5</i>	Reference pages for NetWorker client software	Optional (install on systems running IRIX 5.3 only)
<i>networker4.books.NetWorker_UG</i>	The InSight version of the <i>IRIX NetWorker User's Guide</i> , viewable with <i>insight(1)</i>	Optional on systems running IRIX 5.3; cannot be installed on systems running IRIX 4.0.x
<i>networker4.books.NetWorker_AG</i>	The InSight version of this guide	Optional on systems running IRIX 5.3; cannot be installed on systems running IRIX 4.0.x
<i>networker4.sw.archive.client_5</i>	Optional archive and retrieve software for client	Requires purchase of TurboPak option

4. Perform the installation and exit *inst*:

```
Inst> go
Inst> quit
```

5. Check to see if the hostname alias *nsrhost* has been set to your NetWorker server:

```
jupiter% /usr/etc/ping -c1 nsrhost
PING mars.engr.xyz.com (192.48.155.20): 56 data bytes
...
```

In the command, the `-c` is followed by a numeral `1`. The first line of output should contain the name of the NetWorker server.

If the output doesn't contain the name of the NetWorker server, NetWorker commands such as the *nwadmin* won't automatically connect you to your server; you'll have to use the "Change Server" command on the File pulldown menu to specify your server.

To set the hostname alias *nsrhost* to your NetWorker server, follow the directions in the section "Configuring Clients to Recognize the Server" in this chapter.

Removing NetWorker

If you need to remove the NetWorker software and data files on a server for any reason, follow these steps:

1. Shut down the NetWorker server using the *nsr_shutdown(1M)* command:

```
mars# /usr/etc/nsr_shutdown
```

2. Use the *versions(1M)* command to remove any or all of the NetWorker software. For example, to remove all of the software, enter this command:

```
mars# versions remove networker4
```

3. Remove the data files by entering these commands:

```
mars# cd /nsr
mars# rm -rf .
mars# cd /
mars# rm -f /nsr
```

You can remove NetWorker software on a client with the *versions* or *inst* commands.

Using NetWorker Windows and Menus

This chapter introduces NetWorker windows and menus and explains

- starting the NetWorker programs
- using NetWorker menus
- printing NetWorker window contents
- saving NetWorker window contents
- using NetWorker generic buttons

Starting the NetWorker Programs

The NetWorker graphical user interface makes learning and using NetWorker straightforward. The intuitive windows, pulldown menus, and dialog boxes guide you through the backup and recover processes. To use the NetWorker window interface, you must have the X Window System running on your system.

The NetWorker software consists of three programs:

- *NetWorker Administrator* (*nwadmin*) for managing and configuring
- *NetWorker Backup* (*nwbackup*) for backing up data
- *NetWorker Recover* (*nwrecover*) for recovering files

Note: Your site might also have the optional NetWorker Archive programs (*nwarchive*) and (*nwretrieve*). For information on these programs, see Chapter 5, “Archiving and Retrieving,” in the *IRIX NetWorker User’s Guide*.

Only users listed on the NetWorker administrator list have permission to use and make changes to the administrator version of NetWorker. Users not included on the administrator list can start and open the NetWorker

Administrator program, but many of the commands in the pulldown menus will be grayed out, making it impossible to open the related windows.

If you are the system administrator, you probably want all three programs available on your workstation. Typical users who only need to back up and recover files do not need the administrator program running on their systems.

Note: If you start NetWorker from a remotely mounted directory, the message `Using server server_name as server for client_name.` might appear. 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.

Starting 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 a monitor for server activity.

To start the administrator program, enter the `nwadmin` command in a shell window:

```
# /usr/etc/nwadmin &
```

If the NetWorker Administrator icon already appears on your screen, double-click to open it. The icon shows two workstations and the name of your NetWorker server; see Figure 3-1.



Figure 3-1 NetWorker Administrator Icon

Note: You might have several NetWorker servers at your site. To easily identify each server, NetWorker automatically replaces the NetWorker Administrator caption in the window title bar with the name of the server. To restore the original caption, change the resource `iconServerName` by using the X11 resources. Begin by entering the `xrdb -m` command at the system prompt, and then change the setting to `False`. Refer to the *X Window System's User's Guide* for more detailed instructions.

The NetWorker Administrator window appears, as shown in Figure 3-2.

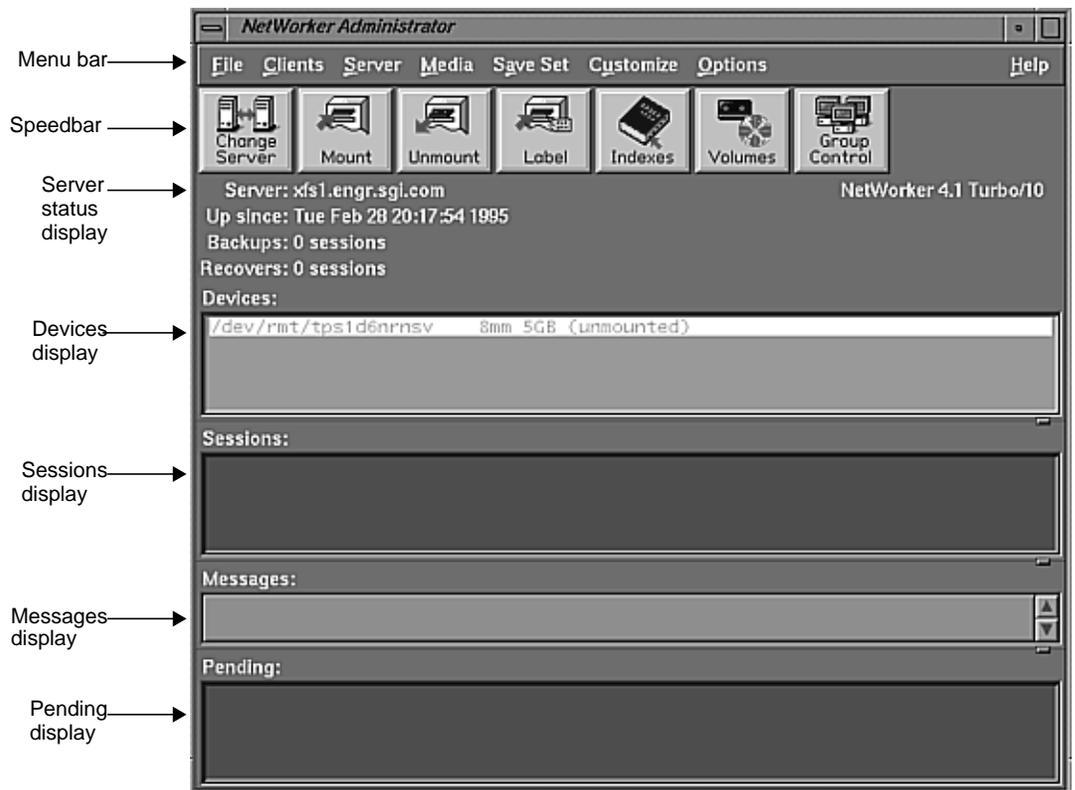


Figure 3-2 NetWorker Administrator Window

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 pulldown menus. These are explained in the section “Using NetWorker Menus” in this chapter.
- Speedbar** Provides convenient buttons for performing common NetWorker tasks; click a button instead of selecting the corresponding command from a pulldown menu.



Change Server opens the Change Server dialog box; use instead of the “Change Server” option in the Server menu.



Mount mounts a backup volume in a backup device; use instead of the “Mount” option in the Media menu.



Unmount unmounts a backup volume in the backup device; use instead of the “Unmount” option in the Media menu.



Label labels a volume in the backup device; use instead of the “Label” option in the Media menu.



Indexes opens the Indexes window; use instead of the “Indexes” option in the Clients pulldown menu.



Volumes opens the Volume Management window; use instead of the “Volumes” option in the Media menu.



Group Control opens the Group Control window; use instead of the “Group Control” option in the Server menu.

Server status display

Shows how long the current server has been online, the number of backups and restores the server has completed since it has been online, and the amount of data contained in the backups and restores. The upper right corner of this display shows the version of NetWorker running on the server.

- Devices display**
Shows the devices known to the server and the backup volume, if any, that is mounted on the device.
- Sessions display**
Provides information on which clients are currently backing up, browsing, or recovering files.
- Messages display**
Provides a history of the NetWorker server activity.
- Pending display**
Shows messages that may require operator attention.

Starting the Backup Program

Use the NetWorker Backup window to back up client files.

Note: This window does not provide any capability for administering or configuring NetWorker. This feature is beneficial to the user who simply wants to back up files and does not want to be concerned with the administration of NetWorker. The backup program icon includes the “NetWorker knight” and the caption “NW Backup.”

Start the software program by entering the *nwbackup* command at the system prompt:

```
# /usr/etc/nwbackup &
```

If the NW Backup icon is already displayed on your screen, as shown in Figure 3-3, select it to open the window.



Figure 3-3 NetWorker Backup Icon

Note: For complete information about the NetWorker Backup program, refer to the *IRIX NetWorker User's Guide*.

Starting the Recover Program

Use the NetWorker Recover window to recover client files. This window does not provide any capability for administering or configuring NetWorker. This is beneficial to the user who simply wants to recover files and does not want to be concerned with the administration of NetWorker. The recover program icon includes the “NetWorker knight” and the caption NW Recover.

Start the software program by entering the `nwrecover` command at the system prompt.

```
# /usr/etc/nwrecover &
```

If the NW Recover icon is already displayed on your screen, as shown in Figure 3-4, select it to open the window.



Figure 3-4 NetWorker Recover Icon

Note: For complete information about the NetWorker Backup program, refer to the *IRIX NetWorker User's Guide*.

Using NetWorker Menus

This section includes information about the pulldown menus in the NetWorker Administration window and how to use them.

To use NetWorker menus, you can point and select with a mouse, or use keyboard shortcuts. To pull down a menu using the keyboard shortcuts, hold down the `<Alt>` key and press the key that corresponds to the underlined letter of the selected menu; the menu appears. After releasing both keys, press the key that corresponds to the underlined letter in the command name.

The File Menu

The File pulldown menu, shown in Figure 3-5, contains only the Exit command for exiting NetWorker.



Figure 3-5 File Pulldown Menu

However, many of the other NetWorker windows have a different version of the File menu that contains the Save, Print, and Exit commands, as shown in Figure 3-6.



Figure 3-6 File Menu in Other NetWorker Windows

Typically, the Save and Print commands are available only when you use a NetWorker window in the tabular view. See “Displaying a Tabular View” later in this chapter for more information on these commands.

To exit NetWorker, select “Exit” from the File menu, or press both the <Alt> key and the ⌘ key at the same time, release them, and then press x .

Clients Menu

The Clients menu, shown in Figure 3-7, contains commands for configuring your NetWorker clients and for monitoring and managing their indexes. “Configuring a Server’s Clients” in Chapter 6 is a complete guide to the commands in the Clients menu.



Figure 3-7 Clients Menu

Server Menu

The Server menu, shown in Figure 3-8, contains commands for setting up, changing, and configuring the NetWorker server. Chapter 6, “Configuring a NetWorker Server,” is a complete guide to the commands in the Server menu.



Figure 3-8 Server Menu

Media Menu

The Media menu, shown in Figure 3-9, contains the commands for using and managing the NetWorker server’s backup volumes. It also includes the commands for configuring and monitoring the devices used for backup and recovery, whether they are stand-alone backup devices or jukeboxes.

Chapter 7, “Configuring and Managing Volumes and Backup Devices,” and Chapter 11, “Using NetWorker with Jukeboxes,” explain commands in the Media menu.



Figure 3-9 Media Menu

Save Set Menu

A save set is a set of files or a filesystem that you define for NetWorker backup. The Save Set menu, shown in Figure 3-10, contains commands for recovering entire save sets and for “cloning,” or making identical copies of save sets. Chapter 9, “Recovering and Cloning Save Sets,” is a complete guide to the commands in the Save Set menu.

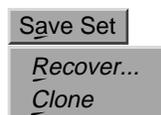


Figure 3-10 Save Set Menu

Customize Menu

The Customize menu, shown in Figure 3-11, contains selections for further customizing and configuring your NetWorker backups. The menu contains commands for configuring schedules, policies, directives, label templates, groups, and notifications. Chapter 6, “Configuring a NetWorker Server,”

and Chapter 7, “Configuring and Managing Volumes and Backup Devices,” explain commands in the Customize menu.



Figure 3-11 Customize Menu

Options Menu

Use the Options menu, shown in Figure 3-12, to make selections for customizing the NetWorker interface. The Options menu includes the Polling Interval, Speedbar, and Tone commands: you can set the polling interval, choose whether or not to display the speedbar, and cause a tone to be sounded when a pending message arrives.

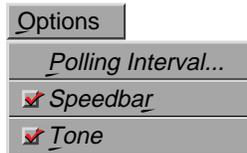


Figure 3-12 Options Menu

Changing the Polling Interval

The Polling Interval command 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 3-13.



Figure 3-13 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.

Enabling or Disabling the Speedbar

NetWorker provides the convenience of a speedbar for the functions you perform most often. By default, the speedbar is enabled. Use the *Speedbar* toggle button in the Options menu to turn the speedbar on and off.

This manual informs you when a speedbar button is available for accomplishing the same task that you can accomplish with the commands from a pulldown menu.



Figure 3-14 Speedbar

Enabling or Disabling Tone

Select the *Tone* toggle button in the Options menu if you wish to have a tone sounded when a pending NetWorker message arrives and when it is leaves. The tone sounds only once when the message first arrives and once after the task requested by the pending message has been completed.

Help Menu

All NetWorker windows contain a *Help* menu or button that offer you a variety of choices for help. Figure 3-15 shows the Help menu.

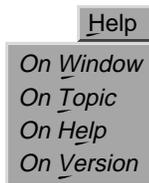


Figure 3-15 Help Menu

The Help menu contains four commands.

- “On Window” displays a help screen with information about the currently selected window.
- “On Topic” displays a list of topics from which you select.
- “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 are located in many of NetWorker's windows. The View menus contain the Tabular, View Details, and Search commands. Use these commands to view the contents of a window in a tabular format or to see more details in a current window.



Figure 3-16 View Menu

Displaying a Tabular View

NetWorker can display information in a tabular view for most of the administrative windows, enabling you to view information or data in a more compact format. The tabular view is especially helpful for large sites that have a lot of data you want to view in a single display.

For example, in the Clients window there are two ways to view client information:

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

For example, to display the tabular view for the Clients window, follow these steps:

1. Select the Client Setup command from the Clients menu. The Clients window appears, as shown in Figure 3-17.

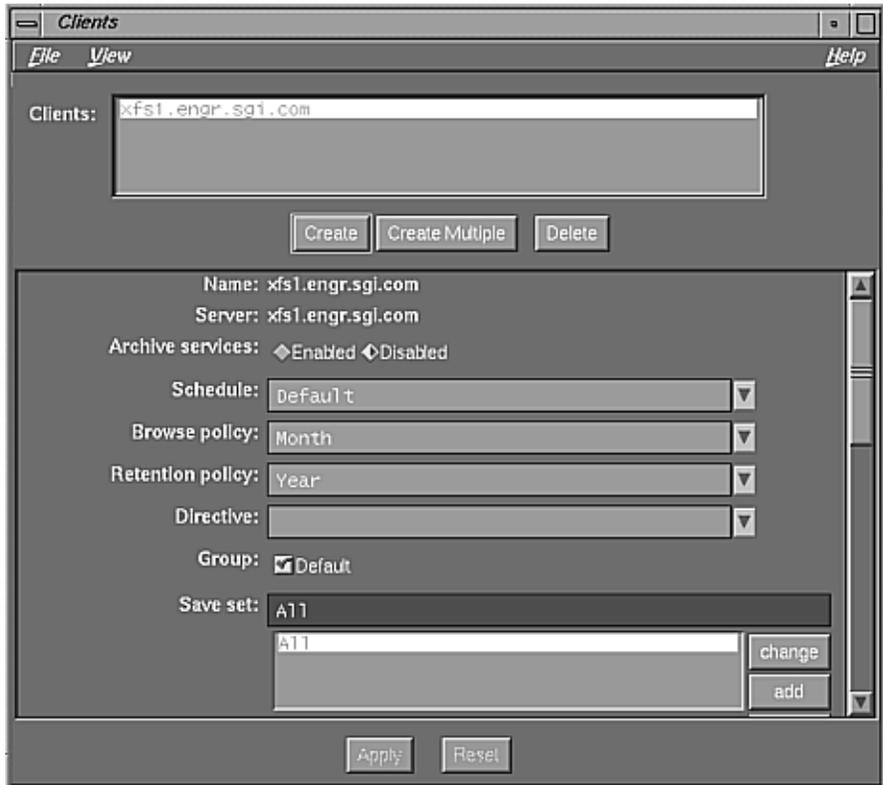


Figure 3-17 Clients Window

2. Select the *Tabular* toggle button in the View menu. The Clients window changes to a tabular view, as shown in Figure 3-18.



name	server	archive services	schedule
marah.wpd.sgi.com	xf44.enqr.sgi.com	Enabled	Default
xf44.enqr.sgi.com	xf44.enqr.sgi.com	Enabled	Default

Figure 3-18 Clients Window: Tabular View

Use the scroll bars or resize the window to see more of the client configuration information.

To return to the normal Clients window, reselect the *Tabular* toggle button from the View menu.

The tabular view is available for all of the administrative windows except for the Save Set Clone and Schedules windows.

Displaying Details

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

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

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

1. Select “Groups” in the Customize menu. The Groups window appears.
2. Select the *Details* toggle button from the View menu, as shown in Figure 3-19.

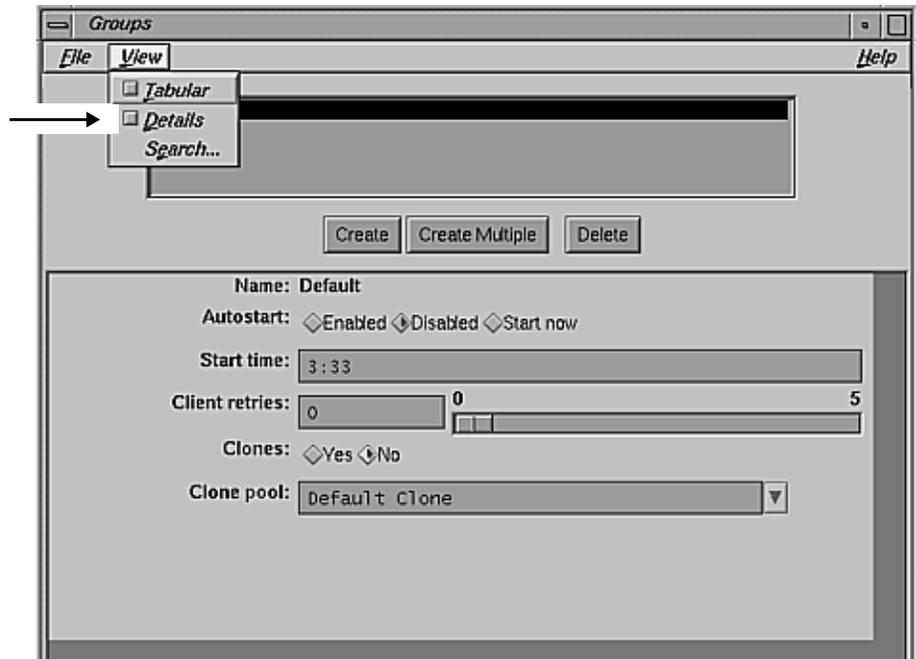


Figure 3-19 Groups Window: Displaying Details

The Groups window changes to show you more detailed information, as shown in Figure 3-20.

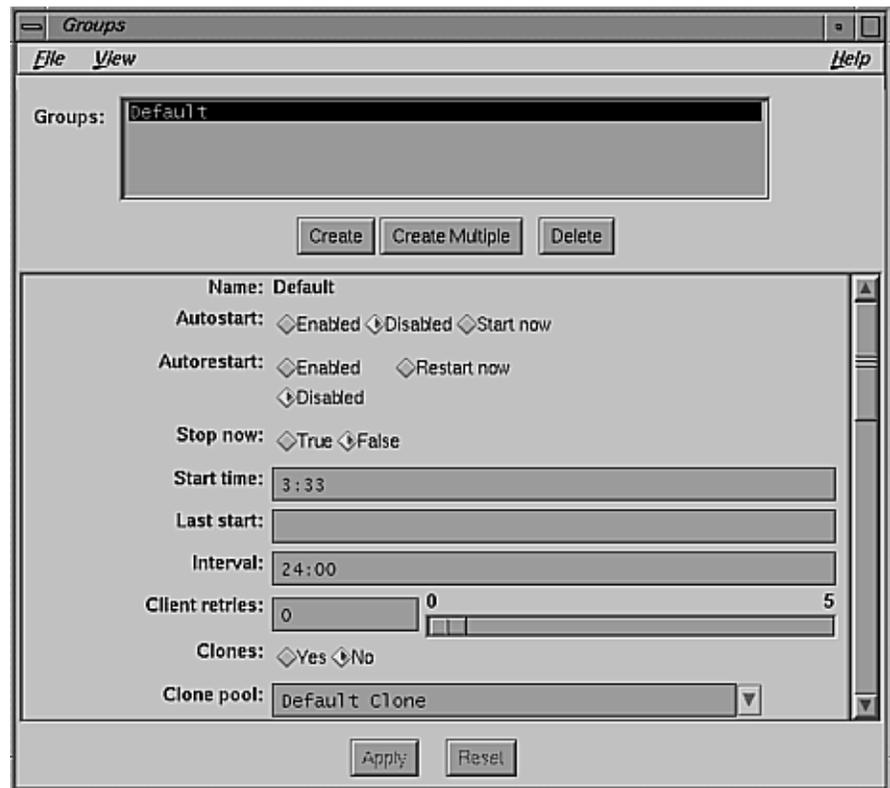


Figure 3-20 Groups Window: Detail View

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

Searching for Resources

The Search command provides you with an alternative method for locating or finding a NetWorker configuration or resource. For example, if you back up a large number of clients, this feature avoids the necessity of searching through the entire scrolling list of client names for a single client.

For example, to search for a NetWorker client, follow these steps after opening the Clients window:

1. Select Search from the View menu. The Search dialog box appears, as shown in Figure 3-21.



Figure 3-21 Search Dialog Box

2. Enter the name of the client you would like to locate in the Search for resource field.
3. Click *Ok*.

Printing Window Contents

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

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

1. Select the Tabular command from the View menu. The contents of the NetWorker window changes to a tabular view.
2. Select the Print command from the File menu. The Print Options dialog box appears, as shown in Figure 3-22.



Figure 3-22 Print Options Dialog Box

This Command field displays the command *lp*. You can modify the command, or you can enter a completely different command. Follow these steps:

1. In the Command field, specify a printer other than the default printer, if desired. For example:

```
lp -dfranklin
```

2. Enter the number of copies you want in the Copies field.
3. Click *Ok* to print.

Having the ability to print the contents of a window provides you with an easy method for maintaining records on NetWorker activities and configurations.

Saving Window Contents

You can use the Save command only 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. Select the Tabular command from the View menu. The contents of the window changes to a tabular format.
2. Select the Save command from the File menu. The Save Options dialog box appears.

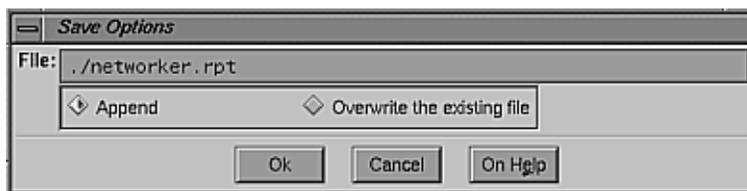


Figure 3-23 Save Options Dialog Box

The File field is set to the environment variable NWREPORT if it exists; otherwise the default is *networker.rpt*.

3. Enter the name of the file in the File field. To update or append the existing file, click the *Append* button. To overwrite the existing file, click the *Overwrite the existing file* button.
4. Click *Ok* when you have completed your selections.

Using NetWorker Generic Buttons

NetWorker uses generic buttons in many of the windows. Most of them are self-explanatory; however, listed below is a description of each one.

- *Cancel* cancels configuration choices or an operation.
- *Create* adds or creates a new configuration.
- *Ok* closes a window or to apply configuration selections.
- *Delete* deletes an item or NetWorker configuration.
- *change* changes a selected item or configuration.
- *add* adds a new item or configuration.
- *Apply* applies your selections.
- *Reset* resets the window to the previous selections in the window.

Create Multiple Button

Use the *Create Multiple* button to create multiple items or configurations all with the same characteristics or selections. Using this button saves you from having to make identical selections over and over again for a variety of NetWorker configurations.

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

1. Open the Clients window by selecting Client Setup from the Clients menu.
2. Click the *Create Multiple* button, as shown in Figure 3-24.



Figure 3-24 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. After adding the client names, select the configurations for all of them at once.
5. When you are finished selecting the configurations for all of the clients, click the *Apply* button.

Print and Save Buttons

Several NetWorker windows contain Print and Save buttons for printing and saving the contents of the window. These buttons work the same way as the Save and Print commands explained in “Printing Window Contents” and “Saving Window Contents” earlier in this chapter.

NetWorker Test Drive

This chapter leads you through the process of “test driving” the NetWorker software. This chapter introduces you to the benefits of key NetWorker features.

In this test drive you will

- change the NetWorker server
- label and mount a backup volume
- perform an unscheduled backup and recover
- view the Schedules window for information about scheduling backups
- view the Clients window for information about your NetWorker clients
- enable the backup in the Groups window
- perform a scheduled backup immediately by pressing the *Start* button in the Group Control window
- verify the backup in the Group Control window
- learn about the NetWorker Support for Jukebox and Concurrent Device Support for NetWorker options

Changing the NetWorker Server

Select “Change Server” to choose another NetWorker server to monitor or configure. Most networks need only one NetWorker server to back up all the systems on the network. If you have more than one NetWorker server, however, choose the one you want to monitor or configure by selecting its hostname in the Change Server window, shown in .



Figure 4-1 Change Server Window

To change NetWorker servers, follow these steps:

1. Select “Change Server” from the File menu in the main window. The Change Server window appears.
2. Click the *Update Server List* button to display all of the NetWorker servers on your network. NetWorker broadcasts to find NetWorker servers; however, the broadcast may not reach all servers on your network, so the list may be incomplete.
3. Highlight the hostname of the desired server in the Server scrolling list, or enter a valid hostname in the Server field.
4. Click *Ok* to monitor and configure the new server in all the NetWorker windows.

Note: All operation and administration tasks apply to the server that has been selected.

Labeling and Mounting the Backup Volumes

NetWorker requires that all volumes have a label before they are used to back up files. NetWorker uses these labels to track and identify the contents of the backup volumes. When you start a recover, NetWorker automatically determines the volume(s) that are needed and asks for them by the label name. This powerful feature of NetWorker means that you can quickly locate the backup volumes required to recover files.

The mounting operation is automatic when NetWorker is used with a jukebox.

Use the Label dialog box to label and mount backup volumes. Follow these steps to label and mount your backup volumes:

1. Find a blank piece of media you can use for a backup.
2. Load it into the NetWorker server backup device.
3. Click the *Label* speedbar button.

The Label dialog box appears, in which you enter the name of the volume.

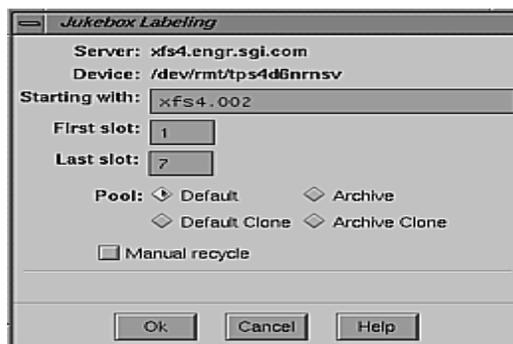


Figure 4-2 Label Dialog Box

4. Enter a name, such as “demo.001” into the Volume name field.

You do not need to select a pool because the Default pool is already preselected.

5. Click the *Ok* button.

The name of the mounted volume appears next to the device name in the Devices display in the NetWorker Administrator window.

See “Labeling Backup Volumes” and “Mounting and Unmounting Backup Volumes” in Chapter 7 for more information on NetWorker labeling and mounting features.

Trying a Backup and Recover

Typical day-to-day operation of NetWorker includes scheduled, unattended backups. NetWorker also supports unscheduled backups of selected files that are initiated by a user from a client system. Users can also initiate recovers.

Tip: You should now read the section Quick Tour in Chapter 2 of the *IRIX NetWorker User’s Guide* to become familiar with unscheduled backups and recovers.

After completing the Quick Tour in the *User’s Guide*, you should be familiar with the NetWorker recover features. Because of the online index, users can browse their backed-up files and quickly identify files to be recovered; a system administrator does not have to worry about locating an important file when the user cannot remember the filename. NetWorker also automatically identifies the volumes where the files are stored. Adding a jukebox to your backups automates the recover process.

NetWorker offers many features to help you recover files in addition to those introduced in the Quick Tour.

Viewing the Backup Schedule

Most of your backup activity occurs unattended through NetWorker's scheduling features. This section describes the Schedules window.

Open the Schedules window by selecting "Schedules" from the Customize pulldown menu. The Schedules window is shown in Figure 4-3.

The Schedules window displays an on-screen calendar that you use to set up a backup schedule, or you can select a preconfigured schedule. These schedules are based on feedback from NetWorker customers about the most popular ways to perform network-wide backups.

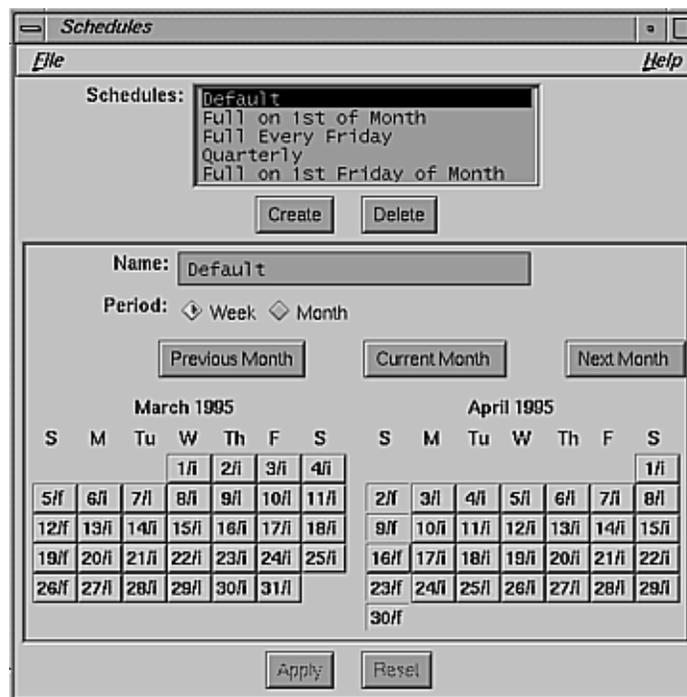


Figure 4-3 Schedules Window

Notice each calendar day displays an "f" or an "i" next to the day. The "f" represents a full backup, and the "i" represents an incremental backup.

Your backup server is also treated as a NetWorker client. Your server is automatically backed up using the Default schedule, which has pre-configured selections. It completes a full backup every Sunday, followed by incremental backups on the other days of the week.

Note: Regardless of the selections in the preconfigured schedules, NetWorker always performs a full backup the first time it backs up files.

For more information about schedules, see “Determining Backup Schedules” in Chapter 5.

Viewing the Clients Window

After NetWorker is installed on your server, the server is backed up as a client of itself using the preconfigured selections in the Clients window, which is shown in Figure 4-4.

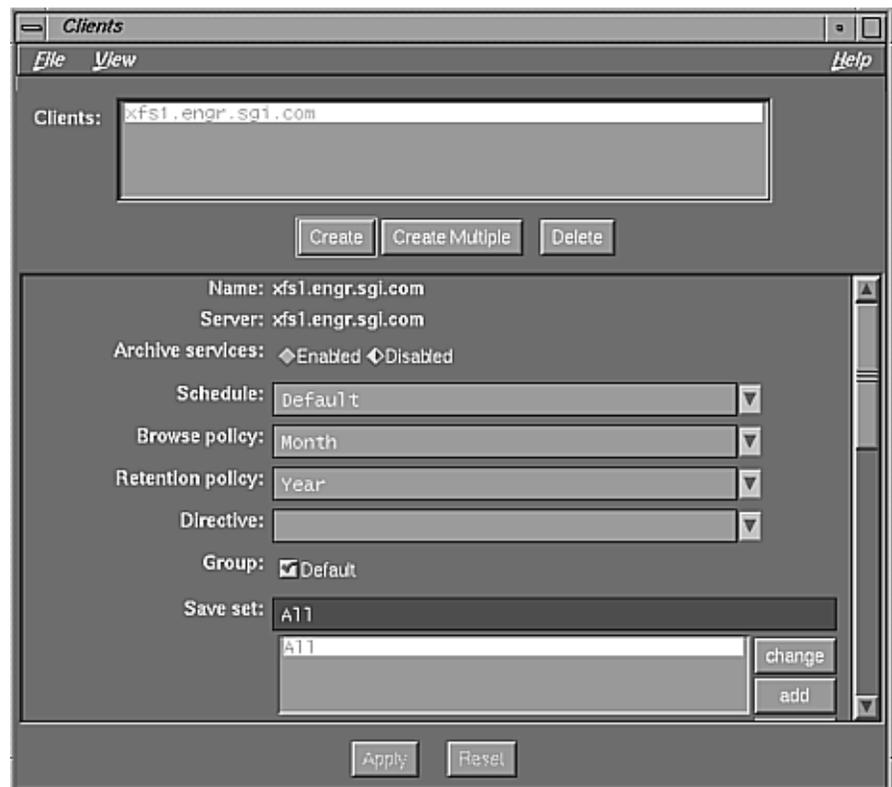


Figure 4-4 Clients Window

The Clients window displays the name of your backup server in the Clients field. Preconfigured selections are represented by the highlighted buttons. All of the configurations you select for your NetWorker client are created in the Clients window.

If you are backing up other client systems to your server, all you have to do (after installing NetWorker on the clients) is configure your clients using the Clients window. NetWorker can then back up your clients using the selected configurations.

You do not need to make any changes in this window for now. You may wish to make changes after you have created more groups and clients and are setting up different backup schedules.

The client configuration features are discussed in detail in the section “Configuring a Server’s Clients” in Chapter 6.

Enabling the Scheduled Backup

Now that you have viewed the Schedule and Clients windows, enable the Default group to which your backup server belongs by choosing Enabled in the Groups window.

By enabling the group Default, you are telling NetWorker to back up all the clients in that group according to their preconfigured selections. Because your NetWorker server is preconfigured to belong to the group Default and the Default backup schedule, it is backed up based on those selections.

Open the Groups window, shown in , by selecting “Groups” from the Customize pulldown menu.

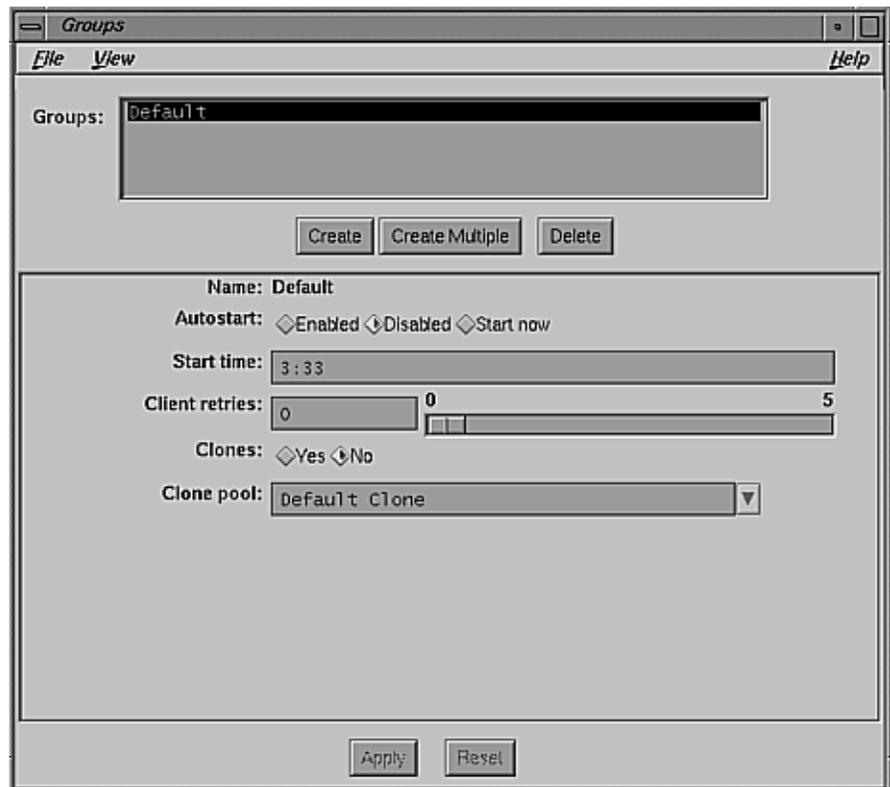


Figure 4-5 Groups Window

Notice that the Start time is 3:33 a.m. (the time is given in 24-hour format). Also notice that the Autostart feature is Disabled as the preselected choice.

To enable the Default group, follow these steps:

1. Select the *Enabled* button in the Autostart choices.
2. Click the *Apply* button to apply the change.
3. Select “Exit” from the File menu to close the Groups window.

Use the Start time field to determine what time you want NetWorker to back up a group of clients. To change the time, simply enter the new start time.

Your NetWorker server is automatically included in the group Default. By enabling the Default group, your server and any other clients you add to the Default group in the future are backed up at 3:33 every morning, starting tomorrow.

The Groups window is described in detail in the section “Configuring Backup Groups” in Chapter 6.

Starting a Backup Now

You have looked at several different windows that have shown you that your NetWorker server is backed up according to preconfigured selections. However, you should start a backup now to see how NetWorker backs up a client.

You can start a scheduled backup immediately in two ways:

- by pressing the *Start* button in the Group Control window
- by pressing the *Start now* button in the Groups window

A backup begins right away for the group selected in the Groups scrolling list—the Default group, in this example. By starting a backup now, you are essentially overriding the predefined backup time for the Default group (3:33 a.m.). However, the Default group still also backs up at 3:33 a.m. as scheduled.

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

1. Select the “Group Control” command from the Server pulldown menu. The Group Control window appears, as shown in .

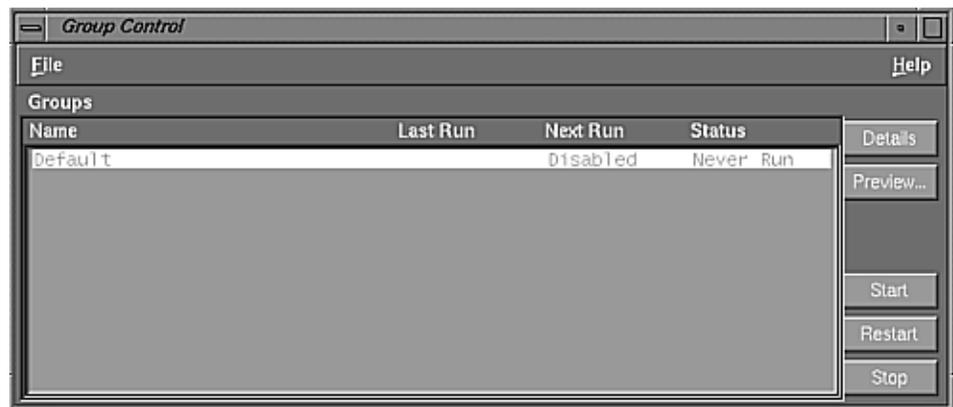


Figure 4-6 Group Control Window

2. Select the Default group in the Groups scrolling list.
3. Click the *Start* button.

NetWorker begins a backup of the clients in the Default group according to the Default backup schedule, and displays *running* in the Status column of the Group Control window.

After clicking the *Start* button, keep the Group Control window open so you can monitor the progress of your backup.

4. For more detailed information about the files being backed up, click the *Details* button. The Group Control Details window appears, as shown in . After NetWorker has successfully backed up your files, the *Finished* message appears in the Status field.

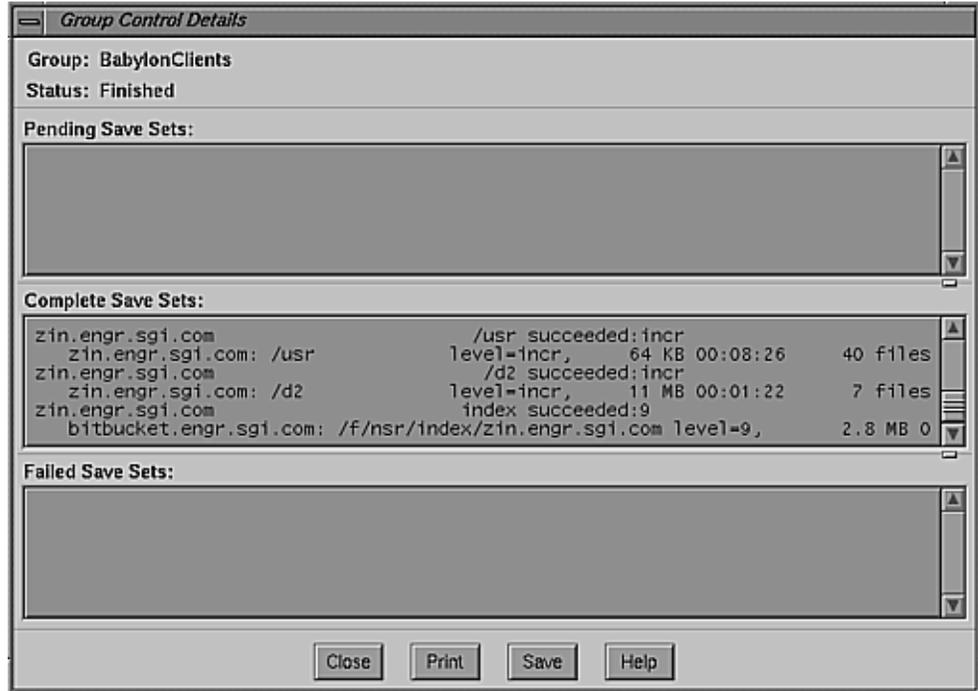


Figure 4-7 Group Control Details Window

For more information about the Group Control window, see the section “Monitoring and Controlling Backups” in Chapter 8. For more information about the Group Control Details window, see the section “Viewing Details of a Scheduled Backup” in Chapter 8.

Using an Autochanger

If you use an autochanger (jukebox or tape library) as your backup device, most of your media handling operations are automatic. Mounting volumes is automatic, and media requests are also filled automatically.

You need to purchase a NetWorker Autochanger Software Module in order to use your autochanger with NetWorker.

See Chapter 11, “Using NetWorker with Jukeboxes,” for instructions on configuring and operating autochangers.

Using Concurrent Devices

A powerful feature that enhances server efficiency during backups is optional *concurrent device support*, included in the TurboPak option. This support is used when you have more than one device attached to the NetWorker server and want to use all of the devices for backups simultaneously. It spreads the backups evenly across the devices.

See the section “Setting Concurrent Device Use” in Chapter 6 for instructions on how to configure your server to operate devices concurrently.

Note: If you are not sure whether your site has this feature, see the release notes.

Displaying NetWorker Licenses

To display installed, enabled NetWorker licenses, open the Registration window by selecting the “Registration” choice from the Server menu. The Registration window appears, similar to the one shown in Figure 4-8.

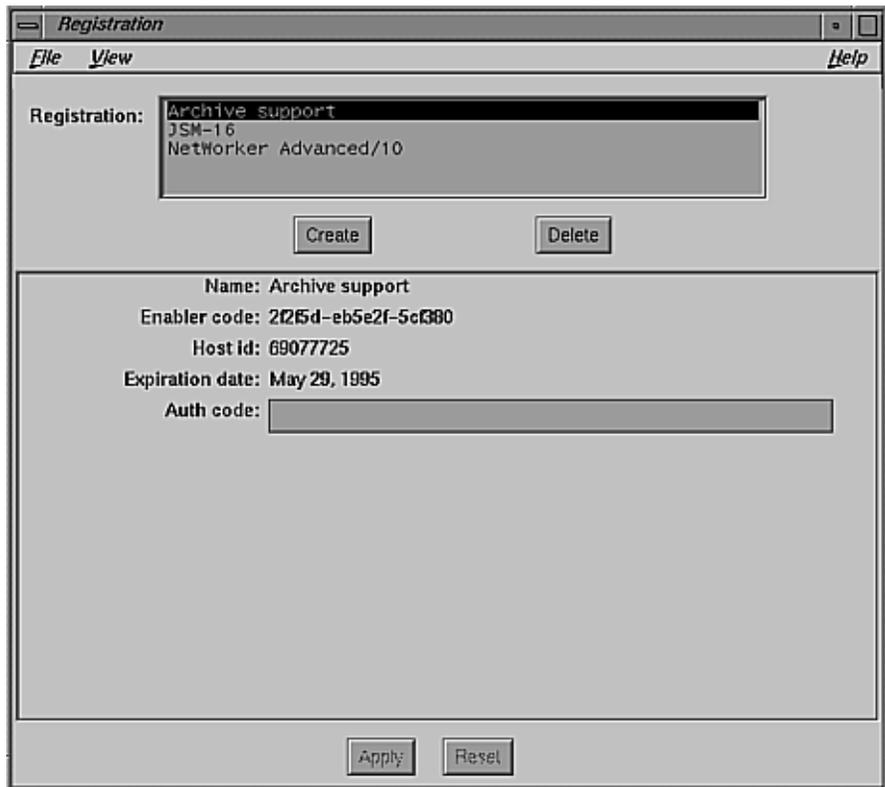


Figure 4-8 Registration Window

The Registration window shows each of the NetWorker licenses that have been installed on the server (using `/var/netls/nodelock`) and enabled (using `nsrkap(1M)`). For information on installing and enabling licenses, see the sections “Enabling Server Software” and “Enabling Autochangers” in Chapter 2. This version of NetWorker does not use the Registration window to install and activate licenses.

The Next Step

Here are some suggestions for further configuration tasks you might want to try:

- Add more clients to your Default backup group and enable the group for an overnight backup. In the morning, check the results of the scheduled backup in your mail.
- Create new groups to conform to your organizational structure and define schedules for each group. For example, create an Administration group with a Monday Full schedule; a Marketing group with a Tuesday Full schedule; an Engineering group with a Wednesday Full schedule; and a Sales group with a Thursday Full schedule.
- Run some recovery tests. After completing one or more backups, delete a file from your system. Using the Recover window and the “Versions” command, recover the correct version of the file.
- Simulate a hard disk crash and then recover your files using the disaster recovery procedures described in Appendix B, “Recovering From a Disk Crash.”

You have now taken a quick “test drive” of IRIX NetWorker. You have seen that NetWorker is a powerful tool for protecting files across entire networks. More importantly, NetWorker is able to recover your files quickly and efficiently in the event of loss.

You have also seen that NetWorker is easy to administer and configure using the X Window system graphical user interface. The user interface allows you to monitor and configure NetWorker from any node on the network.

NetWorker Configuration and Policy Issues

This chapter provides examples and suggestions for you to consider while you are designing your NetWorker environment. It also offers background information to help you understand the logic behind the NetWorker backup schedule and index policy features.

This chapter explains

- determining backup schedules
- determining browse and retention policies (index policies)
- determining jukebox policies

Note: Practical instructions for setting backup schedules, and browse and retention policies are in Chapter 6, “Configuring a NetWorker Server.” Instructions for setting up NetWorker for jukeboxes are in Chapter 11, “Using NetWorker with Jukeboxes.”

Determining Backup Schedules

The NetWorker server backs up each client system across a network according to a backup schedule. Schedules are created in the Schedules window and assigned to individual clients in the Clients window.

Schedules can be simple or 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 and explains the default schedules provided with NetWorker.

When you design backup schedules, consider these questions:

- 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 the network-wide backups?
- Do you want to be able to use just a few backup volumes to recover from an entire disk crash?

The typical capacity of an 8 mm tape cartridge is about 5 GB, and the maximum transfer rate is around 400 KB per second. Systems generally cannot sustain that transfer rate, but if they could, it would take more than 6 hours, for example, to back up 10 GB of data. To back up a network with a lot of data, you can use more than one schedule to stagger the full backups over several days. For example, Silicon Graphics systems can maintain a rate of 350 KB/second on an 8mm cartridge tape and 1.0 MB/second on a DLT.

Think about how many backup volumes you want to keep—this number depends on how often the data changes, and how long you want to keep the online backups.

If you run only incremental backups every night, you need more backup volumes to fully recover from a disk crash. If a site has 10 GB of data, and 5% of all the data is modified each day, that means 500 MB of data needs to be backed up every day. At 400 KB per second, 500 MB takes about 25 minutes to back up, and fills up about one tenth of an 8 mm cartridge tape. If you are maintaining the backups for three months, you will have about 12 tapes of backups to keep on your shelves (or in the autochanger).

You also need to decide on a policy for recovering files. For example, if the users expect to be able to recover any version of a lost file for at least three months, you must maintain all the backup volumes for the three-month period. On the other hand, if the users expect to be able to recover only the latest version of a lost file, you can use level 1–9 backups to decrease the quantity of backup volumes you need to maintain.

The following sections explain

- NetWorker backup levels
- backup time requirements
- using compression
- staggering the backup schedules
- balancing convenience and security
- using save sets
- NetWorker preconfigured backup schedules

NetWorker Backup Levels

A backup schedule specifies the level of backup that NetWorker performs for a client on each day of a weekly or monthly period. NetWorker supports four backup levels:

full

Every file is backed up, no matter when it last changed.

level 1–9

Files are backed up that have changed relative to the most recent lower backup level. Each backup level is represented by a number, 1 through 9, where 1 represents the fullest backup, and 9 represents the most minimal backup.

incremental

Files are backed up that have changed since the last backup, regardless of the level.

skip

No files are backed up; skip the backup. If you do not use an autochanger, you may want to skip a backup on a holiday if you know that no one will be around to change backup volumes on the server.

If you do not need to maintain every version of a backed-up file online, you can use a backup scheme that includes occasional full backups followed by level 1–9 and incremental backups during the cycle. Different backup levels allow you to trade off the number of backup volumes and amount of time

required to complete a backup versus the number of backup volumes and amount of time it takes to recover from a disk crash.

Figure 5-1 diagrams how backup levels work.

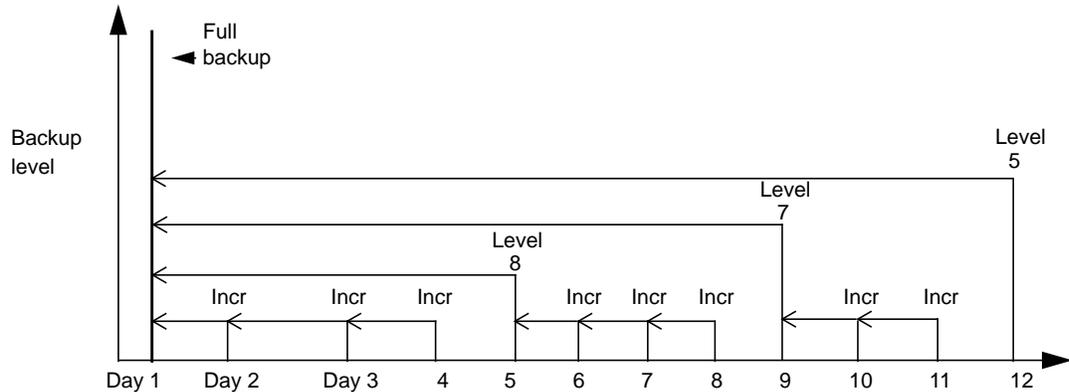


Figure 5-1 Backup Levels

Assume you use a new backup volume for each day's backup. On day 1, a full backup is run. 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. On day 4, the incremental backs up everything that has changed since day 3. At this point, you have four backup volumes. To recover from a disk crash, you need all four of them: the one with the full backup (day 1), 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, all you need is the full backup volume and the level 8 backup volume. If you had to recover from a complete loss of a disk, you have reduced the number of backup volumes you need to only two.

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. Planning your backup strategy carefully should allow you to

recover everything on a disk with a maximum of four backup volumes, assuming that each day's backup can fit on one volume.

The rest of this section compares various backup levels and combinations of levels:

- full backups versus incremental backups
- level backups
- combining full, incremental, and level backups

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 schedule is simple to set up and execute, and it makes recovering from a disk crash easy—you simply need the last full backup volume.

Issues to consider are listed below:

- Full backups take more time to execute than do 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 do incremental or level backups.

You might decide to schedule a full backup at the beginning of the period and then schedule incremental backups the rest of the period. This schedule minimizes the amount of time that the backups take, minimizes the size of the backups, and causes the NetWorker indexes to grow at a slower rate. 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).

Level Backups

You can use level 1 through level 9 backups to moderate between the two extremes described in the preceding section. Level 1 through level 9 backups allow you to set up a schedule for each client, balancing the 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 because 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. Figure 5-2 diagrams a full backup followed by level 9 to level 1.

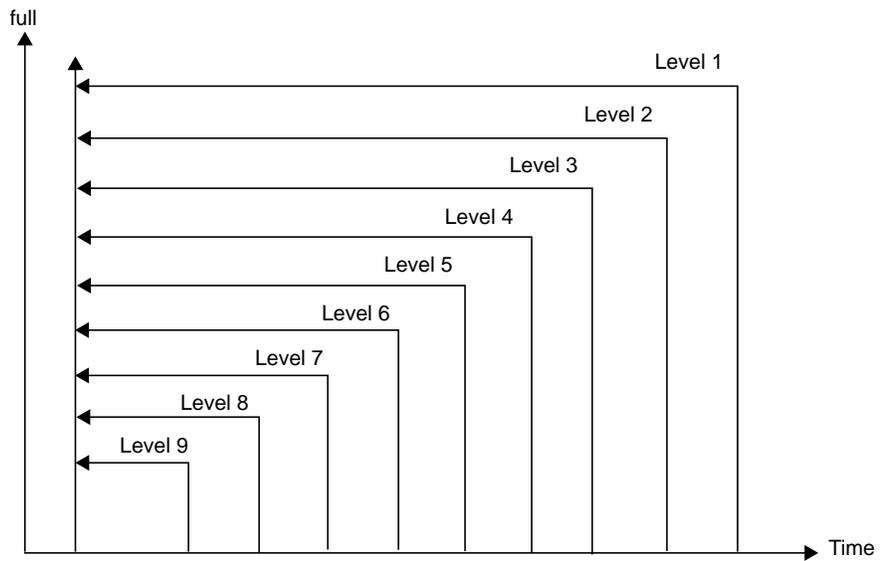


Figure 5-2 Full Backup Followed by Level 9 to Level 1

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 each day, there are more changed files to back up, so the backups take longer to complete.

Figure 5-3 diagrams a backup schedule that also starts out with a full backup, but the level backups that follow are in reverse order: starting with a level 1 on the first day following the full backup, on down to a level 9 backup. Each day, the backup backs up only the files that have changed on that day.

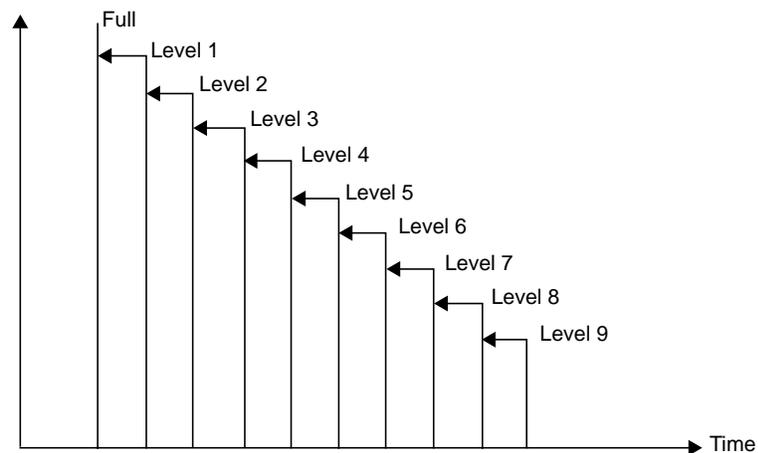


Figure 5-3 Full Backup Followed by Level 1 to Level 9

The advantage of this schedule is that each day's backup is small and completes in a short time. The disadvantage is that recovering from a disk crash requires 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.

Combining Full, Incremental, and Level Backups

Sites with even a few gigabytes of files to back up often choose a monthly schedule based on full, incremental, 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.

This monthly backup schedule minimizes the size of daily backups and also makes it relatively easy to recover in the event of a disk crash. This schedule offers several advantages. First, the level 5 backups simplify recovery. For example, if a disaster strikes on the 24th of the month, all the files needed to recover an entire client system are located on tapes from just five backup sessions:

- incrementals from the 21st, 22nd, and 23rd
- level 5 backup from the 20th
- full backup at the beginning of the month

Second, the incremental backups are relatively small and quick to execute, even for large network environments, and several days of incrementals can fit onto a single tape. This situation further simplifies recovery and also avoids the need to have someone change tapes each day.

Figure 5-4 diagrams level 5 and incremental backups after a full backup.

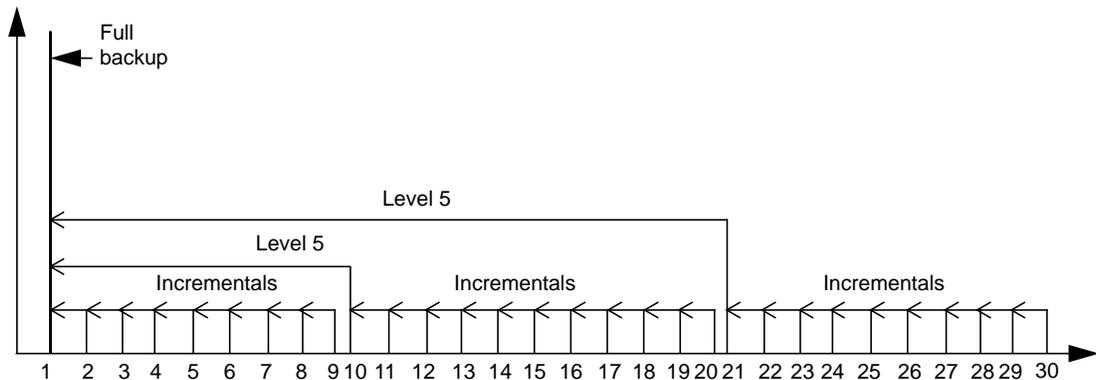


Figure 5-4 Full Backup Followed by Level 5 and Incremental Backups

Backup Time Requirements

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 at night. Although NetWorker is able to back up live filesystems, most administrators want 100% of their network and system capacity ready for users during work hours.

How many files can NetWorker back up in, for example, a four-hour backup window? If your backup server is able to 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), you can 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 must be limited to weekends and holidays, when users are not affected.

A Silicon Graphics system with a 2 GB system disk and an 8 GB logical volume requires at least seven and a half hours for a full backup on an EXABYTE™ 8500 tape drive. For a Silicon Graphics DLT, 10 GB of data takes about 2.8 hours to back up uncompressed onto one tape.

To reduce the amount of time that backups take, follow these recommendations:

- Select a backup server with enough CPU power, memory, and bus bandwidth so that the backup server is not the bottleneck. See the section “Guidelines for Choosing a Configuration” in Chapter 12 for more information on choosing a hardware configuration.
- Leave the NetWorker 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.
- Experiment with compressing files on the client systems to reduce the size of the data that has to be written to tape. 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 the files that are being backed up, a full backup of 8 GB will probably fit onto a single 5 GB backup volume. 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 the ability to skip over specified files during the backup. For example, you could choose to skip over *core* files and *.o* files. The NetWorker *skip* directive provides an easy way to specify that such files be skipped. (See the section “Using Directives” in Chapter 6 for more information.)
- Make liberal use of incremental save levels. These are very efficient, since they take minimal backup media space and run very quickly.
- Add a second backup device to your backup server. For unattended backups, a NetWorker server with two backup devices is worth more than twice as much as a NetWorker server with only one backup device. Often the NetWorker server with two backup devices is more productive than two NetWorker servers with only one device each. With IRIX NetWorker plus the TurboPak option, NetWorker can simultaneously back up to more than one device.
- Using a jukebox and NetWorker Support for Jukebox option software is the most efficient way to complete unattended backups.

Unless you have a jukebox and a NetWorker Support for Jukebox option, you also have to schedule backups based on someone being available to load and unload tapes. Many administrators find that an incremental backup of their network fits onto a single 4-mm or 8-mm or DLT tape, but they must schedule multitape full and level backups for specific nights or weekends when an operator is on duty to load additional tapes. If an operator is not available over a holiday weekend, then 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.

The rest of this section gives details on strategies for optimizing backup time requirements:

- using the compression directive *compressasm*
- staggering the backup schedules
- balancing convenience and security
- using NetWorker preconfigured backup schedules

Using the Compression Directive `compressasm`

Using the directive `compressasm` involves significant **CPU** usage. If you have a large **CPU**, and performance monitors such as `gr_osview` during backup do not indicate that the **CPU** is overused, `compressasm` can compress data before it goes to the server.

`compressasm` typically achieves a 2:1 compression ratio, so network traffic is reduced (your results may vary). There is no harm in using `compressasm` in conjunction with a compressing tape drive, but the drive will probably not achieve much compression on the data at the tape drive.

If you are deciding between `compressasm` and a compressing drive solely to increase the amount of data on a tape, use the compressing drive, because the hardware on the drive can compress faster than NetWorker and it places no load on the **CPU** due to the compression.

Follow these guidelines for compressing data during backup:

- Use `compressasm` to minimize network bandwidth if you have available **CPU** power.
- Use compressing drives, such as the DLT 2700, 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 has no effect, and may even expand the data.
- Do not use `compressasm` if you have a compressing drive and no networked clients.

Staggering the Backup Schedules

Networks with a large number of files can take a very long time to back up completely, and require a lot of loading and unloading of tapes. 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. In NetWorker, you can assign a separate backup schedule to each filesystem. Each filesystem, in essence, is treated as if it were a separate client.

Balancing Convenience and Security

You can 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 it is, or what part of the backup schedule cycle is in effect. The same backup volume might contain full, level 1-9, or incremental backups; to NetWorker, it makes no difference. 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 might lose a day of work. In the worst possible case, if the backup volume with the full backup is destroyed, users could lose all the work done since the last full backup. Therefore, some administrators prefer to remove the backup volume used for a full backup and 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 backup, and the other volumes that contain the most recent level 1-9 and incremental backups.

Using 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 indexes.

By default, NetWorker backs up all of a client's files whenever that client is scheduled to be backed up. The save sets feature allows you to exclude certain files from backup, or to back up some of a client's files on one schedule and other files on a different schedule. This technique is particularly useful for clients with a great deal of data to back up.

NetWorker has several windows relating to save sets:

- create save sets for a client in the Clients window (see “Adding a New Client” and “Scheduling Large Client Filesystems” in Chapter 6)
- set browse and retention policies for save sets in the Policies window (see “Determining Browse and Retention Policies (Index Policies)” later in this chapter)
- view save sets in the Instances window (see “Viewing Save Sets to Determine Resource Usage” in Chapter 8)
- recover and clone save sets in the Save Set Recover and Set Save Clone windows (see Chapter 9, “Recovering and Cloning Save Sets”)

Using NetWorker Preconfigured Backup Schedules

For your convenience, NetWorker is shipped with five preconfigured backup schedules: Default, Full Every Friday, Full on First Friday of Month, Full on First of Month, and Quarterly. If these schedules fit your backup requirements, you can use them “out of the box.” Alternatively, you can delete them and create new ones to accommodate 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.

Note: You cannot 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.

The most efficient way to best 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.

Default

NetWorker provides a preconfigured backup schedule named “Default,” which you are not allowed to 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 useful for a small to medium-sized network where the scheduled backups fit onto one backup volume or where autochangers are used.

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. Figure 5-5 shows the default schedule.

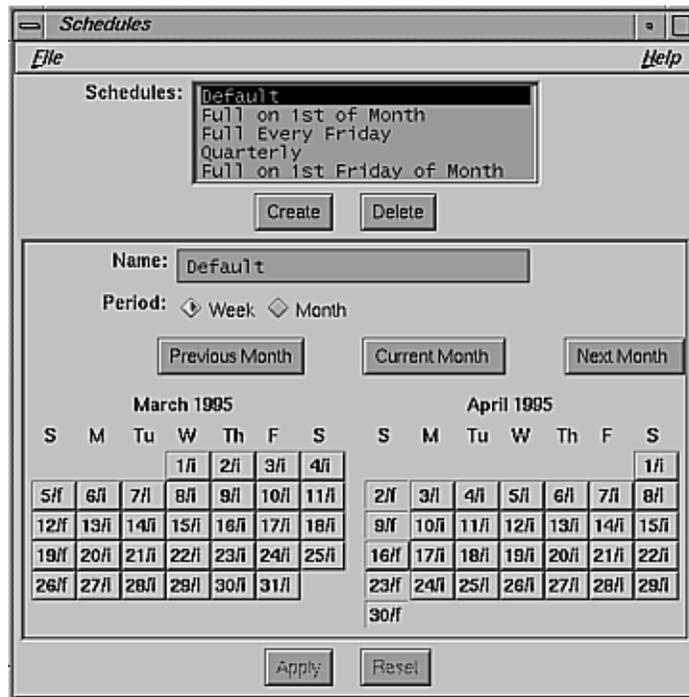


Figure 5-5 Default Schedule

Each time you use the Schedules window to create a new weekly backup schedule, this preconfigured schedule appears in the calendar as your starting point.

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 First 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. Figure 5-6 shows this schedule.

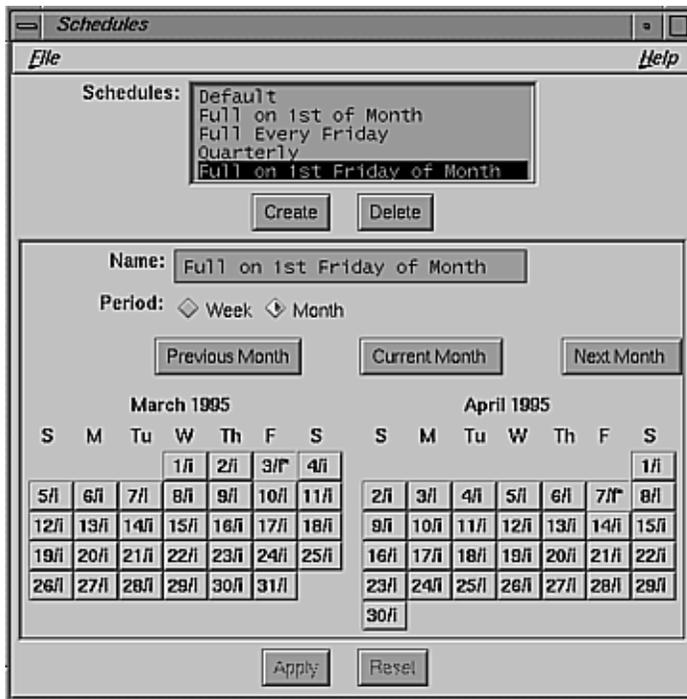


Figure 5-6 Full on First Friday of Month Schedule

If you use this schedule, 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 are at risk of not being 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 (Friday, in this example), you must set the overrides in each month. (Notice the “f*” in the first Friday of each month.) This requirement is because the first weekday (Monday through Friday) in a month can fall on any calendar day from 1 to 7.

Note: When the year ends, you must add the overrides for the next year. In other words, the overrides do not carry over from one year to the next. Preconfigured schedules, however, maintain the overrides for years into the future.

Full on First of Month

This monthly schedule completes a full backup on the first calendar day of the month. Many sites prefer to begin each month with a full backup on the first 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. Figure 5-7 shows this schedule.

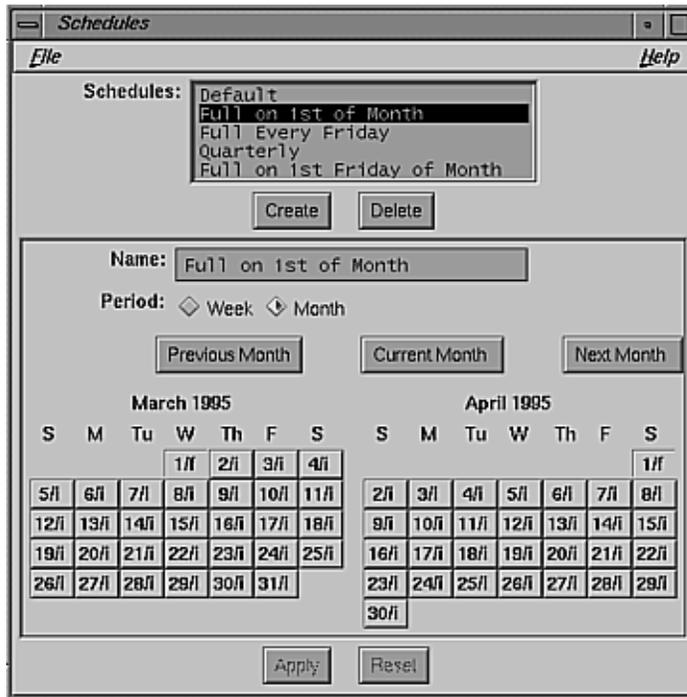


Figure 5-7 Full on First of Month Schedule

Whenever you create a monthly schedule, this schedule is your starting point.

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. Figure 5-8 shows this schedule.

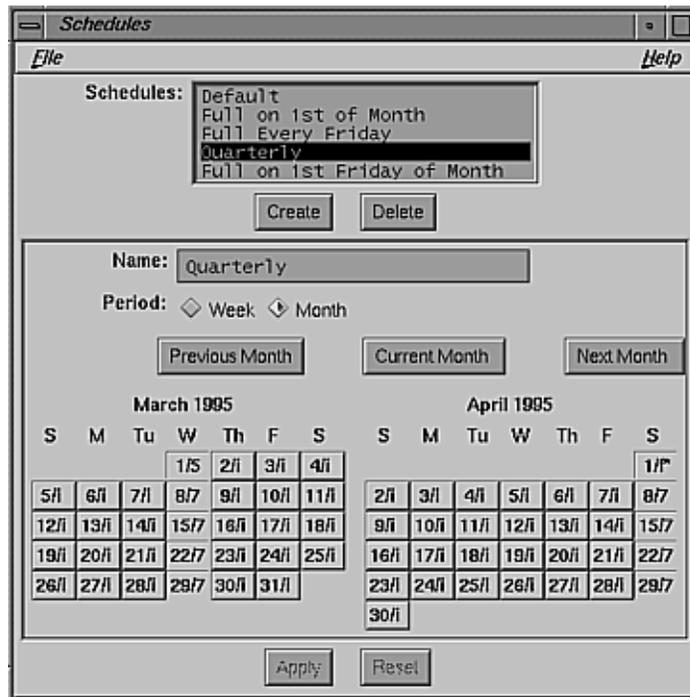


Figure 5-8 Quarterly Schedule

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, segregate the backup volume with the full backup and store it in a safe place. The monthly level 5 backups should also be segregated onto their own backup volumes. On the other days of the week, leave one backup volume in the server, so that the level 7 and incremental backups are stored on it. However, if a week's worth of backups is on one backup volume, and it is destroyed the same day the disk crashed, you could not recover the changes that took place that week. Therefore, it would be better to change backup volumes every day, putting each day's backup on

its own volume. If a daily incremental backup is destroyed and you need to recover from a disk crash, you can recover all but one day's work.

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.

Determining Browse and Retention Policies (Index Policies)

When Worker 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 stores information about files backed up by NetWorker, and the media index maps the saved files to the backup volumes. NetWorker maintains one file index per client and one media index per NetWorker server. NetWorker uses the indexes as databases to locate the files that are marked for recovery.

Each entry in the file index typically includes this 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 changes with each backup, as entries for the newly backed up files are inserted. As long as an entry for a file remains in the file index and the backup volume is not damaged, you can recover the file using the NetWorker Recover window.

The media index maps each file to the backup volume or volumes where it is stored. NetWorker uses the media index to tell you or the jukebox which backup volume to mount during a recover. 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 entries it contains. NetWorker determines which volume to mount for recovering a file by mapping the saved files to their backup volumes.

In the NetWorker Policies, Indexes, and Volume Management windows, you create policies for automatic index management, monitor the contents of the

indexes, select entries for removal, and mark backup volumes as recyclable. NetWorker uses *browse and retention policies* to manage and reduce the size of the online indexes:

- The browse policy determines how long entries for your files remain in the online file index and thus browsable in the Recover window. Entries older than the browse policy are automatically removed from the online file index.
- The retention policy determines how long entries are retained in the media index and thus are recoverable. Entries older than the retention policy are marked as recyclable in the media index.

NetWorker performs four actions on both kinds of indexes: *inserting entries, browsing, 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 is looking through the index for information about your saved files or the contents of your backup volumes; thus, it neither increases nor decreases the size of an index.
 - The file index is browsed when you use the NetWorker Recover window to locate a file, or when you use the Indexes window to browse the save sets that contain the files you see in the Recover window.
 - The media index is browsed when you use the Volume Management window to view the save sets on the backup volume.
- Removing entries frees space in the index. NetWorker uses the free space to insert new entries. The browse and retention policies determine when entries should be removed from the index. You can also remove entries manually by selecting “Remove oldest cycle” in the Indexes window or by selecting “Remove” from the Volume pull-down menu in the Volume Management window.
- Reclaiming space returns empty space to the filesystem. The empty space is created when entries are removed from the index; you remove it by using the *Reclaim space* button in the Indexes window.

Note: “Monitoring and Managing Index Disk Space Usage” in Chapter 8 gives detailed instructions for removing index entries and reclaiming space.

This section explains NetWorker browse and retention policies and the trade-off between providing faster, easier recovery for users and conserving disk space. Topics include:

- how browse policies work
- reclaiming disk space
- recovering files removed from the index
- how media retention policies work
- NetWorker preconfigured browse and media retention policies

How Browse Policies Work

One popular NetWorker feature is its index, which allows a user to browse the 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's online index (about 200 bytes per version). Since disk space is limited, you need to establish a policy of how far back in time you 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 you select, like the backup schedule, can be different for each client.

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 does 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 might expect.

In Figure 5-9, the browse policy is set to one week, which happens to equal one complete backup cycle.

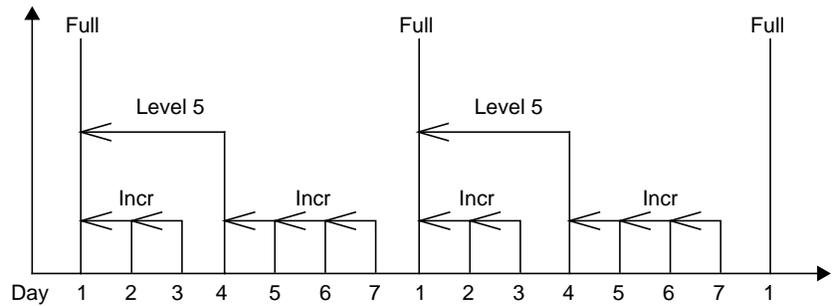


Figure 5-9 Backups for a Browse Policy of One Week

NetWorker does 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 is not removed from the online index in exactly one week, however, because there are unexpired incrementals and a level 5 backup that depend on the full backup. Each incremental backup is removed from the online index one week after it was completed. The level 5 backup is removed one week after the last incremental that depends on it is removed, and then the full backup is 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.

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 that is long enough to cover most recovery requests.

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. Because reclaiming this space requires time, processing power, and swap space, having this process constantly occurring on your backup server would be inefficient. Instead, NetWorker first reuses newly freed space to store information about new files that are backed up. When less than 50% of file index space for a client is being used by files that have not reached the end of their browse period, NetWorker automatically invokes a process that returns the space to your system.

You can also reclaim disk space at any time by using the Reclaim space button in the Indexes window. See “Monitoring and Managing Index Disk Space Usage” in Chapter 8 for instructions on space-saving operations.

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. The recover process is not as convenient as when the entries are still in the online index, however.

If you do not want to rebuild the index, and the save sets you need are still in the media index, and 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*. See “Recovering Files Removed From the Index” in Chapter 8 for instructions.

Tip: Recovery is considerably easier if the file information is still in NetWorker’s online index. Set a browse policy that is long enough to cover most recovery requests.

How Media Retention Policies Work

Your need to conserve disk space may lead you to establish a short browse period. The NetWorker media retention policy complements the browse policy by letting you specify a longer period of time during which files can still be recovered, although with more difficulty. The retention policy is also used by NetWorker to automatically recycle backup volumes.

As mentioned earlier in this chapter, NetWorker maintains a file index for each client system and a much smaller media index that tracks which save sets are 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.

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 Recover window up until the browse policy time is reached and the file information is removed from the file index. Then you can use the more tedious process described to recover your files until the retention policy time is reached and the backup volume is recycled.

NetWorker Preconfigured Browse and Media Retention Policies

NetWorker is shipped with five preconfigured browse and retention policies: Week, Month, Quarter, Year, and Decade. 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 that NetWorker retains entries for every file that is 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 can browse all the file entries for backed-up files dating back to a month. Each month the oldest entries for the client's files are automatically removed from the server's file index. However, the backup tapes that contain the data for the files are still tracked by NetWorker in the media index.

Week Policy

This policy maintains the file index entries or the media index entries for one week. If you use this as a browse policy, the users can only view and mark for recovery files that go back in time for a week. As a browse policy, it is useful when you have a limited amount of disk space and users do not expect to be able to recover versions of their data that are older than a week.

As a retention policy, Week means that your backup volumes turn over quickly, and NetWorker recycles 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 Policy

This browse policy allows users to view and recover versions of files dating back at least a month. The Recover window displays versions for 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 month plus one full backup cycle.

Quarter Policy

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 Policy

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 for at least three quarters, Year is a good browse and retention policy. Realize, however, that NetWorker requires more disk space to maintain all the information online.

Decade Policy

This policy retains the entries in the server's indexes for ten years. It is useful for organizations that are required to keep records for very long periods.

Your NetWorker server requires 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 many gigabytes of disk space. In this case, 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 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.

Note: The limit for each client's file index is 2 GB.

Determining Autochanger Policies

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.

This section discusses

- jukebox capacity for one backup cycle
- jukebox capacity for more than one backup cycle
- choosing a jukebox

Jukebox Capacity for One Backup Cycle

A jukebox is most useful if it has at least enough capacity to complete one entire backup cycle without intervention. Such a jukebox allows backups to run while you are out ill, on vacation, or busy with a user emergency. It also 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 off site 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, the file should be recovered very quickly.

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 your 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 directives to skip files during the backup, for example, *core* files. See “Using Directives” in Chapter 6 for information.
- Shorten the length of the backup cycle.

Although your jukebox may only have enough capacity for 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 prompts 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.

Jukebox Capacity for More Than One Backup Cycle

If the jukebox has just enough capacity for a single backup cycle, you must reload tapes at the end of each cycle. If it has 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. To continue backups virtually indefinitely, the jukebox automatically recycles tapes that contain save sets that have passed their browse and retention times.

Suppose you have 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 EXB-10i jukebox. Each of these combinations of browse and retention times allows the jukebox to operate without intervention for an extended period:

- 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 are kept small. With a browse and retention policy of two weeks, your indexes are larger but users 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 does allow 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, $4 * 12 \text{ GB} = 48 \text{ GB}$ fits 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.

Now 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 allows the jukebox to run unattended on an ongoing basis.

If you want to keep files online in the jukebox longer, you can use the methods listed earlier to reduce the size of the backup cycle. Alternately, you can stretch out the backup cycle. For example, you can perform full backups every other week rather than every week. This system 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. Table 5-1 illustrates that one complete backup cycle will be about 64 GB in size.

Table 5-1 Choosing a Jukebox

Level	Size	Frequency	Total
Full	25 GB	* 1 time /month	25 GB
Level 5	2.5 GB	* 2 times/ month	5 GB

Table 5-1 (continued) Choosing a Jukebox

Level	Size	Frequency	Total
Incremental	1.25 GB	* 27 times/month	34 GB
			64 GB

To determine the size of 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. With some effort you can recover for them files that they had any time during the past six months. Thus you will need six months times 64 GB, which yields 384 GB of capacity.

In practice you will need a little extra jukebox capacity, because there will be a number of “unavailable” volumes as 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 grow, a rule of thumb is that you should purchase a jukebox with about 50% more capacity than your current requirement.

Configuring a NetWorker Server

This chapter covers the system administration tasks for configuring the NetWorker server. These configuration tasks include

- using NetWorker preconfigured settings
- configuring backup groups
- setting up backup schedules
- creating browse and retention policies
- using directives
- configuring a server's clients
- fine-tuning server settings

Tip: Set up the groups, backup schedules, policies, and directives before you configure any clients. After you have created these custom configurations, apply them to specific clients in the Clients window. You can also edit the configurations that are in use by existing clients.

Using NetWorker Preconfigured Settings

NetWorker is shipped with easy-to-use preconfigured settings, which are choices NetWorker has established for backing up the clients on your network. You can use these preconfigured selections, change them, or create your own configurations. Using these preconfigured settings, you can start backing up your data as soon as you install NetWorker.

After it is installed on a server, NetWorker backs up the NetWorker server as a client of itself using these preconfigured settings. If you are backing up other client systems to your server, all you have to do (after installing NetWorker on the clients) is add the clients by using the Clients window.

(Open the Clients window by selecting the “Clients” command from the Customize pulldown menu.) NetWorker can then back up your clients using the preconfigured settings.

Your NetWorker server and newly created clients have these preconfigured settings, unless you select otherwise:

- | | |
|------------------|---|
| Schedule | The client uses the Default schedule of a full backup every Sunday and an incremental backup on the other days of the week. The Default schedule is described in detail in the section “Using NetWorker Preconfigured Backup Schedules” in Chapter 5. |
| Browse policy | The browse policy for entries in the client’s online file index is one month. This means that entries for the client’s files remain in the file index and can be browsed and recovered in the Recover window for one month before they are automatically removed. For information about the browse policy, see “Determining Browse and Retention Policies (Index Policies)” in Chapter 5. |
| Retention policy | The retention policy for the client’s media index is one year. This means that information concerning the names of the backup volumes containing save sets generated by clients is retained for one year. Files can be recovered for up to one year, providing the backup volumes are not overwritten with new backups. |
| Directive | The client uses the Default directive which skips <i>core</i> and <i>.o</i> files and backs up only the relevant information in the NFS [®] based swap files. It also backs up your mail files without ever making any changes to them. |
| Group | The client is in the Default backup group which has a start time of 3:33 a.m. To enable the scheduled backup to take place at 3:33 a.m. every day, change the Autostart feature to Enabled in the Groups window. |
| Save set | The preconfigured selection is All. All of the local client filesystems are backed up. |
| Recover access | This field is blank, meaning that only the client itself can recover its backed-up files. |

The NetWorker server is listed in the Clients scrolling list prior to creating any clients because the server is automatically created as a client of itself when NetWorker is installed. To back up the NetWorker server as a client of itself using the preconfigured settings, all you have to do is enable the Autostart feature of the Default group in the Groups window.

Configuring Backup Groups

NetWorker allows you to bundle clients into backup *groups*. This feature allows you to easily schedule clients to be backed up at different times. A client can be in one or many backup groups. Backup groups allow you to stagger the backup times of groups of clients to avoid overloading the server.

NetWorker is shipped with a preconfigured backup group named Default. When a client is added, it automatically belongs to the Default backup group.

The automatic backup of the Default group starts at 3:33 a.m. If this backup start time suits your needs, you do not need to change it. You must enable the Default group before automatic backups take place. You can make changes to the Default group, but you cannot delete it.

If you have a large number of clients, or if you prefer to start backing up certain systems at different times, you can easily do so. For example, you may want to start the backup of the engineering department systems 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.

After creating a new backup group in the Groups window, you assign clients to the group in the Clients window. Because it is more expeditious to create backup schedules, policies, and directives before configuring any clients, the Clients window is discussed later in this chapter. This section explains

- creating a backup group
- changing a backup time
- deleting a backup group

Creating a Backup Group

Create a backup group and set the group backup start time in the Groups window. Select “Groups” from the Customize pulldown menu (shown in Figure 6-1) to open the Groups window.



Figure 6-1 Customize Pulldown Menu

The Groups window appears, as shown in Figure 6-2.

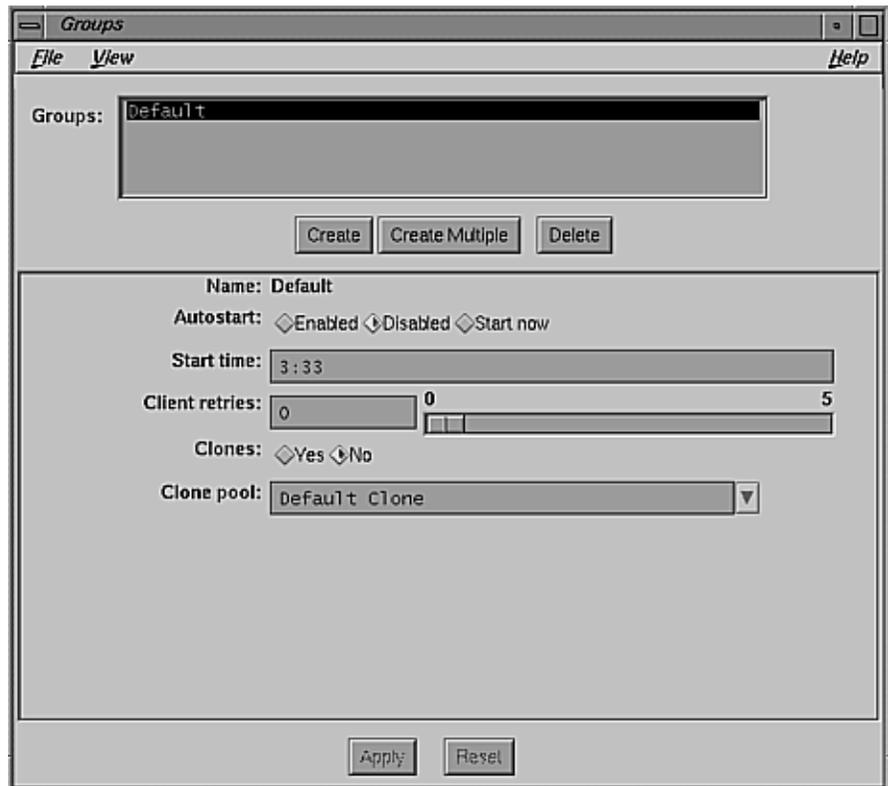


Figure 6-2 Groups Window

Use features of this window as follows:

Groups scrolling list

View names of all backup groups known to the NetWorker server. Selecting a group name in the scrolling list displays its name in the Name field and its backup start time in the Start time field.

Name field

View 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** Disable, enable, or immediately initiate a network-wide backup. To allow NetWorker to back up your clients in the Default group you must select Enabled from the Autostart choices.
- Enabled: Allows the network-wide backup of the highlighted group to start at the scheduled time
 - Disabled: Prevents network-wide backup of the group from starting.
 - Start now: Initiates the network-wide backup of the group immediately, instead of at the scheduled start time.
- Start time field** View the time NetWorker starts the backup of the clients in the group currently selected in the Groups scrolling list. The Default group is preconfigured to back up at 3:33 am.
- Client retries field**
Use the sliding bar or enter the number of times NetWorker should attempt to back up a failed client during a scheduled backup. 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** Click *Yes* to automatically clone the data being backed up for a group. Click *No* if you do not want to clone the data being backed up in a group.
- Cloning a backup during execution means that you are writing simultaneously to two different backup volumes, the normal one and the clone. The data written to each volume for a particular save set is identical.
- Note:** To clone data you need NetWorker TurboPak that provides support for concurrent devices. To determine if your site has this option, see the release notes.
- Clone pool field** Enter the name of a pool or use the arrow button to select a clone pool from the scrolling list to which the cloned data will be backed up.

To create a new group and assign a backup start time to the group, follow these steps:

1. Open the Groups window in the Customize menu.
2. Click the *Create* button. Notice that the Name field displays an editable line.
3. Enter a new group name into the Name field. Choose a name that describes the systems in the group, either by department or schedule.
4. Select “Enabled” from the Autostart choices.
5. Delete the time in the Start time field and enter a new time, using the 24-hour clock convention. (For example, 10 p.m. is entered as “22:00.”)

Note: Normally, you should schedule the start time for each group far enough apart so that one group can complete backing up before the next group starts. This action is not required if you are using a short backup window and many devices.

6. Use the Client retries sliding bar to set a value for the number of times NetWorker should attempt backing up a failed client.

NetWorker normally tries 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 is down, for example), it does not try again unless you configure it to do so.

Select the sliding bar with the mouse and move it to the right or left. The number of times NetWorker tries again to back up the failed client appears in the Client retries field.

7. Click *No* or *Yes* in the Clones field. If you click *Yes*, select or add the name of the clone pool in the Clone pool field.
8. Click the *Apply* button to apply the changes. The new group name appears in the Groups scrolling list.

If you make a mistake, clear your changes and start over by clicking the *Reset* button.

Note: The error message `user user_name needs to be on administrator's list` means that you do not have permission to make configuration changes. See “Adding or Changing Administrators” later in this chapter for more information.

When you are finished creating a backup group and setting the start time for it, the group appears as a checkbox in the Clients window. You then assign clients to the group using the Clients window. See the section “Configuring a Server’s Clients” later in this chapter for more information.

Changing a Backup Time

To change the backup start time for a group, follow these steps:

1. Open the Groups window in the Customize menu.
2. Select the group from the Groups scrolling list for which you wish to change the backup time.
3. Delete the time in the Start time field and enter a new time, using the 24-hour clock convention. (For example, 11 p.m. is entered as “23:00.”)
4. Click the *Apply* button to apply the changes.

If you make a mistake, clear your changes and start over by clicking the *Reset* button.

Deleting a Backup Group

NetWorker does not let you delete a backup group if any clients are assigned to it. Before you try to delete a group, make sure no clients are assigned to the group in the Clients window.

To delete a group, follow these steps:

1. Open the Group Control window in the Server menu to make sure that no clients are assigned to the group you want to delete.

If clients are assigned to the group, use the Clients window to delete the all clients, or to change each client’s group affiliation.

2. Open the Groups window in the Customize menu.
3. Select the name of the group you want to delete from the Groups scrolling list.
4. Click the *Delete* button. An *Ok to delete?* message appears.

5. Select *OK* to delete the group. If no clients belong to the group, the group name disappears from the Groups scrolling list.

If any clients belong to the group you are trying to delete, you see the message *The group **group** is still used by some clients. Select *OK* to dismiss the message. To delete the group, you must first remove the clients from the group using the Clients window, then repeat this procedure.*

Setting Up the Backup Schedules

NetWorker automatically backs up all the clients in a group, starting at the time you specify in the Groups window. The level and schedule of the backup is determined by the individual backup schedule of the client.

Note: For background information on schedules, see “Determining Backup Schedules” in Chapter 5.

You create backup schedules for the clients in the Schedules window. To open the Schedules window, select “Schedules” from the Customize pulldown menu (shown in Figure 6-1 earlier in this chapter). The Schedules window appears, as shown in Figure 6-3.

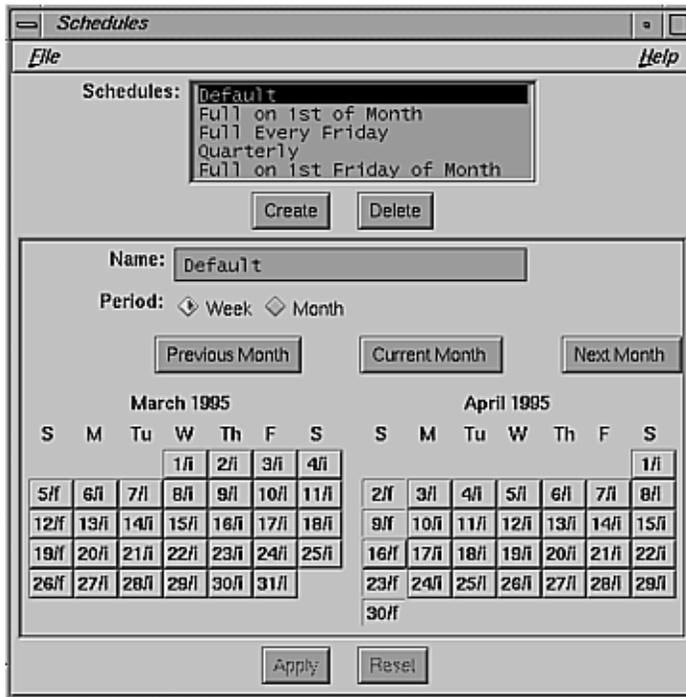


Figure 6-3 Schedules Window

The features of this window have these functions:

Schedules scrolling list

Displays the names of the schedules known to the NetWorker server, including the five preconfigured backup schedules and any schedules that you create. The five preconfigured schedules are described in “Using NetWorker Preconfigured Backup Schedules” in Chapter 5.

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 Lets you set the calendars to a weekly or monthly period. If you select *Week*, you set the schedule for one week, and that same schedule appears in all the calendar weeks. If you select *Month* and set the schedule for one month, the same schedule applies to all the calendar months.

- *Previous Month* buttons
Click to view the previous month's backup schedule in the calendars. Each time you click the *Previous Month* button, NetWorker displays two months further back in the calendar year.
- *Current Month* button
Click to view this month's and next month's backup schedules in the calendars.
- *Next Month* button
Click to view the upcoming month's backup schedules in the calendars. Each time you click the *Next Month* button, NetWorker displays the next two months further ahead in the calendar year.

Calendars Display the day of the month and the backup level that is scheduled for that day. An asterisk (*) next to a backup level means you have forced a different level (an override) on that day. For example, suppose 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.

The date on each calendar day is followed by a forward slash (/) and either a number or character signifying the level and kind of backup for that day:

- f: full backup
- 1-9: level one through nine backups
- i: incremental backup
- *: skip scheduled backup

Before using this window to set up a schedule, decide what schedule best fits the needs of your network, based on the backup levels that NetWorker supports. See "Determining Backup Schedules" in Chapter 5.

Note: You can make changes to an existing schedule, but you cannot change its name. If you want to change the name of a schedule, you must delete it and recreate it, giving it a new name.

Creating a New Backup Schedule

To create a backup schedule, follow these steps:

1. Choose “Schedules” from the Customize pulldown menu.
2. In the Schedules window, click the *Create* button to create a schedule.
3. Enter a name for the new schedule into 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 pull down the Backup level menu. Set the backup level for each day (or week, or month):
 - Press the **F** key for a full backup on a calendar day.
 - Press the **I** key for an incremental backup on a calendar day.
 - Press the **S** key to skip a backup on a calendar day.
 - Press any number key between **1** and **9** for a level backup.
 - Select the “Overrides” command from the Backup level menu to bring up the Overrides menu. Then press the corresponding speed key for an override on a calendar day.
 - To remove an existing override, pull down the Overrides menu and press **R**.
6. When you are done, click the *Apply* button to create the new schedule.

You can set up as many NetWorker backup schedules as you want, as long as each has a unique name. To assign a backup schedule to a client, use the Clients window, as explained in “Configuring a Server’s Clients” later in this chapter.

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 makes recovery from a disaster easier. Also, someone can change backup volumes and monitor the server on Mondays. The other days of the week, incremental backups protect the daily changes. You are going to name the backup schedule “Monday Full” to remind yourself that this schedule does a full backup every week, on Mondays.

Follow these steps:

1. Open the Schedules window.
2. Click the *Create* button.
3. Enter **Monday Full** into the Name field.
4. Select the *Week* button, if it is not already selected.
5. In the current month calendar, select any Sunday and pull down the Backup Level menu, as shown in Figure 6-4.

March 1995

S	M	Tu	W	Th	F	S
			1/i	2/i	3/i	4/i
5/i	6/i	7/i	8/i	9/i	10/i	11/i
Overrides ▶			15/i	16/i*	17/i	18/i
full			22/i	23/i	24/i	25/i
Incr			29/i	30/i	31/i	
1						
2						
3						
4						
5						
6						
7						
8						
9						
Skip						

Figure 6-4 Selecting a Backup Level

6. Select “incr” from the menu.

Notice all the Sundays of the calendar change to a backup level “i,” which means incremental backup.

Now set the calendar to do a full backup every Monday:

1. In the current month calendar, select any Monday, and pull down the Backup level menu.
2. Select “full” from the menu.
3. Click *Apply* to save the schedule.

Note: The error message `user user_name needs to be on administrator's list` means that you do not have permission to make configuration changes. See the section “Adding or Changing Administrators” later in this chapter for more information.

Notice that all the Mondays in both calendars now display “f” next to the date, indicating that a full backup takes place instead of the incremental backup.

Scroll through all the months, and notice that the schedule you set up for the current month is maintained throughout the entire calendar.

Example 2: A Monthly Schedule

This monthly schedule is for the clients with files that do not change very much, and therefore need a full backup only once a month. A full backup takes place on the first day of the month. In the middle of the month, a level 5 backup backs up all the files that have changed since the full backup. All the other days of the month, an incremental backup protects any daily changes.

You name this schedule “Monthly Fulls” to remind you that this schedule does a full backup only once a month. You are going to skip the level 5 backup with an override on January 16, because it is a company holiday, and force the backup on January 17 instead.

Follow these steps:

1. Open the Schedules window.
2. Click the *Create* button.
3. Enter `Monthly Fulls` into the Name field.
4. Click the *Monthly* button, if it is not already selected.

Notice that 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. Likewise, if you set up a weekly schedule, it applies for all calendar weeks. This system allows you to plan all the weeks or months in one session, without having to repeat the process every week or month.

Now set the month for a level 5 backup on the sixteenth day:

1. In the current month, select the sixteenth day (16/i).
2. Select “5” from the Backup level menu.

Scroll through the months, and notice that 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.

Now you are going to override the scheduled level 5 backup on January 16 with a level “skip,” to skip the backup for that day, and schedule the level 5 backup on January 17. Follow these steps:

1. Click the *Previous Month* or *Next Month* buttons to display the January calendar.
2. In the January calendar, select the sixteenth day.
3. Pull down the Backup level menu, shown in Figure 6-4, and select “Overrides.”
4. Slide the cursor to the right, to display the Overrides menu, as shown in Figure 6-5.

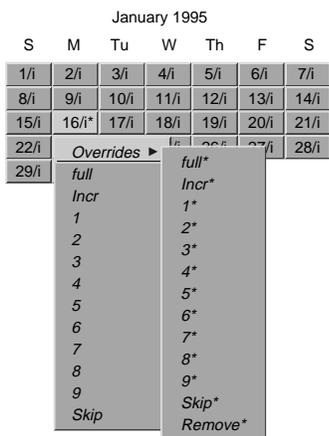


Figure 6-5 Overrides Menu

5. Select “skip*” from the Overrides menu.
6. Select the seventeenth day in the January calendar.
7. Select “5*” from the Overrides menu.
8. Click *Apply* to save the schedule.

Notice that the date on January 16 displays an “s*,” indicating you have forced a “skip backup” instruction for that day. The date on January 17 displays a “5*” to indicate a level 5 backup takes place instead of the incremental.

Note: Any override must be scheduled explicitly each time. In other words, an override does not automatically repeat in a monthly or weekly period.

Creating Browse and Retention Policies

This section describes how to create browse and retention policies that automatically manage 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 single policy (preconfigured or one you create yourself) can be used as either a browse or retention policy.

Note: For background information on browse and retention policies, see “Determining Browse and Retention Policies (Index Policies)” in Chapter 5. You can also manage the indexes manually using the Indexes and Volume Management windows. Manual index management and index management concepts are described in detail in “Determining Browse and Retention Policies (Index Policies)” in Chapter 5; procedures are explained in “Monitoring and Managing Index Disk Space Usage” in Chapter 8.

If you create new policies for automatic index management using the Policies window, the policies appear as selections for both the browse and retention policies in the Clients window. Every client requires both a Browse and retention policy. When you select browse and retention policies in the Clients window, you are telling NetWorker the minimum time period for data to remain browsable and recoverable for that client.

Note: NetWorker includes five preconfigured policies: Week, Month, Quarter, Year, and Decade. The section “NetWorker Preconfigured Browse and Media Retention Policies” in Chapter 5 describes these policies.

To create policies for automatic index management, open the Policies window by selecting “Policies” from the Customize pulldown menu (shown in Figure 6-1). The Policies window is shown in Figure 6-6.

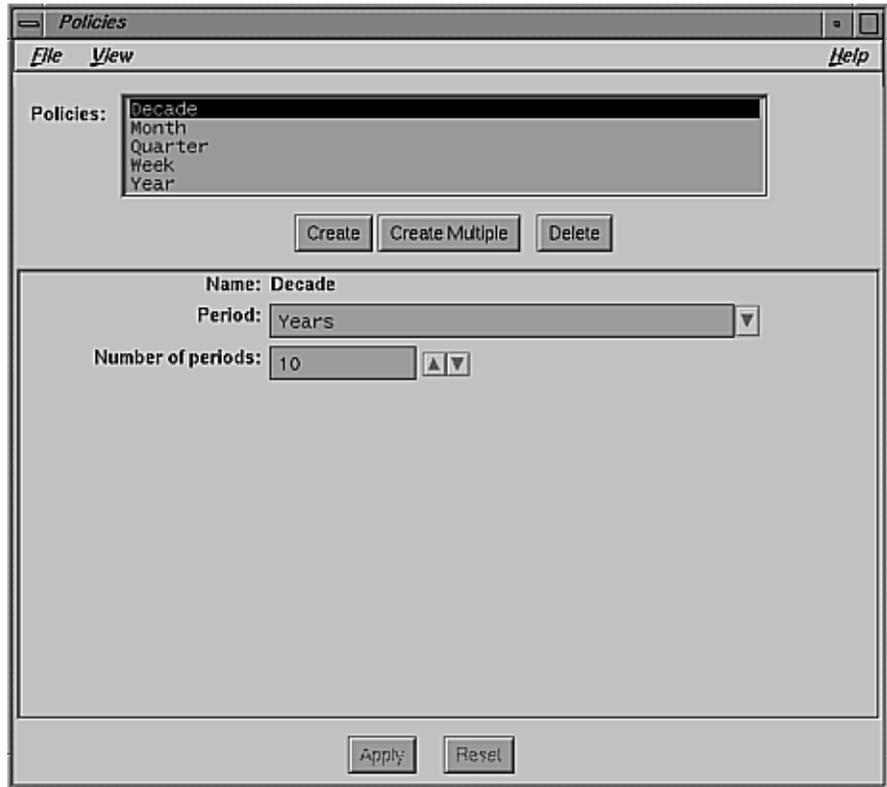


Figure 6-6 Policies Window

The features of this window have these functions:

Policies scrolling list

Lists all 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. If you have not created any new policies, the scrolling list displays five preconfigured policies: Decade, Month, Quarter, Week, Year.

Name field

Displays the name of the currently selected policy. Use the Name field to enter the name of a new policy 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, in which you choose a time period.

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, if you choose Years for the Period, and 10 for the Number of periods, and you apply it as a browse policy in the Clients window, NetWorker automatically removes index entries from the online file index that have been in the index for more than ten years.

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 counter or entering a number in the field.
5. Click the *Apply* button to finish the process.

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 reset the window to where it was before the last *Apply*.

Enter any name for the policy that you want. It is useful to describe the time period to help 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.

Follow these steps to delete a policy:

1. Highlight the desired policy in the scrolling list.
2. Click the *Delete* button. A warning message appears asking if it is OK to delete the selected policy.
3. Click *Yes* if you wish to delete the policy.

Tip: You are not allowed to delete either policies that are currently in use or the preconfigured policies.

Follow these steps to change a policy:

1. Highlight the desired policy in the scrolling list.
2. Select the Period and Number of periods for the policy. The name cannot be changed.
3. Click the *Apply* button to change the policy.

Note: A policy cannot be changed so that the client browse policy exceeds its retention policy: an entry must be removed from the file index before it can be marked recyclable or removed from the media index.

Using Directives

In a typical directory tree, many files do not need to be backed up: for example, *core* files that are generated when a process terminates abnormally, or object files (“*.o*” files), which are easily regenerated from source code. You probably do not want to waste backup volumes or network bandwidth backing up these files. Also, if your network is slow or busy, you can compress client files to use fewer backup volumes and less network bandwidth backing up its files.

NetWorker can process individual client files in very efficient methods through *directives*. A directive contains instructions that assist the backup process. Sometimes the directives require execution of special commands called Application Specific Modules (ASMs).

Note: NetWorker works fine if you use the preconfigured directives without changing them. If you are an advanced user, and want to add custom directives, please read the *nsr* reference page as well as this section.

You have lots of flexibility in using directives. For example, you can apply backup instructions to specific directories in the filesystem. You can use pattern matching to apply the instructions to specific files in the directory tree. For example, using “skip: *.o” skips the backup on any file that has a suffix of “.o” in the directory to which the “skip: *.o” directive applies.

NetWorker is shipped with four sets of preconfigured directives:

DOS standard directives

Use when you are backing up DOS clients.

NetWare standard directives

Use when you are backing up NetWare clients.

Unix standard directives

Use for most of your backups and when you do not need one of the other specialized directives.

Unix with compression directives

Use when you want to compress your backup data. Compressing client files saves you media space and network bandwidth, but it takes more time and CPU cycles on the client. Overall, the entire network may back up faster if all the clients are compressing their files, and the parallelism is set appropriately.

Caution: If your backup device compresses data, do not use both NetWorker's *compressasm* directive and the device's data compression feature at the same time.

Each set of directives includes the most important and most useful backup instructions. For example, they apply the directive "+skip: core" to the *root* directory (/), thus skipping the backup of all *core* files throughout the directory tree. They contain a *swapasm* directive to back up the relevant information about all NFS® based and local swap files, but not the data in them. They also contain a *mailasm* directive to ensure that your mail files are backed up, yet not marked as read.

To display the Directives window, select "Directives" from the Customize pulldown menu. The Directives window is shown in Figure 6-7.

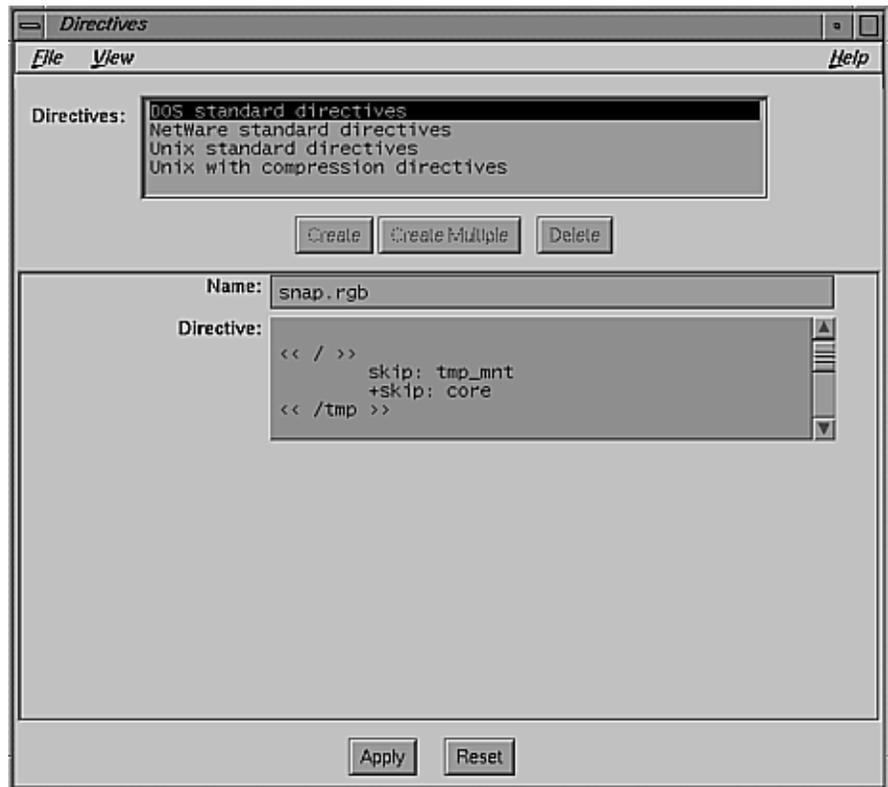


Figure 6-7 Directives Window

The features of this window have these functions:

Directives scrolling list

Lists directives known to the NetWorker server

Name field

Displays the name of the currently selected directive. Use the Name field to enter the name of a new directive, after you click the *Create* button.

Directive field

Displays syntax for the currently selected set of directives. Each directive is preceded by the directory to which it applies (for example, << /tmp >>).

To create a set of directives or add a single line of instruction, follow these steps:

1. Click the *Create* button. The Directive scrolling list changes to allow you to change the name and contents of the directive.
2. Type your instructions directly into the Directive field.

The correct syntax for directives is explained in more detail in the *nsr* reference page.

To change a directive, press the *Change* button and edit the window.

3. After you add a directive or set of directives, click the *Apply* button to apply the change.

For more information about directives, and the algorithm used to apply directives when both server and clients have directives, see the *nsr* reference page.

Configuring a Server's Clients

Before a client can communicate with the NetWorker server, it must have the NetWorker client software installed and the client must recognize the server. See “Installing NetWorker Software on Clients” in Chapter 2 for instructions. After the software has been successfully installed on the client and the client recognizes the server, you can configure the NetWorker server for that client.

NetWorker is shipped with easy-to-use preconfigured settings described in the section “Using NetWorker Preconfigured Settings” at the beginning of this chapter. When you configure a client manually, you choose:

- NetWorker server, if you have more than one
- backup schedule
- browse policy and a retention policy
- directive
- backup group (none, one, or many)
- client files NetWorker should back up, if not all of them

- other systems, if any, that are allowed to recover the client's files

The Clients pulldown menu, shown in Figure 6-8, contains commands for adding, configuring, and monitoring NetWorker clients and their associated save sets.

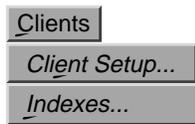


Figure 6-8 Clients Menu

Use the Clients window, shown in Figure 6-9, to apply the configurations (which include schedules, browse and retention policies, directives, and groups) directly to each client. The Clients window also provides you with the flexibility of scheduling large client filesystems individually and the ability to restrict access to client files. Figure 6-9 shows this window with NetWorker preconfigured settings.

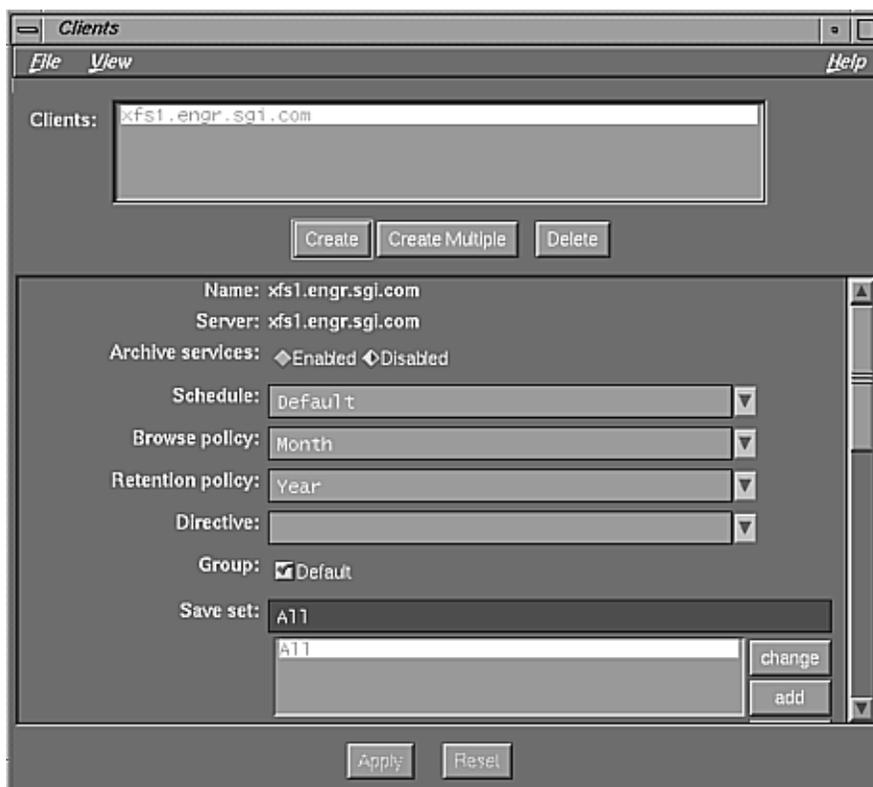


Figure 6-9 Clients Window

Note: When you install NetWorker, the server automatically becomes a client of itself. NetWorker automatically adds the server's hostname to the Clients window. You simply need to enable the Autostart feature in the Groups window for NetWorker to back up the server according to the preconfigured settings.

The features of this window have these functions:

Clients scrolling list

Lists all the NetWorker clients configured for the NetWorker server. The NetWorker server is automatically listed in the Clients scrolling list before any clients are created.

Name field Displays the hostname of the currently selected client.

Server field Displays the hostname of the currently selected NetWorker server.

Archive services choices

If you have purchased the NetWorker Archive product, you can enable or disable archiving. All NetWorker backup clients automatically become archive enabled after the archive application is installed.

Schedule field Displays all the schedules known to the NetWorker server, including the five preconfigured schedules shipped with NetWorker. Any other schedule you create using the Schedules window also appears here as choice.

Browse policy field

Displays the chosen schedule for the selected client. The preconfigured selection is the Default schedule, which performs a full backup every Sunday and an incremental backup every other day of the week. Use the arrow button to access the Schedule scrolling list. The list displays the preconfigured schedules and any new schedules created in the Schedules window.

Retention policy field

Displays the chosen retention policy for the selected client. The retention policy determines how long the save sets on a volume are recoverable before the volume is available for recycling. The preconfigured selection for a client is the Year retention policy: 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. Use the arrow button to access the Retention policy scrolling list that contains the preconfigured policies and any new policies previously created in the Policies window.

Note: NetWorker does not allow a client's Browse policy to exceed the Retention policy because if a volume is recycled, the online file index entries are no longer useful. For example, choosing one month for the Browse policy and one week for the Retention policy generates an error message.

Directive field Displays the directives recognized by the NetWorker server, including the four sets of preconfigured directives shipped with NetWorker, and any new ones you create using the Directives window.

The preconfigured selection is the Default directive, which skips *core* and *.o* files and backs up only the relevant information in the NFS-based swap files. It also backs up mail files without making any changes to their access time. Use the arrow button to access the Directive scrolling list that contains the preconfigured directives and any new directives previously created in the Directives window.

Group choices Displays all the backup groups known to the NetWorker server.

Select a group for each client. By default the client belongs to the Default group, which has a start time of 3:33 a.m. To enable the scheduled backup to take place at 3:33 a.m. every morning, make sure the Autostart feature is Enabled in the Groups window.

Save set field Lists individual client save sets (which typically comprise one or more client filesystems) for the selected NetWorker client. Enter the name of a client filesystem if you want to back up only that filesystem's data. The preconfigured selection is "All," which backs up all of the client's local filesystems. The Save set scrolling list displays all of the filesystems entered for a NetWorker client.

Remote access field

When this field is blank, only the client itself can recover its backed-up fields. For specific users or another client to have access to a client's data, enter their *user@hostname* in this field. You may also enter a NIS *netgroup* name.

The lower part of the Clients window contains the fields shown in Figure 6-10.

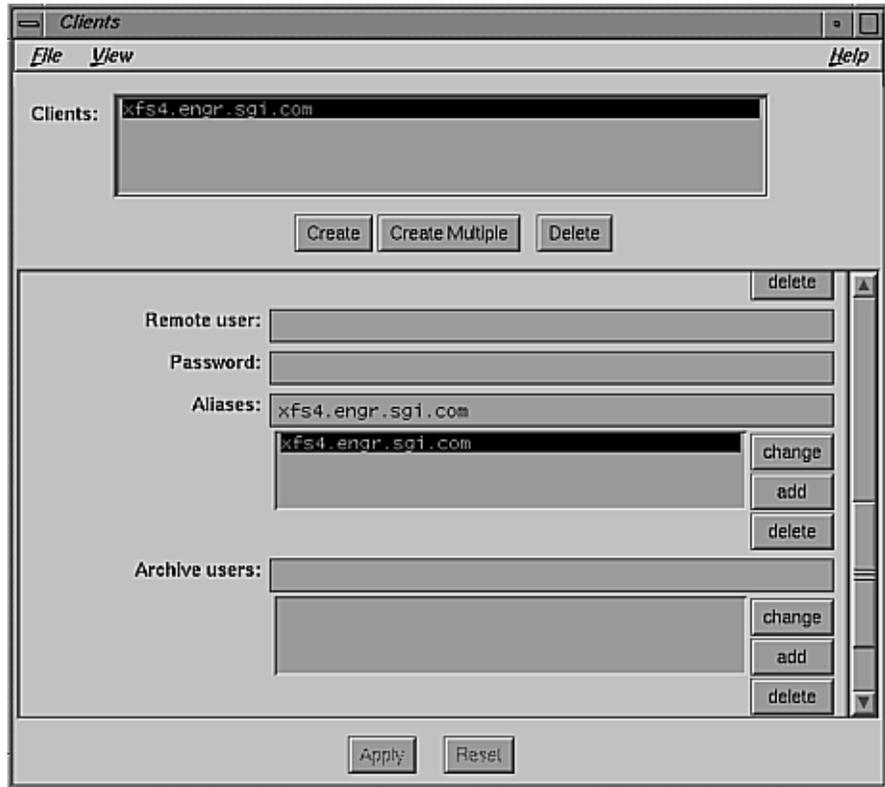


Figure 6-10 Clients Window, Lower Part

Use these fields as follows:

Remote user field

Restrict a client's permissions by entering the client's 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. Enter the password used by the NetWorker server to log into the client system, so that it can back up its file during a scheduled backup. If you enter a password in this field, you must also enter a user name in the Remote user field.

Aliases field Use this field to enter an unlimited number of aliases for a selected NetWorker client. List all of a client's aliases to help eliminate failed saves due to unrecognized or mismatched client names on the network.

Archive users field
Enter user names in this field only if you purchase the NetWorker Archive product. If you are not sure whether your site has NetWorker Archive, consult the *node.lock* file.

NetWorker clients can manually back up and recover files from a NetWorker server and can browse the online file index entries for files. A client may not browse and recover any other client files, unless specifically configured to do so. Normally clients can also be backed up on a regularly scheduled, network-wide backup.

The rest of this section explains

- adding a new client
- scheduling large client filesystems
- removing a client

Adding a New Client

To configure a client, make a selection for

- 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 the client's filesystems)

To configure a new NetWorker client, follow these steps:

1. Open the Clients window.
2. Click the *Create* button.

3. Enter the new client hostname into the Name field.
4. Add the client to one or more backup groups by checking 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 can use the preconfigured policies of Quarter for the Browse policy and Year for the Retention policy.
7. Select a set of directives for the client.
8. Click *Ok*.

If you try to add more clients to your NetWorker server beyond the number you have purchased, this error message appears:

```
Too many clients - maximum is number.
```

Note: Every time you add a new client to NetWorker, it is a good idea to check if NetWorker can access the client in order to back up its files. Use the *savegroup -p* command at the system prompt to see a preview of a group backup without actually backing up any files. See “Previewing a Group Backup” in Chapter 8 for information.

Selecting Client Files to Be Backed Up

Choose which files you want NetWorker to back up for the client. If you want all the client’s files to be backed up, leave the word “All” in the Save set field. If you want only a specific filesystem to be backed up, follow these steps:

1. Select “All” in the Save set scrolling list.
2. Click the *delete* button.
3. Enter the pathname for the filesystem that you want backed up into the Save set field. Figure 6-11 shows an example.

Figure 6-11 Example Pathname for Filesystem to Be Backed Up

4. Click the *add* button. The pathname appears in the Save set scrolling list.

Tip: Use this feature to schedule filesystems to be backed up at different times. For example, if a client has a large amount of data, you can schedule its filesystems separately for backup. For more information, see the section “Scheduling Large Client Filesystems” later in this chapter.

All NetWorker clients are preconfigured so that only the client itself can browse or recover its own files. If your company is concerned about security, do not change the recover access.

Giving Other Systems Access to a Client's Files

To give other systems access to the client's files, follow these steps:

1. In the Remote access field, enter the hostname of the system or netgroup name to which you are giving access to the client's files, as shown in Figure 6-12.

Figure 6-12 Remote Access

2. Click the *add* button next to the Remote access scrolling list.

The hostnames or netgroup names (if you are using NIS) of any systems allowed to browse and recover the client's files appear in the Remote access scrolling list for that client.

3. If desired, make entries in the remaining fields in the Clients window: Remote user, Password, and Aliases. See earlier in this section for explanations of these fields.
4. Click the *Apply* button to apply the selections 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 the previous set of choices.

Scheduling Large Client Filesystems

At a conservative backup rate of 300 KB per second for an 8 mm tape, a full backup of a client with 5 GB of data in several filesystems would take about five 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 takes.

You can, however, schedule the client filesystems to be backed up at different times. You may want to consider this when you have several very large filesystems, or a large number of small filesystems. By separating the client's filesystems into different backup groups, you can still save the client's files without trying to do a time-consuming full backup that requires many changes of the backup volumes.

The Save set scrolling list in the Clients window, shown in Figure 6-9, displays the filesystems that NetWorker backs up for the client. All the filesystems are backed up at the same time, as long as the keyword All is left in the list.

To back up filesystems individually, you 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, and configure the second (duplicate) client to back up the other half of its local filesystems with a second backup schedule in another group.

The example below illustrates how to schedule a client's filesystems separately. First, find out the names and sizes of the locally mounted client filesystems with the `df(1)` command:

```
% df -lk
Filesystem                Type kbytes    use    avail %use Mounted on
/dev/root                  efs   31270    12549   18721  40% /
/dev/usr                   efs  1125799   546991  578808  49% /usr
/dev/dsk/lv6               efs  4581709  4318432  263277  94% /repo
/dev/dsk/lv4               efs  4581709  3678982  902727  80% /online
/dev/dsk/dks1d3s6         efs   929736   846320   83416  91% /d2
/dev/dsk/dks2d3s7         efs  1357229  1190607  166622  88% /d1
/dev/dsk/dks1d3s0         efs   15737     204   15533   1% /tmp
```

The largest filesystems are `/repo` and `/online`, with 4581709 KB of space.

Next, estimate how long it would take to back up one of these filesystems on 8 mm tape:

- $4581709 \text{ KB} / 400 \text{ KB/second} = 11454 \text{ seconds}$
- $11454 \text{ seconds} / 60 \text{ seconds/minute} = 190.9 \text{ minutes}$

You can repeat the same exercise for all the filesystems in order to decide how to separate them into different backup schedules.

Next, create two new backup schedules in the Schedules window:

- Create the first schedule, "Monday Full," which performs a full backup on Mondays and incremental backups the rest of the week.
- Create a second schedule, "Wednesday Full," which performs a full backup on Wednesdays and incremental backups the rest of the week.

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

Next, use the Clients window to configure the client:

1. Open the Clients window.
2. In the Clients scrolling list, select the client whose filesystems you are dividing 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, backup and retention policies, directive, and the "Monday Full" schedule.

4. Delete “All” from the Save set list and add the filesystems you want to back up on the “Monday Full” schedule: `/`, `/usr`, `/repo`, and `/tmp`.
5. Click the *Apply* button.

Now follow the same procedure for the remaining filesystems, which you back up on the “Wednesday Full” schedule.

1. Click the *Create* button to add the client a second time with different filesystems scheduled.
2. Configure the client by choosing the backup group, backup and retention policies, and directive; make sure you choose the “Wednesday Full” schedule for it.
3. Delete “All” from the Save set field, and enter each filesystem you want to back up on the “Wednesday Full” schedule: `/online`, `/d2`, and `/d1`.
4. Click the *Apply* button.

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

Caution: When you use the filesystem scheduling feature, make sure you do not leave any filesystems off the Save set list. Any filesystem left off the list is never backed up.

Removing a Client

To remove a client’s access to the server, follow these steps:

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

Deleting a client means that the client cannot back up its files to the NetWorker server or recover its files. The backup history for the client still remains in the file index and media index, until you specifically remove the entries. See “Removing a Backup Volume” in Chapter 7 for instructions on how to remove the backup history for a client from the indexes.

Fine-Tuning Server Settings

Setting up your server requires a minimal amount of configuring after you first install and run NetWorker. Preconfigured settings 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. Use the Server window to change server parallelism, use devices concurrently, add or change NetWorker administrators, and provide important company and product information.

To open the Server window, select “Server Setup” from the Server pulldown menu, as shown in Figure 6-13.



Figure 6-13 Server Menu

The Server window is shown in Figure 6-14.

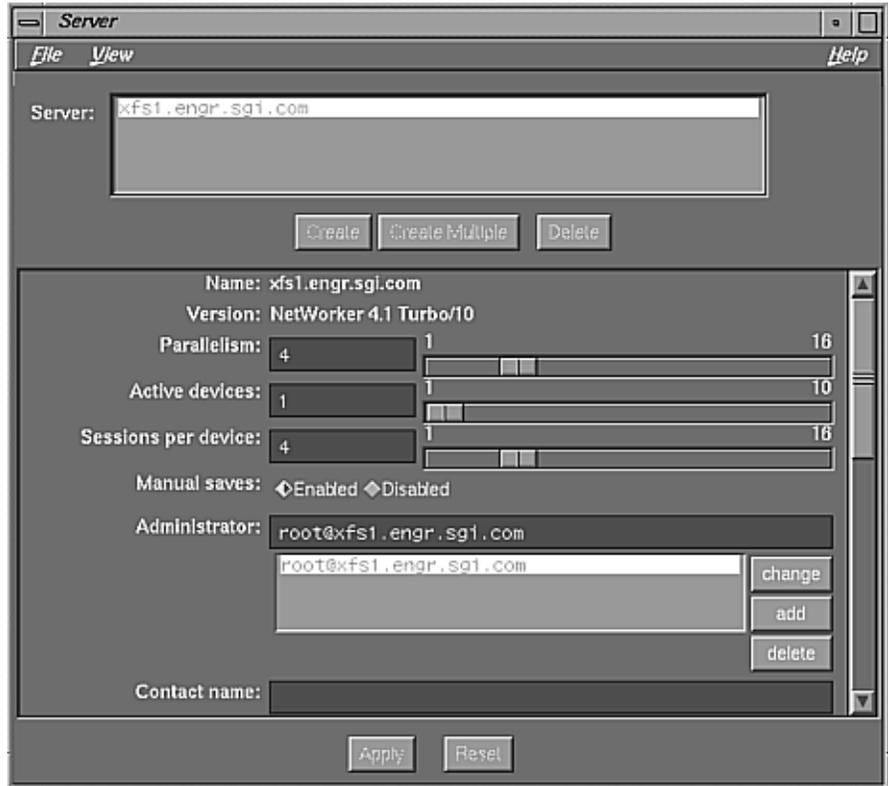


Figure 6-14 Server Window

If necessary, use the scroll bar to the right side of the window to view the entire contents of the Server window.

The features of this window have these functions:

- Name field Displays the name of the current NetWorker server. If you have more than one NetWorker server on the network, and you wish to change servers, use the Change Server dialog box.
- Version field Displays the version of the software currently installed on the NetWorker server.

Parallelism field

Displays the number of clients backing up in parallel. Set the parallelism by using the mouse to slide the bar to the left or right, or enter the number in the Parallelism field. If you have purchased the TurboPak option, the maximum value for parallelism is 64; if you have not, the maximum value is 16.

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 value for the maximum number of concurrent devices is 16.

Note: The Active devices and Sessions per device fields are available only if you purchase NetWorker TurboPak. See “Setting Concurrent Device Use” in this chapter 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 recoveries only. For example, if Parallelism is 12, and Active devices is 3, you can set the Sessions per device to 6. These settings mean that two of the devices handle all of the incoming save sets and the third device remains available for performing a recovery.

Manual saves choices

Use these choices 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 a manual backup, giving the administrator more control over the backup devices and location of the backed-up data. When the Enabled toggle button is selected, any NetWorker client can perform a manual backup.

Administrator scrolling list

Lists 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. Silicon Graphics does not use these fields for this purpose.

The rest of this section explains

- changing parallelism
- setting concurrent device use
- adding or changing administrators
- changing the NetWorker server

Changing Parallelism

A NetWorker server can back up files from many clients in parallel and can multiplex the files onto the same backup volume. This feature supplies a steady stream of files to the server, enabling the backup devices to operate efficiently. Use the Parallelism field in the Server window to optimize the efficiency of your server by controlling the number of clients that back up their files at any one time.

Note: If you have purchased the TurboPak option, the maximum value for parallelism is 64; if you have not, the maximum value is 16.

Tip: If the server takes too long to back up all the clients, increase the parallelism. If the server is so busy backing up clients it is not available for any other tasks, decrease the parallelism. If the server runs out of swap space or memory, decrease the 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 selecting “Server Setup” from the Server pulldown 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.

When one client finishes a backup, another client can start a backup. However, the number of clients allowed to back up at one time cannot exceed the parallelism setting. Presume a case in which there are eight NetWorker client and one server, and parallelism is set to four. In this case, after the first client finishes a backup, the fifth client will start, so that a total of four clients are backing up at one time. That is, any four clients may back up at any one time.

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.

Setting Concurrent Device Use

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. To do so, use the Parallelism, Active devices, and Sessions per device fields in the Server Setup window. Figure 6-15 shows these fields. The values in these fields control the backup traffic from the NetWorker server to the media devices.

Note: You must purchase NetWorker TurboPak to back up to devices concurrently.

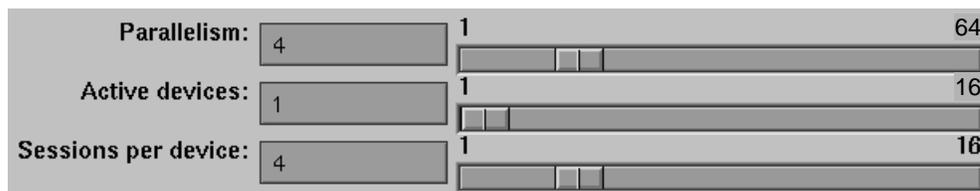


Figure 6-15 Parallelism, Active devices, and Sessions per device Fields

Enter the maximum number of active devices and the desired number of sessions per device. A *save session* is generated when a client starts a backup.

- The Active devices value sets the maximum number of backup devices NetWorker uses concurrently for backups. An active device is one the NetWorker server uses for backups or recovers. The maximum number of concurrent devices is 16.
- 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, the server sends the next session to the next underutilized backup device. 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.

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

1. Open the Server window by selecting “Server Setup” from the Server menu.
2. Set the Parallelism value. (See “Changing Parallelism” earlier in this section.)
3. Select the Active devices slide bar and move it to the right or left to display the value you want.
4. Select the Sessions per device slide bar and move it to the right or left to display the value you want.
5. Click the *Apply* button.

Note: To distribute your client backups evenly across the available backup devices, 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 entered 8 in the Parallelism field, and you entered 2 in the Active devices field, enter the number 4 in the Sessions per device field.

Adding or Changing Administrators

NetWorker is shipped with preconfigured settings so that only `root@server_name` on the server has permission to change system configuration. 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 or netgroup names to the Administrator list.

To change the administrator setting, follow these steps:

1. Open the Server window by selecting “Server Setup” from the Server menu.
2. Enter the new user name or NIS (Network Information Service) netgroup name into the Administrator field. The netgroup names must be preceded by an ampersand (&) in the Administrator field.

For example, you can restrict administrator access to a particular machine with the command `root@server_name`.

3. Click the *add* button. (The *change* button lets you change the currently selected name.)



Figure 6-16 Changing Administrators

4. Click the *Apply* button to apply the changes.

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.

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

Note: The error message `user user_name needs to be on administrator's list` means you do not have permission to make configuration changes.

Changing the NetWorker Server

This section describes how to use the Change Server dialog box to change to another NetWorker server on your network, if you have more than one server. Follow these steps:

1. To open the Change Server dialog box, select “Change” from the Server pulldown menu, or click the *Change Server* speedbar button, as shown in Figure 6-17.

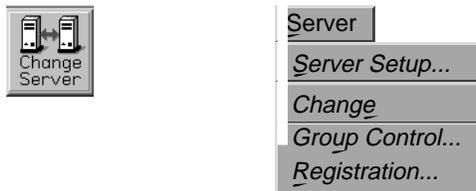


Figure 6-17 Opening the Change Server Dialog Box

Figure 6-18 shows the Change Server dialog box.



Figure 6-18 Change Server Dialog Box

2. Use the *Update Server List* button to display all of the NetWorker servers on your network, if necessary.
3. Highlight the hostname of the desired server in the Server scrolling list, or enter a valid hostname in the Server field.
4. Click *Ok* to complete the operation and close the window.

Configuring and Managing Volumes and Backup Devices

This chapter covers the system administration tasks for managing media and backup devices:

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

Note: This chapter does not apply to autochangers, such as jukeboxes or tape libraries. For information on managing jukeboxes, see Chapter 11, “Using NetWorker with Jukeboxes.”

Labeling Backup Volumes

You must give a backup volume a *volume label* before you can use it for backing up files. NetWorker uses the volume label to maintain a record of each backup volume in the online *media index*. A record of each file on every backup volume is maintained in the NetWorker *file index*. NetWorker uses these two indexes to tell you (or the jukebox) which backup volume to mount when you want to recover lost data.

Every backup volume belongs to a *volume pool*. You can use the preconfigured pools or create your own. If you choose not to separate your backup volumes into pools, NetWorker automatically uses a pool named Default for all your backup volumes.

Each volume pool has a matching *label template* associated with it. 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 a custom label template in the Label Templates window.

NetWorker automatically labels the backup volumes with the next sequential label from a template unless you choose to override it with one not associated with a pool. Volume pools and label templates are described in detail in “Using Label Templates” and “Using Volume Pools” in this chapter.

Each backup volume should have an adhesive label attached to it which matches the internal NetWorker label. For jukebox volumes, use bar code labels if appropriate. There are no restrictions for labeling the backup volumes, except that each volume requires a unique label. If you choose not to use a preconfigured template, you can label the volumes with names appropriate to your network. No matter what naming scheme you choose, keep in mind that a convention should be easily understood by operators and administrators.

Tip: In the NetWorker environment, 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 used for the backup and a number. For example, atlas.001 for the backup volume from the NetWorker server named atlas, and mars.010 for the backup volume from the NetWorker server named mars.

Note: Backup volumes are not overwritten, but are appended to each time they are used. Therefore, you can leave the same backup volume(s) in the server backup device(s) no matter what backups are scheduled. When one backup volume is full, NetWorker requests a new one to be mounted either by you or by the jukebox until the accumulated backups have completed. Since NetWorker maintains a record of the data stored on each volume in the indexes, it is not necessary to calculate tape capacity requirements.

Using the Label Dialog Box

Use the Label dialog box to label backup volumes. To open the Label dialog box choose the “Label” command from the Media pulldown menu, or click the *Label* speedbar button, as shown in Figure 7-1.

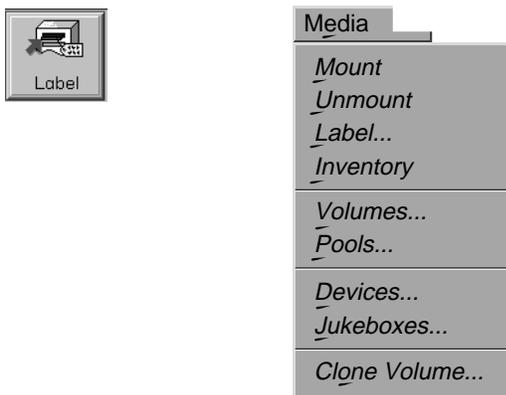


Figure 7-1 Opening the Label Dialog Box

Figure 7-2 shows the Label dialog box.

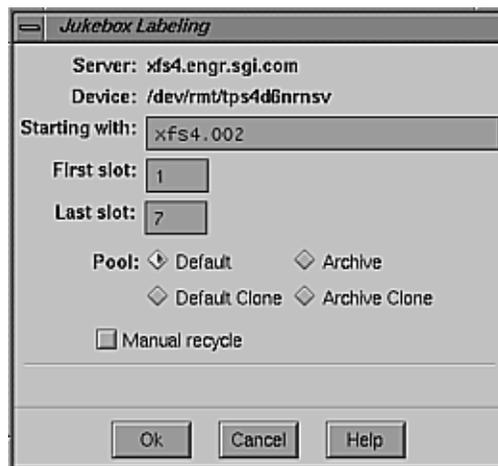


Figure 7-2 Label Dialog Box

Note: If you are using a jukebox, the volume labeling and mounting operations can be automated. See Chapter 11 for more information.

The features of this window have the following functions:

Device field Contains the name of the selected device. To select another device, click *Cancel* to close the Label dialog box, highlight the new device in the NetWorker Administrator window, and select the “Label” command again.

Volume name field
Contains the label that will be assigned to the volume and displays the name of a 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.

Tip: To prelabel as many volumes as you like at one time without mounting them, deselect the *Mount after labeling* toggle button.

Pool choices Choose a volume pool from the Pool choices. Only pools that have been enabled in the Pools window (Media window) appear in the Label dialog box. All preconfigured pools that appear in the dialog box, including Default, Default Clone, Archive, and Archive Clone, are by default enabled. NetWorker automatically uses the label template associated with the selected pool to label the backup volume.

Note: For more information on volume pools, see “Using Volume Pools” later in this chapter.

Manual recycle toggle button
Select if you want to override the volume retention policy. The volume mode will appear as man (manual recycle) in the Volume Management window. Manual recycle mode is explained in “Viewing the 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.

Ok button

Click to label a backup volume, mount it, or both. If you make a mistake, click the *Cancel* button.

When you have completed labeling, the backup volume name appears in the Devices display of the NetWorker Administrator window beside the path of the device. The name of the volume appears only if it has been mounted.

If 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.

Caution: Relabeling a backup volume destroys all records of its contents under the old label. Therefore NetWorker asks for confirmation with the message `About to relabel`.

Labeling and Mounting a Backup Volume

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

Note: Use these instructions only if you are labeling and mounting backup volumes in a backup device that is not associated with a jukebox.

1. Place a blank backup volume in the NetWorker server backup device.
2. Open the Label dialog box by selecting “Label” from the Media menu or by clicking the *Label* speedbar button.
3. Choose a volume pool from the Pool choices.
4. If you want a different label from the one chosen by the label template, delete the name from the Volume name field and enter the new name.
5. Select the *Manual recycle* toggle button if you want to override the backup volume’s retention policy.
6. Click the *Ok* button when you are ready to label the volume, mount it, or both. If you make a mistake, click the *Cancel* button to cancel the operation.

Note: If 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. Relabeling a backup volume destroys all record of its contents under the old label.

Mounting and Unmounting Backup Volumes

Before you can back up files to a backup volume or recover files from it, you must mount it in the server backup device. If you are using a jukebox, NetWorker automatically mounts the backup volumes when they are needed.

If you have more than one device for backups, all the server backup devices used for backing up and recovering files appear in the Devices display of the main window.

Use the “Mount” and “Unmount” commands in the Media pulldown menu or the associated speedbar buttons to mount and unmount backup volumes. See Figure 7-3.

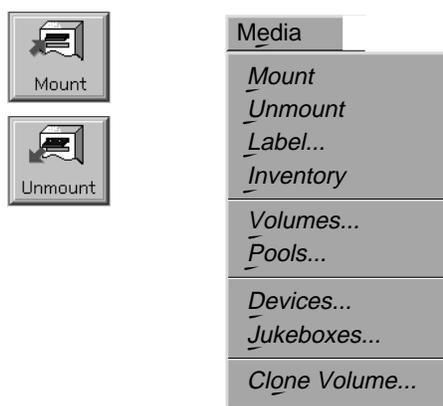


Figure 7-3 Mounting or Unmounting a Backup Volume

To mount a backup volume in a specific device, select the device name in the Devices display of the main window and click the *Mount* speedbar button.

The Devices display changes to show the name of the backup volume mounted in the device.

Tip: NetWorker allows you to premount backup volumes, so that you can complete unattended network-wide backups using backup devices in a series or concurrently.

To unmount a backup volume from a specific backup device, select the name of the device in the Devices display and click the *Unmount* speedbar button. The Devices display changes to show the backup volume unmounted in the backup device.

If you try to back up files when no backup volumes are mounted, NetWorker requests a writable volume with this message in the Pending display (a writable volume is one that is labeled and has room for data):

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

One of three messages is displayed, 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 DLT tape volume name
```

If you need more than one backup volume to recover the file(s), NetWorker gives you a list of all the backup volumes in the order they will be needed. During the recovery process, NetWorker requests each backup volume it needs, one at a time. If you are using a jukebox with one of the NetWorker Support for Jukebox options, volume mounting is automatic for backup volumes already stored in the jukebox.

Finding a Backup Volume Name

If the physical adhesive label on the backup volume is missing or illegible, you can determine its name in one of two ways:

- Mount the volume in the backup device by clicking the *Mount* speedbar button; the volume's name appears next to the device in the Devices display of the NetWorker Administrator window.

- Load the volume in a device and select the *Label* speedbar button. The Label dialog box shows the name of the label in the Volume name field, as shown in Figure 7-2.

Automatically Managing Backup Volumes

NetWorker provides flexibility for managing backup volumes and their save sets. Use the Volume Management window to manage backup volumes manually, and use the Policies window to manage them automatically.

NetWorker uses browse and retention policies to manage backup volumes and their save sets automatically. However, you can manually override the automatic policies by using various commands in the Volume Management window.

Note: For more information about browse and retention policies, see the section “Determining Browse and Retention Policies (Index Policies)” in Chapter 5.

The Volume Management window contains information about the media index, and also information about the NetWorker server’s collection of backup volumes and the save sets they contain. To view information about your backup volumes, open the Volume Management window by selecting Volumes from the Media pulldown menu, or click the *Volumes* speedbar button, as shown in Figure 7-4.

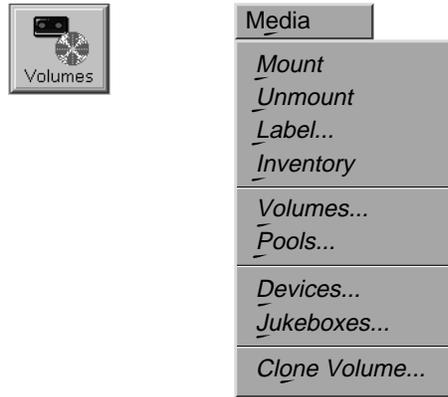


Figure 7-4 Opening the Volume Management Window

The Volume Management window appears, as shown in Figure 7-5.

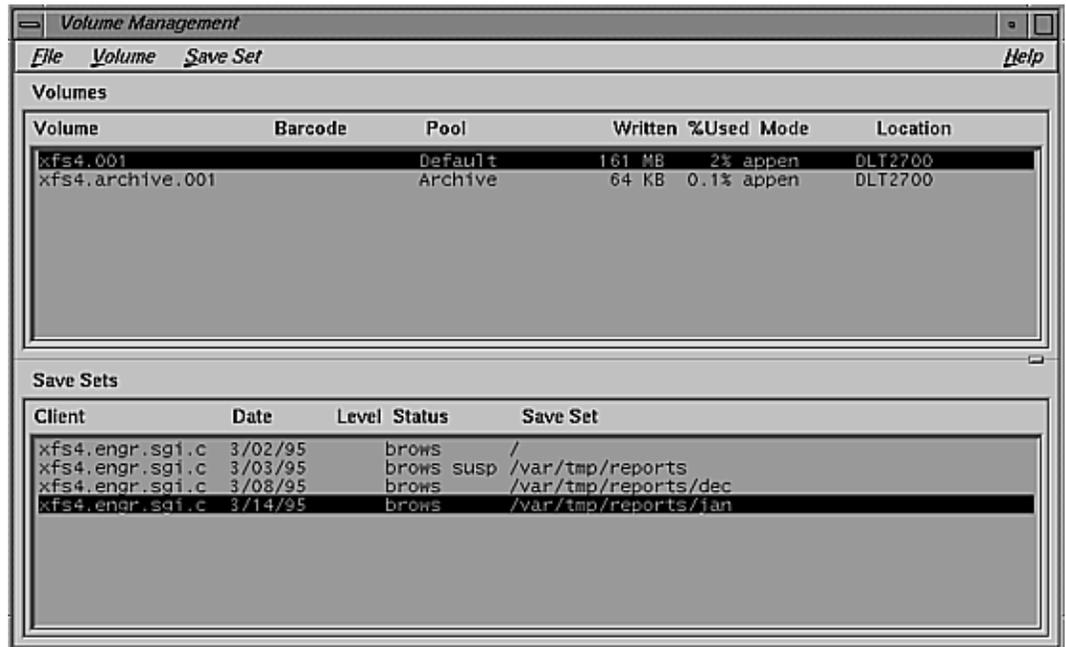


Figure 7-5 Volume Management Window

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

Viewing the Volumes Scrolling List

The Volumes scrolling list provides the following information about each backup volume:

Volume	Backup volume name.
Barcode	Bar code label, if one exists.
Pool	Name of the pool to which the backup volume belongs.
Written	Amount of data written on the volume.
% Used	Percentage of the backup volume used compared to its estimated total capacity. If a volume is full, NetWorker displays either 100% or Full.
Mode	Status or kind of backup volume. If both the volume mode and %Used are marked Full, NetWorker will display only Full in the %Used column. <ul style="list-style-type: none">• appen (abbreviation for appendable): The backup volume has room for more data.• full: The backup volume has no room for more data. The save sets on the volume have not yet passed the time period specified by the Browse and Retention policies, which would make the volume available for recycling.• recyc (abbreviation for recyclable): All the save sets on the volume have passed both the Browse and Retention policies, and the volume is ready to be relabeled. Once the volume has been relabeled, it is eligible to be overwritten with new data.• man (abbreviation for manual recycle): means that both the volume and its save sets must be manually recycled. This mode is always tagged on as an addition to one of the other three volume modes.

Location: Location of the backup volume, which is entered by the user as a reminder of a volume's location.

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

Viewing the Save Sets Scrolling List

The Save Sets scrolling list provides the following information about each save set stored on a 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 does not exist, the save set was generated by a manual backup.
Status	Type of save set: <ul style="list-style-type: none">• brows (abbreviation for browsable): The save set has an entry in the file index for the client because it has not passed the time period specified by the Browse policy for the client. The files contained in a browsable save set can be browsed and marked for recovery using the NetWorker Recover window.• recov (abbreviation for recoverable): The save set entry has been removed from the file index, either automatically due to the client Browse policy or manually using the Remove oldest cycle button in the Indexes window. To recover files contained in this save set, use the Save Set Recover window if you want only to recover the data. If you want to recover both the file index entries and the data, use the scanner command. (See the <i>scanner</i> reference page for information on using the <i>scanner</i> command.)

- **recyc** (abbreviation for recyclable): 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** (abbreviation for scanned-in): The save set was scanned in using the scanner command. You must remove these save sets manually from the index. The standard index management policies do not apply to scanned data.
- **in-pro** (abbreviation for in-progress): The save set is in the process of being backed up
- **abort** (abbreviation for aborted): The backup was aborted manually by the administrator or the system crashed
- **susp** (abbreviation for suspect): A previous recovery attempt failed. This status is always displayed as an addition to one of the other six 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

You can manually mark entire backup volumes and their save sets as recyclable, using the Volume Management window. You can set a location for a volume, change the mode of a volume, or remove a volume from the media index. This section explains

- changing the status of a save set
- setting the location of your backup volumes
- changing the mode of a backup volume
- determining backup media capacity

- recycling volumes
- removing a backup volume

Changing the Status of a Save Set

Use the “Change Status” command in the Save Set pulldown menu in the Volume Management window to change the status of a save set, as shown in Figure 7-6. NetWorker marks a backup volume suspect if an error occurred trying to read the media.

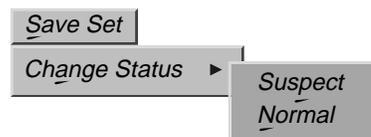


Figure 7-6 Changing the Status of a Save Set

You might want to change the status of a save set to normal if you know that the save set’s data is not really suspect, but NetWorker had difficulty reading the data for other reasons. For example, the media drive heads might have been dirty, causing NetWorker to assign a status of suspect to the save set.

Select “Suspect” if you made a mistake changing the save set from suspect to normal, and you want to change the save set back to suspect again.

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

1. Open the Volume Management window.
1. In the Save Sets scrolling list, select the save set whose status you want to change.
2. Choose the appropriate Change Status command from the Save Set menu (see Figure 7-6).

Setting the Location for Your Backup Volumes

To help you find a volume when it is needed for recovering a file, you can set a location to your backup volumes. A volume's location might include a vault, shelf, jukebox 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. Select "Set Location" from the Volume menu. The Set Location dialog box appears, as shown in Figure 7-7.

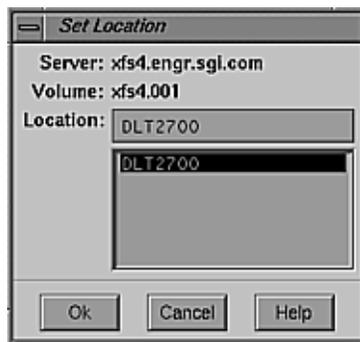


Figure 7-7 Set Location Dialog Box

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

After you apply a location to a backup volume, it is listed in the Set Location dialog box and can be applied to any backup volume in the Volume Management window.

Note: Each time a volume's location changes within a jukebox, NetWorker automatically updates the volume's location with the name of that jukebox. This feature is useful because if you use volumes that were previously

assigned another location, their new location becomes updated automatically. If you move volumes offsite or to a different facility, you must manually change their assigned location to accurately reflect their new location.

Tip: When NetWorker requests that you mount a backup volume during a recovery, you can find the volume by referring to its location in the Volume Management 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. Four possible modes for a backup volume exist. See the section “Viewing the Volumes Scrolling List” earlier in this chapter for a description of three of the modes: appendable, full, and recyclable. This section introduces you to the fourth mode, Read Only.

The Volume pulldown menu (see Figure 7-8) contains the Change Mode command that allows you to change the currently selected volume to one of the following: Appendable, Read Only, or Recyclable.

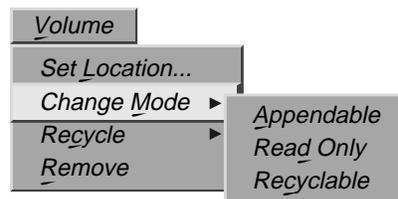


Figure 7-8 Changing Backup Mode

Use Table 7-1 to determine which existing volume modes can change.

Table 7-1 Changing Volume Mode

Volume's Current Mode	Change to Appendable	Change to Read Only	Change to Recyclable
Full	Yes*	Not necessary	Yes
Recyclable	Yes*	Yes	No
Appendable	No	Yes	Yes

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

A volume mode of Read Only or Full means that NetWorker cannot write data to that volume. NetWorker automatically assigns Full mode to a volume that has filled with data. An administrator manually assigns Read Only to a volume so data cannot be written to it, whether it is full or not. When you change a volume's mode to Read Only, the volume is marked with an (R), and the previously assigned mode is removed.

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

1. Highlight a volume in the Volumes scrolling list. Notice that the save sets stored on the volume appear in the Save Sets scrolling list.
Note: If the status of all of the save sets is recyc, then NetWorker automatically marks the volume as recyclable.
2. Select "Change Mode" from the Volume menu in the Volume Management window.

A confirmation box appears (see Figure 7-9), asking if you want to change the volume to the mode you have chosen.

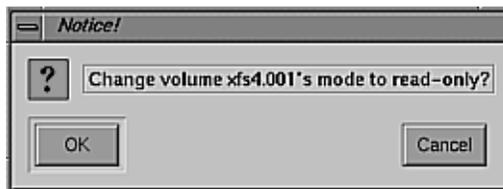


Figure 7-9 Changing Volume Mode: Confirmation Box

3. Click the *Ok* button 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.

Caution: If a cloned volume does not exist when you manually change the mode of a volume to recyclable, all of its save sets change to a recyclable status, regardless of their relationship to the browse or retention policies specified in the Clients window. If a cloned volume does exist, the status of the save sets will remain unchanged. Do not change the mode to recyclable if any possibility exists that you might still need the save sets on the volume.

Determining Backup Media Capacity

Occasionally NetWorker marks backup volumes as “full” when they are not really full. (The Volume Management window and the `mminfo -m` command display the details of the backup volumes.)

NetWorker marks magnetic tape as being 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 is marked as “full” has a bad spot on 13% of the length in the beginning of the tape. A tape like this can still be used for recoveries, but cannot be used for any more backups.

Tip: If you see this “bad spot” behavior on many of the backup volumes, it may indicate that the device needs cleaning or maintenance.

Tapes are also marked “full” when they are recovered after being deleted from the media index.

Recycling Volumes

You can override a volume's retention policy by changing it to manual recycle. Select "Recycle" from the Volume pulldown menu; the Manual and Auto commands appear, as shown in Figure 7-10.

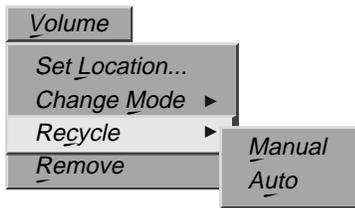


Figure 7-10 Recycling Volumes

You might want to change to manual recycle because you need to keep the save sets on a volume longer than its retention policy specifies. You can also change a volume marked manual recycle back to automatic recycle, so the volume will once again use its original retention policy.

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

1. Select a volume from the Volumes scrolling list.
2. Select "Recycle" from the Volumes menu in the Volume Management window; the Manual and Auto commands appear.
3. Select "Manual" from the Recycle menu. NetWorker identifies the volume by adding "man" to the volume's mode in the Volumes scrolling list.

To remove manual recycle from the volume, highlight the volume again and select "Auto" from the Recycle menu.

Removing a Backup Volume

Remove a backup volume from the media index only if the volume has become physically damaged or unusable. To remove a volume from the media index, select “Remove” from the Volume pulldown menu in the Volume Management window, as shown in Figure 7-11.



Figure 7-11 Removing a Volume

When you select a volume for removal, NetWorker checks to see if the volume has a *clone*, that is, a copy of the save sets on another volume. If the volume has a clone, NetWorker does not remove the entries from the online file index. If the volume does not have a clone, NetWorker removes the entries from the online file index.

Note: See the next section in this chapter, “Cloning Backup Volumes,” for information on using clones.

Tip: If a backup volume is lost or destroyed, you should remove its entries from both the file and media indexes. Otherwise, NetWorker might request the volume for recovering data.

To remove a volume from the media index, follow these steps:

1. Highlight the desired volume in the Volumes scrolling list. The volume’s save sets appear in the Save Sets scrolling list.
2. Select “Remove” from the Volume menu. A window appears, as shown in Figure 7-12.

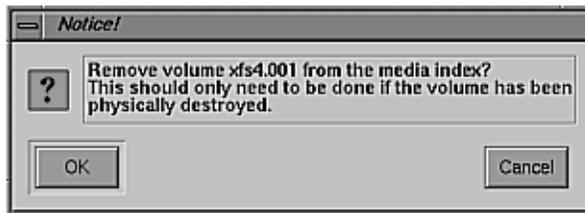


Figure 7-12 Warning Message for Removing a Volume

3. Click the *Ok* button in this window 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.

Note: If you accidentally delete a backup volume from the media index, you may still recover the data by using the *scanner* program, as long as you have not relabeled the volume.

Cloning Backup Volumes

NetWorker protects the data on your network by backing up files on a daily basis. *Cloning* provides additional protection by allowing you to make identical copies of the data on your backup volumes on an as-needed basis. When you clone backup volumes, you are actually copying all of the save sets that have been backed up to that volume; you do not create an identical copy of the backup volume.

Cloning backup volumes is beneficial for several reasons:

- Creating clones of your media enables you to store identical copies of your data offsite or in a vault for protection. Cloning backup volumes also allows you to easily share or transfer important data to other company facilities.
- 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 reach their expiration date to continue to assure the safety of your data.

- Cloning also helps you determine the integrity of backed-up data because NetWorker displays an error message if it has difficulty reading the data.

Note: During the cloning process, NetWorker does not create an additional set of index entries for either online index. Both the original and cloned volumes use the same set of file and media index entries. Consequently, if you eliminate either the cloned or original volume from the media index, NetWorker will not let you remove the index entries as long as one of the volumes still exists.

Cloning entire backup volumes can only be performed manually; you cannot automatically clone backup volumes. To clone data (save sets) automatically, use the Groups window to enable cloning for individual groups.

Note: For more information about cloning data (save sets), see “Cloning Save Sets” in Chapter 9.

Because save sets span multiple backup volumes, NetWorker might request more than one source volume to complete the cloning operation. Furthermore, because media vary in size and capacity, you may need more destination volumes to hold the cloned data than the source volumes required. Therefore, it is helpful to have extra media available while cloning backup volumes.

If you have a jukebox, NetWorker automatically selects the next available volume marked “appendable” from the clone pool for cloning. If you do not want to mix cloned data on the same destination volume, load, label, and mount a blank volume in the backup device before cloning. See “Example: Cloning Backup Volumes” later in this section.

Recovering from Cloned Volumes

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

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

NetWorker considers volumes for recovery in the following order:

- a volume mounted in the drive
- a volume stored in the jukebox
- a volume stored someplace other than in the jukebox

Note: If you send cloned volumes to another company location that also uses NetWorker, they must use the *scanner* command to retrieve the data because the cloned volumes do not have an entry in their indexes. See the reference page for more information about the *scanner* command.

Using the Clone Windows

The original backup volume is the *source* volume, and the volume to which you clone the data is the *destination* volume. The destination volume is always a different volume than the source volume. NetWorker separately tracks cloned volumes (destination volumes) and marks them as clones in the media database.

The requirements for cloning the save sets on your backup volumes with NetWorker are as follows:

- IRIX NetWorker release 4.1.1 or greater (release 4.1.1 also clones volumes created with earlier releases of NetWorker software)
- NetWorker TurboPak
- multiple backup devices
- source and destination volumes

Note: NetWorker can also clone individual save sets, manually and automatically. See “Cloning Save Sets” in Chapter 9.

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, select “Clone Volume” from the Media pulldown menu, as shown in Figure 7-13.



Figure 7-13 Opening the Clone Volume Window

The Clone Volume window appears, as shown in Figure 7-14.

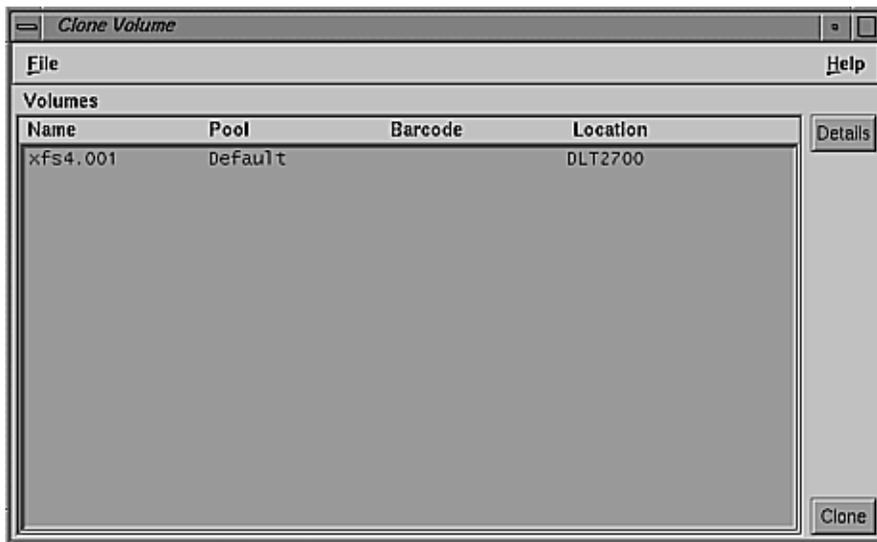


Figure 7-14 Clone Volume Window

The features of the Clone Volume window have the following functions:

Volumes scrolling list

Contains a list of all the backup volumes known to the server. The list includes the volume and pool name, bar code label, and its location. You may clone the save sets on an unlimited number of backup volumes at the same time. However, consider cloning your backup volumes after a scheduled backup has completed, so you do not overload your NetWorker resources and backup devices.

Details button Displays backup volume details. Figure 7-15 shows the Clone Volume Details window.

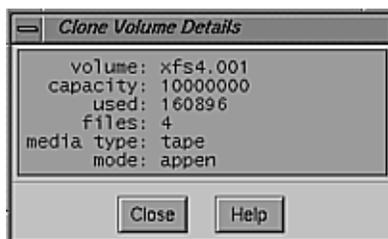


Figure 7-15 Clone Volume Details Window

The features of the Clone Volume Details window have the following functions:

volume	Name of 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 backup volume: appen (appendable), full (full), or recyc (recyclable). See the section “Viewing the Volumes Scrolling List” in this chapter for an explanation of the volume modes.
<i>Clone</i> button	Opens the Clone Volume Status window to begin cloning the selected volumes. The Clone Volume Status window appears, as shown Figure 7-16.

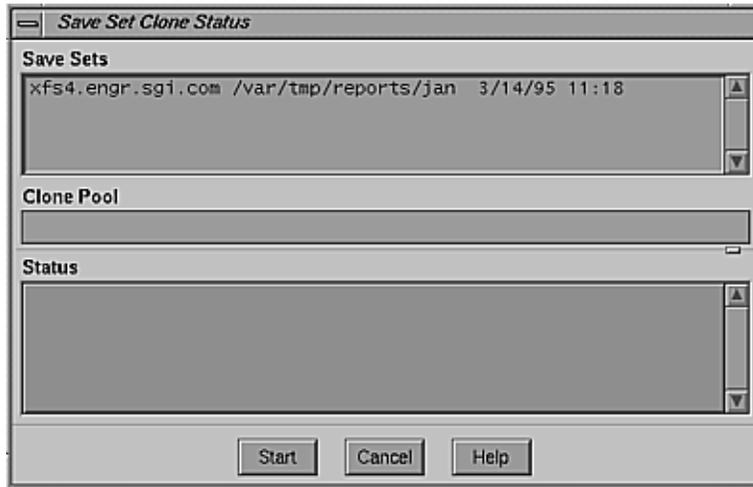


Figure 7-16 Clone Volume Status Window

Use features of the Clone Volume Status window as follows:

Volumes scrolling list

View the name of the volume selected for cloning and the pool to which it belongs.

Because save sets span multiple backup volumes, NetWorker may request more than one source volume to complete the cloning operation. Furthermore, because media varies in size and capacity, you may need more destination volumes to hold the cloned data than the source volumes required. Therefore, it is helpful to have extra media available while cloning backup volumes.

If you have a jukebox, NetWorker automatically selects the next available volume marked “appendable” from the clone pool for cloning. If you do not want to mix cloned data on the same destination volume, load, label, and mount a blank volume in the backup device before cloning. See “Example: Cloning Backup Volumes” later in this section.

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.

Start button Click to begin cloning

Status scrolling list

View the progress of the cloning operation. The same messages appear in the NetWorker Administrator window.

Example: Cloning Backup Volumes

In this example, the network operator needs to clone the full backups from the engineering department to store offsite. 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 needs to 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 will receive a message in the Pending display of the NetWorker Administrator window requesting that you load the volume.

Follow these steps to clone a backup volume:

1. Make sure you have plenty of backup media for the volume to be cloned. NetWorker may request more than one source volume to complete the cloning operation. Furthermore, because media varies in size and capacity, you may need more destination volumes to hold the cloned data than the source volumes required.
2. Load a blank volume in the media drive.
3. Label the volume.
4. Mount the volume in the backup device.

Note: If you do not load, label, and mount a volume and you back up to a jukebox, NetWorker selects an “appendable” volume from the clone pool entered in the Clone Status window. If you back up to standalone devices, a message appears in the Messages display of the NetWorker Administrator window, requesting that you mount a volume.

5. Open the Clone Volume window by choosing “Clone Volume” from the Media menu.
6. Select the volumes that contain the full backups from engineering in the Volumes scrolling list.
7. Click the *Clone* button in the Clone Volume window to open the Clone Volume Status window.
8. Enter the name of the clone pool in the Clone Pool field.
9. Click the *Clone* button in the Status Clone Volume window to begin the cloning process.

The Status scrolling list displays messages that monitor the cloning process.

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 a custom label template using the Label Templates window.

Note: You can also label jukebox volumes with bar code labels. See “Labeling Volumes and Using Optional Bar Codes” in Chapter 11 for information.

NetWorker uses label templates to label and help 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” in this chapter.

It does not matter to NetWorker how the volumes are labeled, as long as each one has a unique name. NetWorker 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.

This section explains

- using the preconfigured label templates
- choosing a label template
- creating a label template
- label template tips and examples

Using the 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. NetWorker also includes preconfigured volume pools with corresponding names (except for Two Sided). The preconfigured volume pools automatically use the preconfigured label template by 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 formats for each preconfigured label template are as follows:

Archive The Archive label template has three fields, separated by periods. The first field contains the name of the NetWorker server, the second field is “archive,” and the third field contains a number. The Archive label template is used only for clients that need to archive data.

Archive Clone The Archive Clone label 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.

Default	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.
Default Clone	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.
Full, NonFull, and Offsite	In each of these three preconfigured label templates, the name of the label template appears in the first field, and the second field contains a number.
Two Sided	The Two Sided template is for use with two-sided media such as optical media, for which you must label both sides. The first field contains the name of the server, the second field contains a number, and the third field contains either an “a” or “b” to differentiate between the two sides of the media.

Table 7-2 shows the format and examples for each preconfigured label template.

Table 7-2 Preconfigured Label Templates

Label Template	Format	Examples
Archive	<i>server.archive.number</i>	space.archive.001 space.archive.099 atlas.archive.325
Archive Clone	<i>server_c.archive.number</i>	moon_c.archive.001
Default	<i>server . number</i>	space.675 space.800 atlas.054
Default Clone	<i>server_c.number</i>	moon_c.001
Full	<i>label_name . number</i>	Full.076
NonFull	<i>label_name . number</i>	NonFull.003

Table 7-2 Preconfigured Label Templates

Label Template	Format	Examples
Offsite	<i>label_name . number</i>	Offsite.120
Two Sided	<i>server . number . side</i>	phoenix.001.a phoenix.001.b

Choosing a Label Template

Before you choose a label template for labeling your volumes, follow these steps:

1. Choose the device that has tapes mounted for labeling from the Devices display in the NetWorker Administrator window.
2. In the Media pulldown menu, select “Pools”; choose a label template in the Pools window.
3. Make sure the Enabled button is selected for the label template you want, as shown in Figure 7-17.

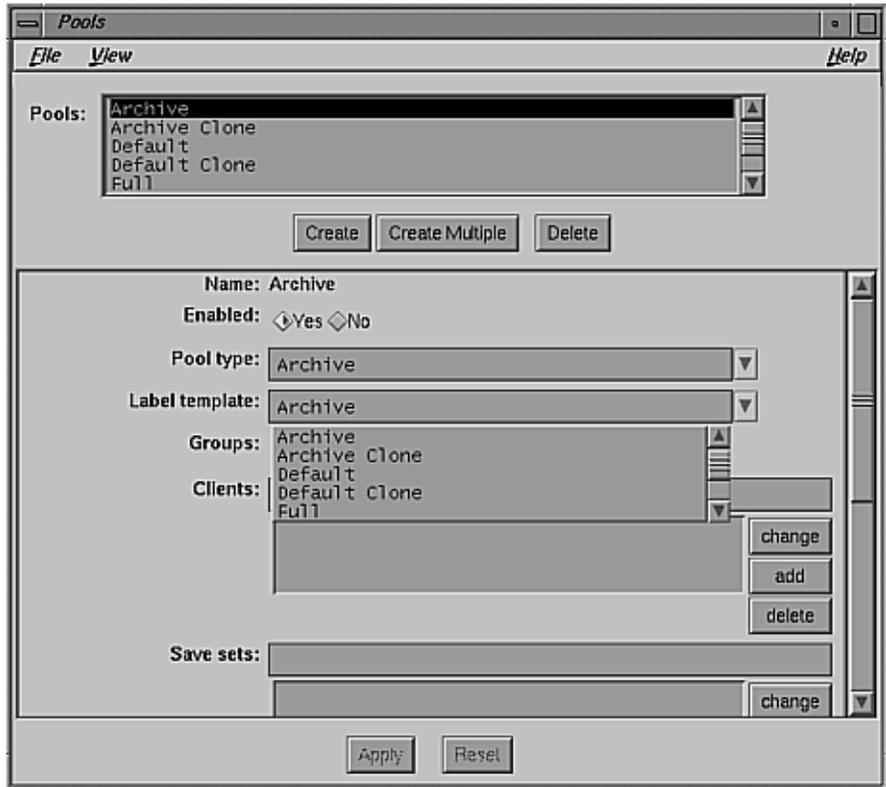


Figure 7-17 Pools Window

Note: Whether you use a preconfigured template or a custom template to label the backup volumes, remember to attach an external adhesive label to each volume with a name that corresponds to the template you have chosen.

4. Select "Label" from the Media pulldown menu.

Creating a Label Template

To create and add a label template, open the Label Templates window by selecting “Label Templates” from the Customize pulldown menu. The Label Templates window is shown in Figure 7-18.

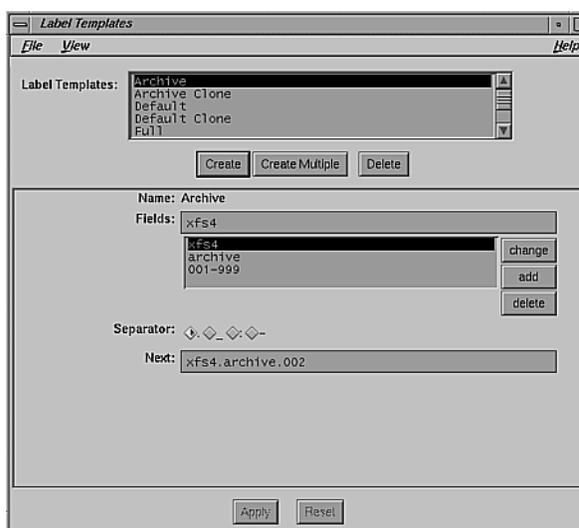


Figure 7-18 Label Templates Window

The features of this window have the following functions:

Label Templates scrolling list

Lists label templates known to the NetWorker server. The preconfigured label templates and any new templates you create appear in this list.

Name field Displays the name of the currently selected label template.

Fields field	<p>Displays the components of the label template. The components appear as one component per line in the display:</p> <ul style="list-style-type: none">• range of numbers, for example, 000–999• range of lowercase letters, aa–zz• range of uppercase letters, AA–ZZ• character string, for example, Accounting and 001 <p>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 can have as many components as you want, but it is best to keep the template simple, with few components. The total length of the label must be less than 64 characters.</p> <p>If you enter a range of numbers or characters, they will change incrementally with each new label. For example:</p> <table><tr><td>first label</td><td>Engineering.001</td></tr><tr><td>second label</td><td>Engineering.002</td></tr><tr><td>third label</td><td>Engineering.003</td></tr></table> <p>Note: When you create a label template, the order in which you enter each component of the template is important. That is, NetWorker applies each component in a left-to-right order, starting with the first one you enter. The components are displayed from top to bottom in the Fields scrolling list, with the first component on the top.</p>	first label	Engineering.001	second label	Engineering.002	third label	Engineering.003
first label	Engineering.001						
second label	Engineering.002						
third label	Engineering.003						
Separator choices	<p>Allows you to choose a separator character (period, dash, colon, or underscore) to separate each component of the label template. If no button is selected, the components will have no separators; for example: AA00aa.</p>						
Next field	<p>Displays the next label NetWorker uses to label the backup volume according to the template. To force a label or start the labeling scheme at a certain point, enter the starting</p>						

label into this Next field. NetWorker continues to generate labels from that point on, according to the rules of the template.

Label Template Tips and Examples

Making your label template match your backup media storage system can help you organize and locate the media in storage. For example, if you store backup media in bins or racks, you can include a place for the bin or rack number in the label template. See “Example 1: A Storage-Oriented Template later in this section.

It is also helpful if the label template name matches the name of the corresponding pool. Table 7-3 shows examples.

Table 7-3 Label and Pool Names

Name	Label	Pool Name
AcctFull	AcctFull.001 AcctFull.002 AcctFull.003	Accounting
EngTest	EngTest.001 EngTest.002 EngTest.003	Engineering Test

This section concludes with three example templates: a storage-oriented template, a sequential label template, and a two-sided media label template.

Example 1: A Storage-Oriented Template

In this example, assume you have only one NetWorker server and one jukebox. 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.

In this example, as diagrammed in Figure 7-19, three racks are 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, arranged left to right.

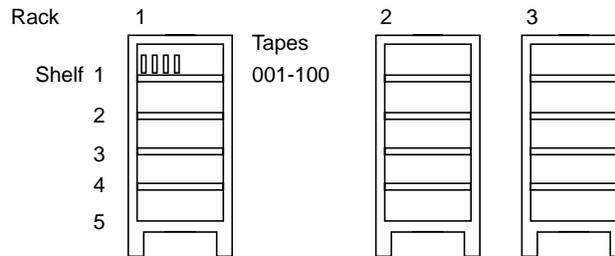


Figure 7-19 Example Racks of Magnetic Tapes

Using this labeling scheme, a tape labeled “2-3-54” is stored in the second rack, on the third shelf, and is positioned between the fifty-third and fifty-fifth tapes on the shelf. NetWorker labels the tapes in the jukebox 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 it labels the tapes on the second shelf of the first rack, numbering them from 1-2-001 to 1-2-100, and so forth, until it labels the last tape in the third rack 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 would be wiser to create a template allowing for 100 tapes and not use all of them, than to create a template for only 10 tapes and soon use all 10 labels.

Follow these steps to create the storage-oriented label template:

1. Open the Label Templates window by selecting “Label Templates” from the Customize menu.
2. Click the *Create* button.
3. Enter the name of the labeling scheme into the Name field (for example, “Engineering backups”).

4. Enter the component of the first field into Fields. Remember, this field holds the rack number, so you enter 1-3.
5. Click the *add* button.
6. Enter the second component of the label template into Fields. This field holds the shelf number, so enter 1-5 into the field.
7. Click the *add* button.
8. Enter the third component of the label template, the tape position numbers. Enter 001-100 into Fields.

You must enter the same number of characters for any numbers in the label template. For example, use 01-20 not 1-20, and 1-9 not 01-9.

9. Click the *add* button.
10. Click the dash (-) *Separator* button to choose it as the separator for each component of the label template
11. Click the *Apply* button to apply the new template to the NetWorker server. The Label Templates window now looks similar to Figure 7-20.

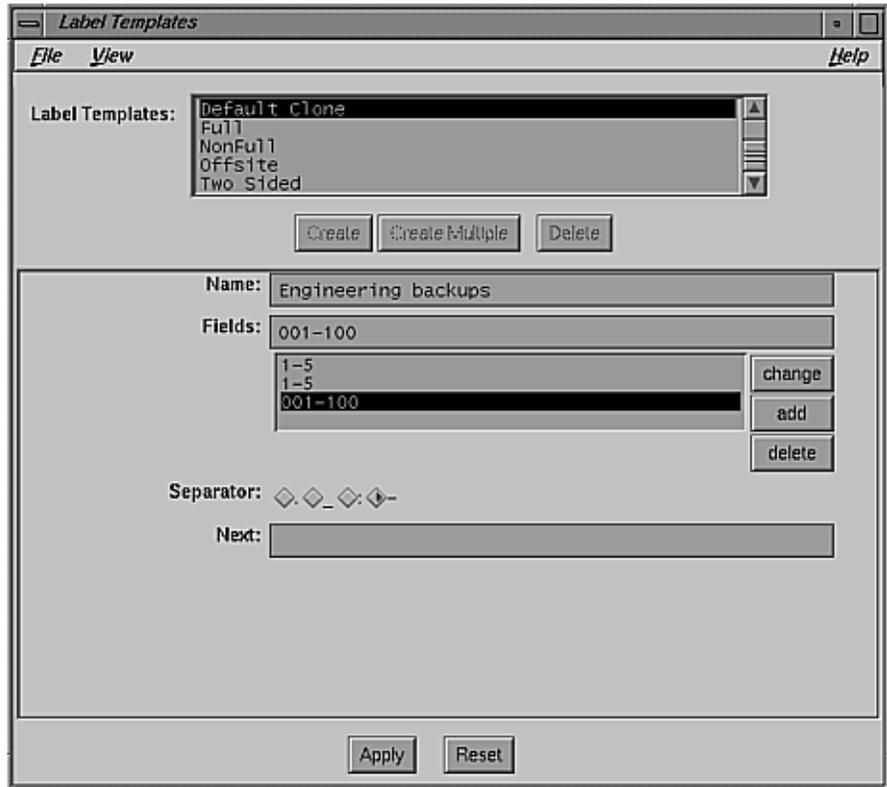


Figure 7-20 Example Label Template

Notice your new template appears in the Label templates scrolling list, the components of the label template 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 is for a sequential labeling scheme. In this scheme, there is no storage system pattern to explain the labels; the labels are simply a way for NetWorker to identify the backup volumes. This label template generates lots of different 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). Examples of labels that NetWorker generates are as follows:

aa.00, aa.01, aa.02... aa.98, aa.99,
ab.00, ab.01,... ab.99,
ac.00, ac.01,... ac.99,
...
az.00,... az.99,
ba.00,... bz.99,
...
zz.00,... zz.99,

Follow these steps to create the sequential label template:

1. Open the Label Templates window by selecting “Label Templates” from the Customize menu.
2. Click the *Create* button.
3. Enter the name of 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 dual character strings, so you enter **aa-zz** into the field.
5. Click the *add* button.
6. Enter the second component of the label template into Fields. This field contains double digits, so enter **00-99** into the field.
7. Click the *add* button.
8. Click the period (.) *Separator* button to select it as the separator for each component of the label template.
9. 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 and not worry about having to come up with 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. The label template for two-sided media has 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. Open the Label Templates window by selecting “Label Templates” from the Customize menu.
2. Click the *Create* button.
3. Enter a name for the label template in the Name field (for example, “TwoSidedLabel”).
4. Add the following components to the Fields scrolling list:

```
server_name  
001-999  
a-b
```

Note: Use a-b to distinguish between the two sides of the media.

5. Select a period for the Separator.

The Label Templates window should now look like the window in Figure 7-21.

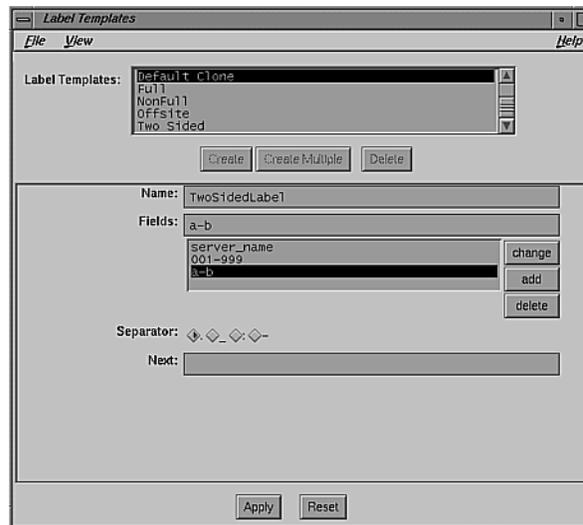


Figure 7-21 Example Two-Sided Media Label Template

6. Click *Apply* to create the label template.

In this example, NetWorker labels the first side of the optical disk as moon.001.a and the second side as moon.001.b.

Note: 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. The components are displayed from top to bottom in the Fields scrolling list, with the first component on the top.

Using Volume Pools

Volume pools establish a logical and systematic method for tracking and organizing your NetWorker backup volumes. For example, you can create a separate pool for each department in your company or a pool which contains only files with confidential information. Or you can 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.

The *volume pools* feature 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. All NetWorker volumes belong to a pool, either the Default pool or a pool you have chosen. Cloned backup volumes must belong to a *clone pool*, just as archived volumes must belong to an *archive pool*.

You can sort data into a volume pool by selecting the type of backup data to be included in the pool. NetWorker uses the choices you make in the Pools window to sort the backup data to specific backup volumes.

You can create as many pools as you need for organizing and sorting your data to backup volumes effectively. There is no limitation on the number of new pools you can create.

All NetWorker pools require a label template, as described in the preceding section “Using Label Templates.” Before creating a new pool, first create a label template with the same name.

Tip: It is strongly recommended that you keep the pool names and their corresponding label templates consistent. If you do not use the exact name, at least use names that have a logical connection.

For each pool you must select either a group, client, or save set. Unless you have data to sort from one of these choices for a pool, NetWorker does not allow you to create the pool. For example, you can sort by:

- backup group
- backup client
- list of save sets (filesystems)

- backup levels: fulls, levels 1– 9, incremental, or manual
- backups for archival only (applies only to clients that archive data)

You can further organize your backup data by

- choosing a special label template
- sending all the backup data to a particular backup device
- choosing to store entries for the files in the online index (or not, if you are sending the backup media to an off-site vault)

NetWorker backs up all the clients in a group according to a schedule, and mixes the clients' data onto backup volumes. NetWorker tracks the backed-up data and the volumes to which the data is backed up. The data, however, is essentially combined on the backup volumes.

Note: You cannot create a new pool if NetWorker is busy doing a backup. Nor can you delete a volume pool if any backup volumes still belong to that pool.

This section explains

- shortcut: creating a volume pool
- using the Pools window
- using preconfigured NetWorker volume pools
- creating a volume pool
- volume pool examples

Shortcut: Creating a Volume Pool

Use these directions if you have reviewed the examples or if you are an experienced NetWorker user.

1. Create a matching label template using the Label Templates window.
2. Open the Pools window by selecting "Pools" from the Media menu.
3. Click the *Create* button and enter the name of the new pool into the Name field.

4. Select “Yes” from the Enabled choices. (Yes is the preconfigured choice.)
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.

Using the Pools Window

Use the Pools window to create a new volume pool. To open the Pools window, select the “Pools” command from the Media pulldown menu, as shown in Figure 7-22.



Figure 7-22 Opening the Pools Window

The Pools window appears, as shown in Figure 7-23.

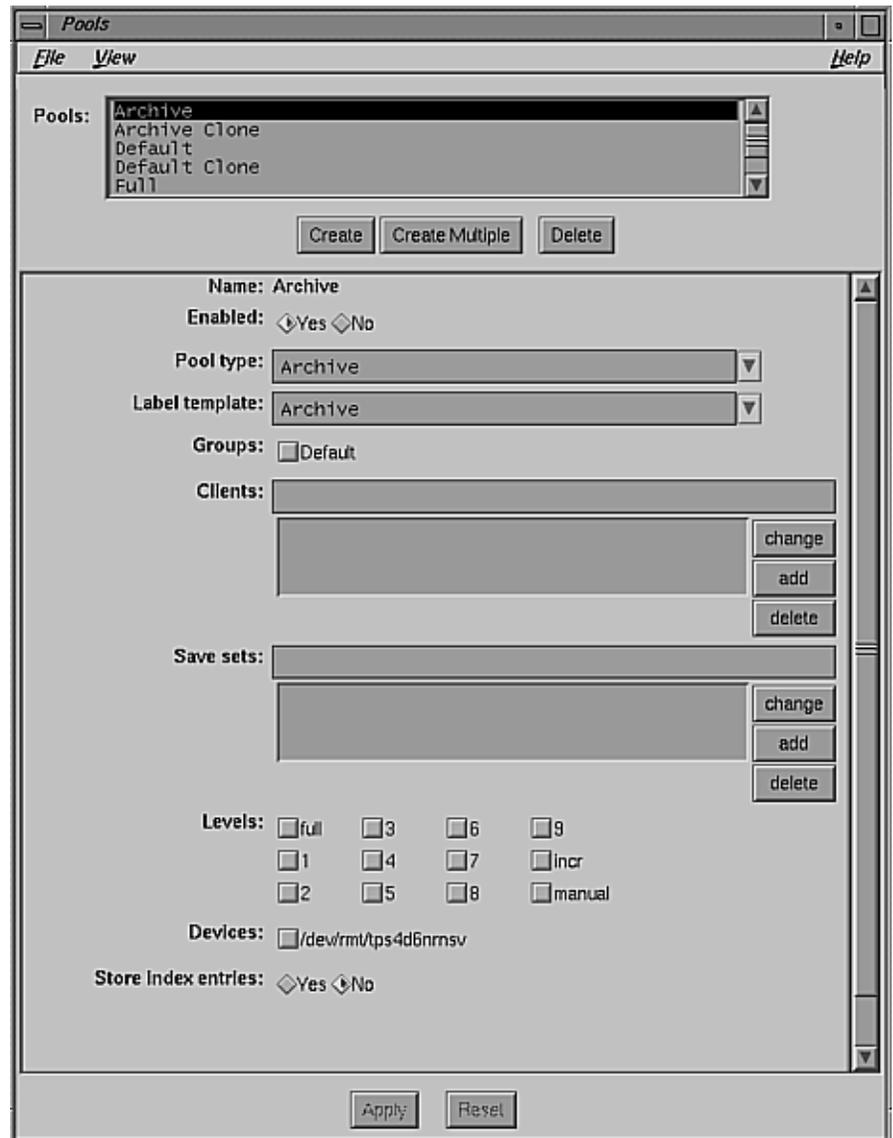


Figure 7-23 Pools Window (Expanded View)

If necessary, use the scroll bar or resize the window to view the entire Pools window.

The Pools window contains all of the choices necessary for sorting your client's data to selected backup volumes. Following is an explanation of each of the elements in the window and instructions on how to use them for modifying and creating custom pools.

You must make at least one selection from the following choices: Groups, Clients, Save sets, or Levels. If you leave a choice unselected in the Pools window, that means that NetWorker will not sort data by that choice. For example, if you do not choose a level from the Levels choices, NetWorker will include all levels of backups for that pool. NetWorker will sort data based on the other selections you have made for that pool.

Use features of the Pools window as follows:

Pools scrolling list

View all available pools, including preconfigured pools and those you have created.

Create button Click to add a new pool.

Create Multiple button

Click to create multiple pools with the same selections. For more information, see "Create Multiple Button" in Chapter 3.

Delete button Click to eliminate a pool. You cannot delete a pool as long as it still contains backup volumes; first remove all of the backup volumes that belong to the pool by using the Volume Management window. Certain preconfigured volume pools cannot be deleted.

Name field View the name of the currently selected pool. It is also where you enter the name of a new pool after pressing the *Create* button.

Enabled choices

Enable or disable the pool. If you want NetWorker to use a pool during the backup process, choose "Yes" from the Enabled choices. To exclude a pool during the backup process, select "No."

Pool type choices

Select the appropriate type of pool for your backup volumes. The four existing choices are Backup, Backup Clone, Archive, and Archive Clone.

Label template choices

Select a label template, either a preconfigured template or one you have created in the Label Templates window. It is highly recommended that you use the same name for both the label template and the selected pool.

Groups choices Choose the groups you want included in the pool. The displayed groups represent all of the groups known to the NetWorker server. There is no limitation 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 will not sort data based on a group.

Clients field

Use the Clients field for these purposes:

- Add one or more clients to a pool without selecting a group.
- Select a specific client belonging to a group. Note that if you pick a specific NetWorker client from a group, you exclude all the other clients from that group. Also note that if you have a group selected, the client you select must belong to that group.

Save sets

Use the Save sets field to:

- send a specific filesystem to a pool (see “Example 4: Specific Filesystem” in this section)
- send a specific filesystem from a specific client to a volume pool without selecting a group (see “Example 2: Secure Device” in this section)
- send a specific filesystem from a specific client of a group to a volume pool (see “Example 1: Separate Departments” in this section)

Note: Leaving a choice for Groups, Clients, Save sets, or Levels unselected in the Pools window means that NetWorker does not sort the data by that choice. For

example, if you do not choose a level from the Levels choices, NetWorker does not sort the data by levels for that pool. NetWorker sorts the data based on the other selections for that pool.

Levels choices Indicate the backup levels that are included in this pool.

Devices choices

Views the devices recognized by the NetWorker server. Choose the devices you want 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 index entries

Determine whether or not to include the index entries from the backup in the online index. If you choose “No,” the index entries are not included in the online index. For example, if you are using the preconfigured pool Offsite, you might not want the index entries online. However, an entry for the save set is still made in the media database.

Note: You cannot select “Yes” for the Backup Clone, Archive, or Archive Clone pools.

Apply button Click to apply your selections 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. Use this window for all administrative tasks relating to volume pools.

Using Preconfigured NetWorker Volume Pools

NetWorker is shipped with seven 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 seven 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 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 seven preconfigured pools are as follows:

Note: The first four preconfigured volume pools—Archive, Archive Clone, Default, and Default Clone—cannot be deleted or modified and contain no selections for you to make.

Archive	Use for archiving client data only. The preconfigured settings are Enabled: Yes; Label template: Archive; Pool type: Archive; Store Index entries: Yes.
Archive Clone	Use for cloning archive data only. The preconfigured settings are Enabled: Yes; Label template: Archive Clone; Pool type: Archive Clone; Store Index entries: No.
Default	If you decide not to use the volume pools feature, NetWorker automatically places all of your backup volumes in this pool. The preconfigured settings are Enabled: Yes; Label template: Default; Pool type: Default; Store Index entries: Yes.
Default Clone	If you decide not to use the pools feature, NetWorker automatically places all of your cloned backup volumes in this pool. The preconfigured settings are Enabled: Yes; Label template: Default Clone; Pool type: Backup Clone; Store Index entries: No.

Note: The remaining three preconfigured volume pools—Full, NonFull, and Offsite—are intended for sorting data by levels.

Full	Use this pool for full backups only. This pool separates all full backups from the incremental and level backups. The Full pool makes it easy to 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 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	Use 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 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.

Note: Enable the pools you wish to have in effect during the scheduled backups by selecting “Yes” from the Enabled choices.

Creating a New Pool

Suppose you want to create a pool named “Test” that selects all of the full backups from a group called QA. The QA 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. This is data that you do not expect to need to recover, but plan to keep for a year.

Follow these steps to create a matching label template for the volume pool:

1. Open the Label Templates window by selecting “Label Templates” from the Customize menu. See the section “Using Label Templates” earlier in this chapter for complete instructions for creating a label template.

The Label Templates window will look similar to “Creating a New Pool” after you create the Test label template.

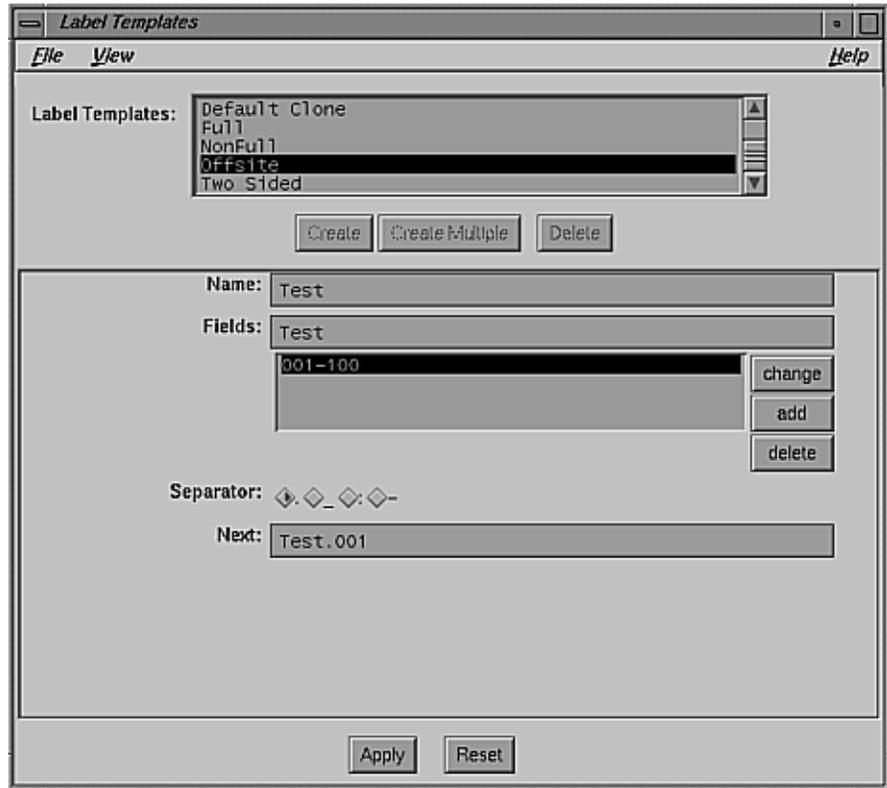


Figure 7-24 Example Label Templates Window

2. Open the Pools window by selecting "Pools" in the Media menu.
3. Click the *Create* button and enter `test` into the Name field.
Note: You cannot create a new pool if NetWorker is busy doing a backup. This ensures that NetWorker will send the backup data to the correct pool.
4. Select "Yes" from the Enabled choices. (Yes is the preconfigured choice.)
5. Select "Backup" from the Pool type choices.
6. Choose the label template Test, which you previously created.
7. Click the *Create* button and enter "Test" into the Name field.

8. Select “Yes” from the Enabled choices.
9. Choose the label template Test, which you previously created.
10. Choose the group QA, because it contains the clients you want to include in the Test pool.

Do not enter anything in the Clients or Save sets fields, because you do not need to selectively include clients or save sets.

11. Select “full” from the Levels choices.
12. Select the device to which you wish to back up from the Devices choices, if you have more than one backup device.
13. Select “No” from the Store index entries display. Selecting “No” means that you do not include the index entries in the online index, saving space on your hard disk.
14. Click the *Apply* button to apply the selections. The Pools window looks like Figure 7-25 after creating the pool “Test.”

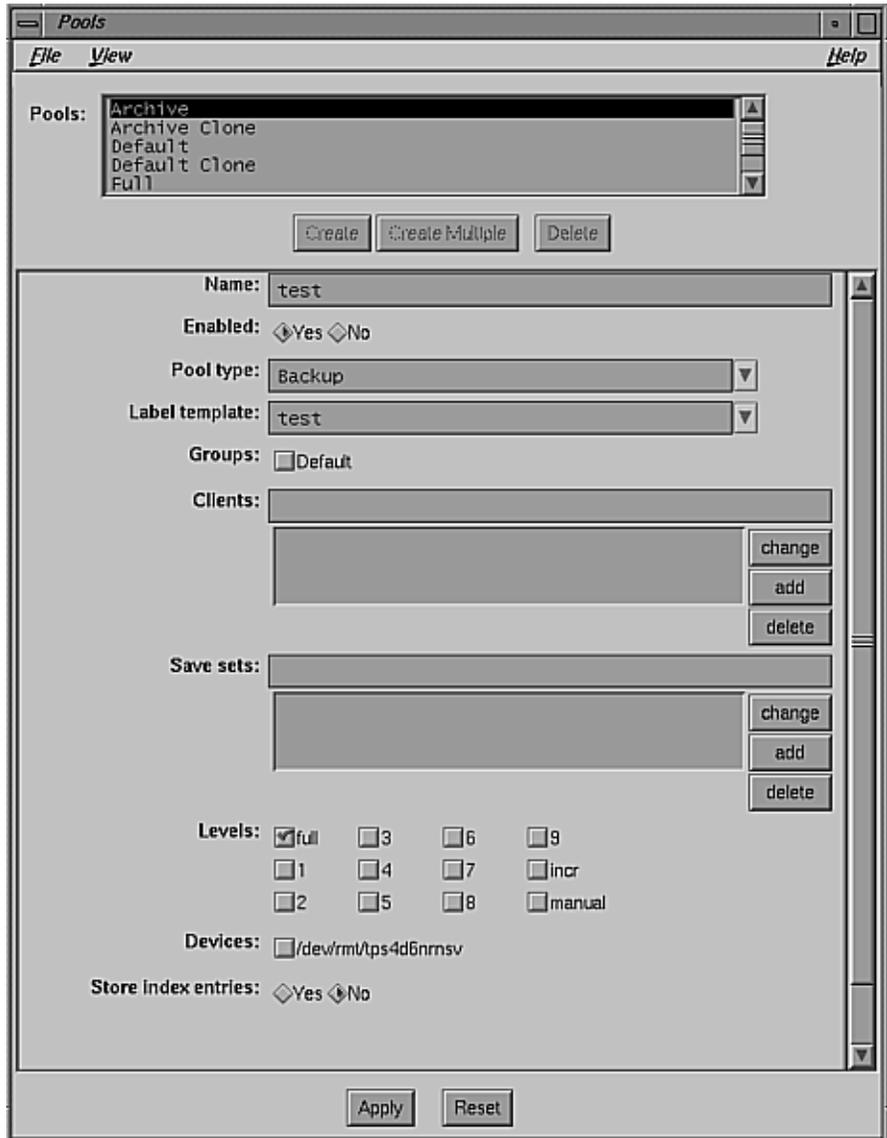


Figure 7-25 Example Pools Window

Once you create a new pool, you can easily modify it by choosing it from the Pools scrolling list and changing the current selections. When you have

completed making the modifications, click the *Apply* button. If you have changed your mind and do not want to implement the new selections, click the *Reset* button to restore the Pools window to its original configuration.

Volume Pool Examples

This section contains four examples of volume pools.

Example 1: Separate Departments

Your company needs to track department resources for accounting requirements which include the related costs of each department's backups. To sort each department's data, create a separate pool for each one. Sorting the data into different pools for each department makes it easy for Accounting to determine how many tapes are being used, which backup devices are being used, and how much time is being spent on administering each department's backup needs.

To create a pool for each department, follow these steps:

1. Create a label template with the name of the department in the Label Templates window.
2. Use the Pools window to create a new pool with the same name as the label template.
3. Select the groups that belong to each department.

Example 2: 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, just a select set of files from a small number of clients; you want to put these files into a pool named "Confidential."

To create the "Confidential" pool, follow these steps:

1. Create a label template named "Confidential."
2. Create a pool with the same name.

3. Enter the name of each client in the Clients field.
4. Enter the filesystems for each of the clients in the Save sets field.
5. Select the device they will be backed up to from the Devices choices.

Example 3: Targeted Backups

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

To back up the graphics files on the jukebox, follow these steps:

1. Create a label template named “Graphics Full.”
2. Create a new pool with the same name.
3. Choose the group(s) that need to be included.
4. Select full from the Levels choices.
5. Choose the jukebox device from the Devices choices.

Example 4: Specific Filesystem

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

To create a pool named “E-mail,” follow these steps:

1. Create a label template named “E-mail.”
2. Create a pool with the same name. You do not need to select a specific group or client because you are backing up all of the client systems’ electronic mail.
3. In the Save sets field, enter:

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

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

Adding and Deleting Backup Devices

This section explains how to use the Devices window to add and delete devices on a server. NetWorker supports multiple devices either in a series, concurrently, or in a jukebox. If your server is configured to support more than one backup device in a series, NetWorker uses the backup volume in one device until the backup volume is full and then uses the backup volume in another device until that volume is full, and so on. Using the IRIX NetWorker TurboPak option spreads the backups evenly over all the backup devices. Using a NetWorker Support for Jukebox option eases the automated handling of backup volumes for unattended backups. See “Backup Devices and Media” in Chapter 1 for a list of supported backup devices.

To view current devices or to add a new device, select “Devices” from the Media pulldown menu, as shown in Figure 7-26.



Figure 7-26 Opening the Devices Window

The Devices window appears, as shown in Figure 7-27.

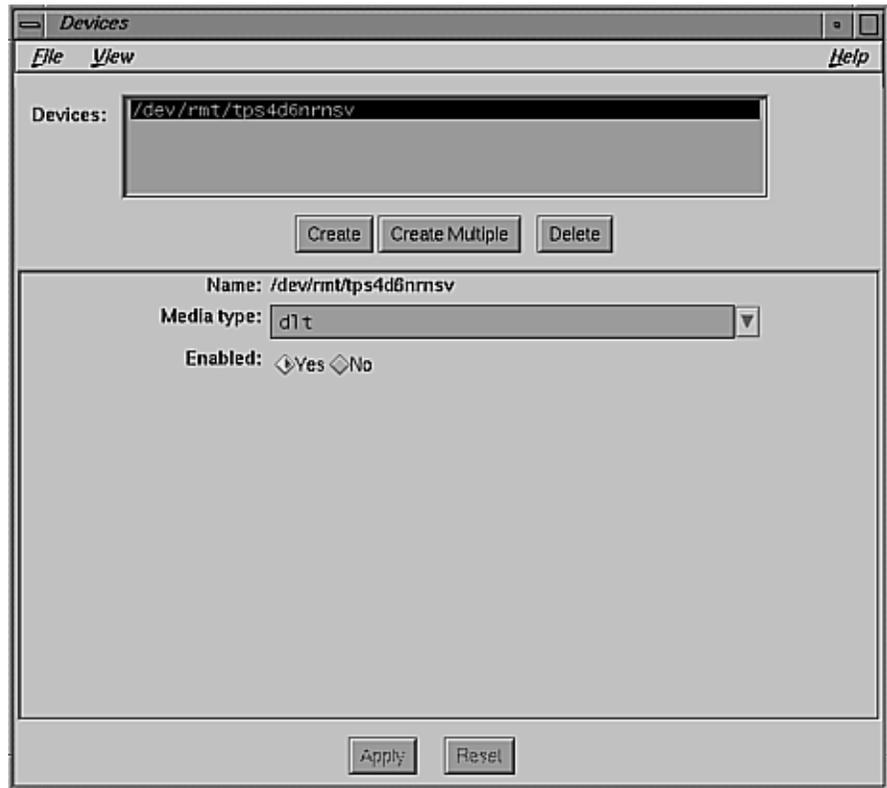


Figure 7-27 Devices Window

The features of the Devices window have the following functions:

Devices scrolling list

Displays device pathnames of the backup devices connected to the NetWorker server, standalone devices, or devices located inside a jukebox. Use these devices for backing up and recovering client files.

Name field

Displays the pathname of the currently selected device or accepts a pathname of a new device.

Media type scrolling list

Displays the media type of the currently selected device. The Media type choices are also used to select the type of new devices. The choices represent

- dlt: digital linear tape
- hint: half-inch magnetic tape drive
- qic: quarter-inch cartridge tape drive
- 4mm: four-millimeter tape drive
- 8mm: eight-millimeter tape drive
- 8mm 5GB: eight-millimeter five gigabyte tape drive
- 3480: 3480 tape drive
- vhs: VHS tape drive
- optical: optical disk drive

These choices are inactive until you click the *Create* button to add a device.

The media types displayed in the Devices window vary depending upon what type of backup devices are supported by the current NetWorker server.

Caution: A no-rewind tape backup device is required (except for optical 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.

Enabled choices

Allows you to enable (select “Yes”) or disable (select “No”) a backup device. If the backup device is not functioning, select “No” so that NetWorker does not attempt to use it for backups. When you disable a backup device, it is listed as disabled in the Devices display of the NetWorker Administrator window.

Adding a Backup Device

Before adding a device, determine this information:

- the pathname of the device
- the media type of the device (what type of media it supports)

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

1. Determine the pathname of the device you want to add (for example, */dev/rmt/tps1d6nrnsv*).
2. Determine the media type of the device you want to add. Use a no-rewind tape backup device (except for optical drives).
3. Select “Devices” from the Media pulldown menu. The Devices window appears, as shown in Figure 7-27.
4. Click the *Create* button.
5. Enter the device pathname for the new device in the Name field.
6. Select the appropriate Media type choice for the type of device.
7. Click the *Apply* button. Any additions you made appear in the Devices window.

Deleting a Backup Device

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

1. Select “Devices” from the Media pulldown menu. The Devices window appears, as shown in Figure 7-27.
2. Select the device you want to delete in the Devices scrolling list.
3. Click the *Delete* button. An “ok to delete?” confirmation box appears.
4. Click *Ok* to delete the device.

When you disable a backup device, it is listed as disabled in the Devices display of the NetWorker Administrator window.

NetWorker Operations

This chapter describes the everyday tasks for operating NetWorker. Generally, five daily tasks are required:

- checking the NetWorker server status
- monitoring and controlling backups
- monitoring backup devices
- monitoring and managing index disk space usage
- checking NetWorker notices

This chapter concludes with instructions for a management task that might need to be performed occasionally: recovering files that have been removed from the index.

Checking the NetWorker Server Status

The NetWorker Administrator window displays the status of the NetWorker server in five displays. If you do not already have the Administrator window on your screen, enter the *nwadmin(1M)* command at the prompt in a shell window to start the NetWorker Administrator program:

```
# /usr/etc/nwadmin&
```

Figure 8-1 shows the window. For more information on window features, read Chapter 3, “Using NetWorker Windows and Menus.”

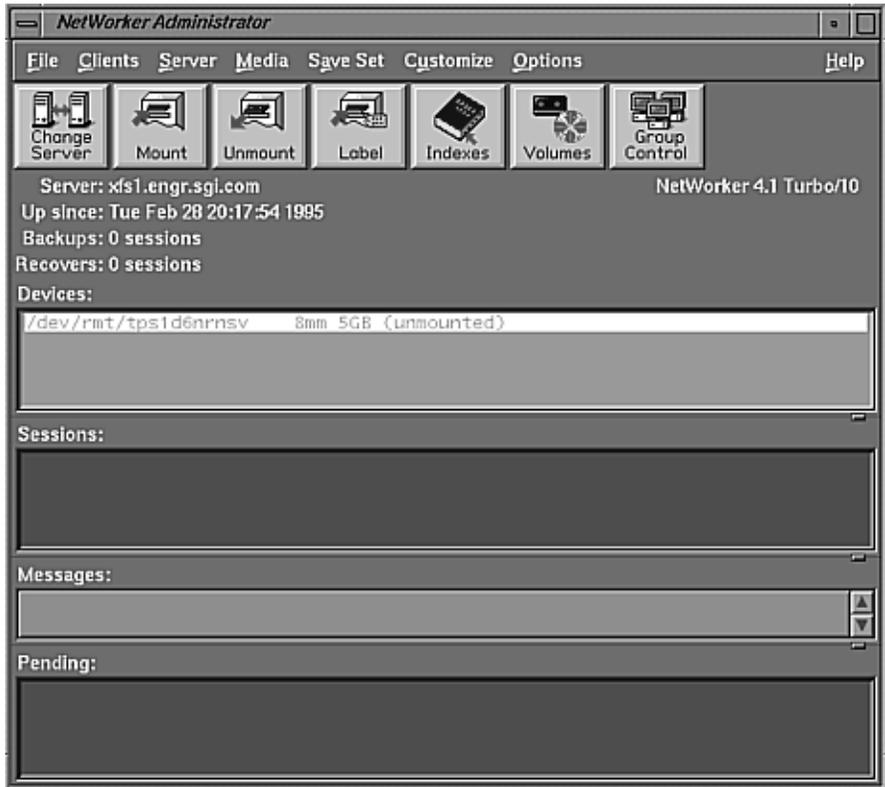


Figure 8-1 NetWorker Administrator Window

To check the status of the NetWorker server, follow these steps:

1. In the Messages display, scroll through the available messages. Error messages in this display indicate problems that should be investigated. Pull down the window to display more messages at a time, if necessary.
2. Look at the Pending display. If there are messages in this display, there may or may not be a problem:
 - If there are active save sets running (denoted by changing messages in the Messages display), NetWorker is operating normally.
 - If there are no active save sets running, there may be a problem and it should be investigated.

3. Open the Group Control window by selecting “Group Control” from the Server pulldown menu. For each group, select the group, then click the *Details* button. Messages in the Failed Save Sets display indicate problems. For more information, see “Viewing Details of a Scheduled Backup” later in this chapter.
4. Check the contents of `/var/adm/SYSLOG` for the last 24 hours. Look for
 - messages output by NetWorker as it works: check for error messages
 - error messages output by the kernel that might affect the operation of NetWorker; for example, messages about SCSI hardware errors

Monitoring and Controlling Backups

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

- 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

Open the Group Control window by selecting “Group Control” from the Server pulldown menu or by clicking the *Group Control* speedbar button, as shown in Figure 8-2.

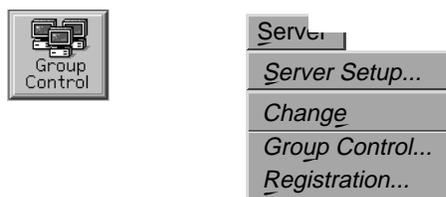


Figure 8-2 Opening the Group Control Window

Figure 8-3 shows the Group Control window.

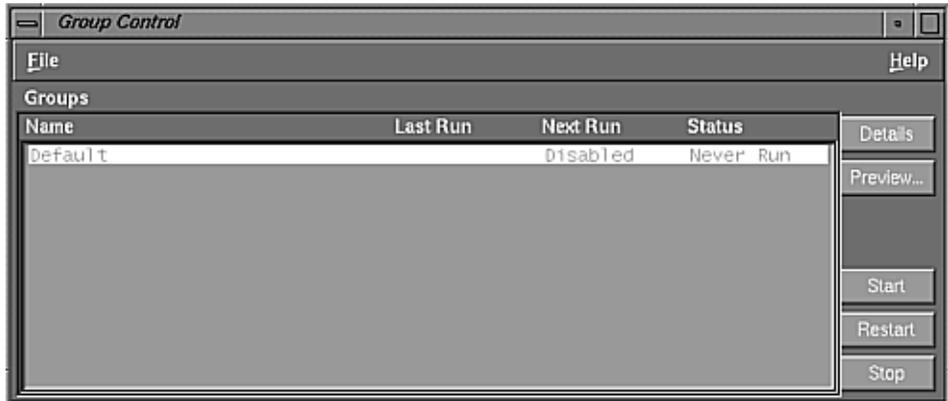


Figure 8-3 Group Control Window

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 | <p>Displays the progress of the backup. NetWorker displays these four progress messages:</p> <ul style="list-style-type: none"> • Never Run: the group has not been backed up since the last reboot or restart of NetWorker service • Finished: the backup has finished • Not Finished: the backup is still running or the backup has exited prematurely without finishing • Preview Run: the preview of a backup is running |
| <i>Start</i> button | Starts a backup immediately. |
| <i>Restart</i> button | Restarts a backup that had been stopped. |
| <i>Stop</i> button | Stops a scheduled backup. |

Details button Displays details of completed or currently running backups.

Preview button Previews an upcoming backup.

Note: Use the Groups window to create new groups and to enable and set the backup time for a group.

Viewing Details of a Scheduled Backup

To view more detailed information about a group's completed or currently running backup, click the *Details* button in the Group Control window. NetWorker displays the Group Control Details window, as shown in Figure 8-4.

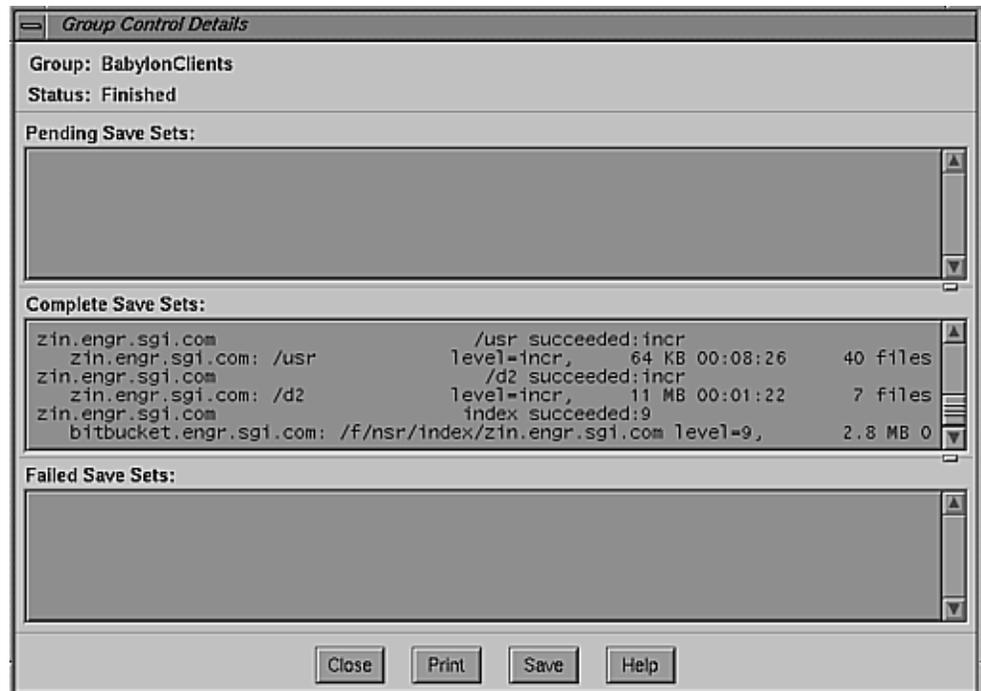


Figure 8-4 Group Control Details Window

The Group Control Details window displays save sets in the process of being backed up and save sets that have completed their backups. Use the window to determine which client's save sets backed up successfully and which save sets may have failed.

The fields in the Group Controls Details window have the following functions:

Pending Save Sets

Displays the filesystems that have not yet been backed up.

Complete Save Sets

Displays the filesystems that NetWorker has successfully backed up.

Failed Save Sets

Displays the filesystems that NetWorker was unable to back up and error messages concerning failures.

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 the backup
- the client crashed during the backup
- the network connection failed during the backup

Tip: If you previously set the Client retries to a number greater than zero in the Groups window and NetWorker fails to back up a client, filesystems might be temporarily listed in the Failed Save Sets and Pending Save Sets scrolling lists while NetWorker tries again to back up the client.

Previewing a Group Backup

You can preview a group backup by clicking the *Preview* button in the Group Control window. When you click the *Preview* button, the Group Control Details window appears, but instead of displaying past information about completed group backups, you see future information about how a group will perform during its next scheduled backup.

NetWorker simulates a backup for the group you select in the Groups scrolling list in the Group Control window. By using the preview feature, you can determine if there will be any problems with an upcoming group backup, and take the necessary steps to correct them.

You can also preview a backup group by checking if NetWorker can access the client in order to back up its files every time you add a new client to NetWorker. Use the `savegroup -p` command at the system prompt to see a preview of a group backup without actually backing up any files. For example:

```
# /usr/etc/savegroup -p group_name
```

This command shows you a backup preview of the clients assigned to the backup group `group_name`. If NetWorker cannot access a client in the backup group, an error message appears.

If you see an error message, make sure `nsrexecd(1M)` is running and that it lists the server's hostname in the command line. To make sure that `nsrexecd` is running, you can use the command `ps(1)` on the client. See the section "Installing NetWorker Software on Clients" in Chapter 2 for more information on `nsrexecd`.

Starting a Scheduled Backup Immediately

To start a backup immediately from the Group Control window, select the group you wish to back up from the Groups scrolling list and click the *Start* button.

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 immediately start a backup in the Groups window with the *Start now* button.

Stopping a Backup Group

If you need to stop a scheduled backup in progress, use the *Stop* button in the Group Control window. Figure 8-5 shows this window.

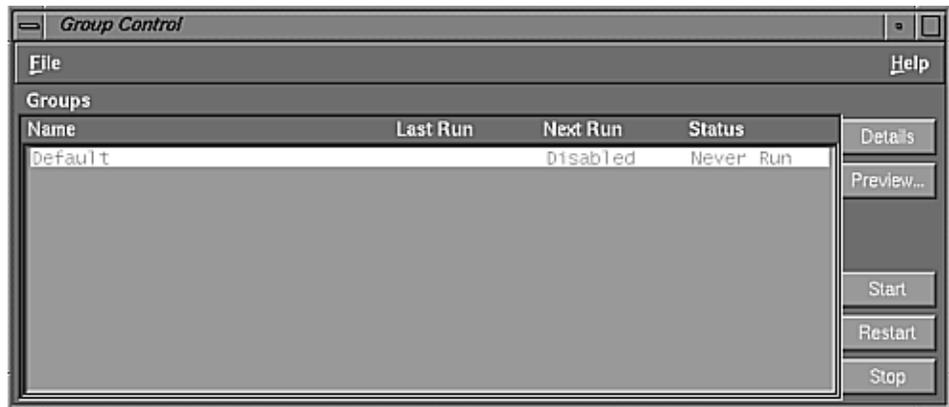


Figure 8-5 Group Control Window

To stop a backup, select the group you wish to stop from the Groups scrolling list and click the *Stop* button.

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.

The next network-wide backup starts as scheduled in the Start time field of the Groups window, or you can restart the backup by clicking the *Restart* button in the Group Control window.

Restarting a Backup Group

If you wish to restart a scheduled backup that you have stopped, use the *Restart* button in the Group Control window. To restart a backup, select the name of the halted backup group from the Groups scrolling list in the Group Control window and click the *Restart* button.

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 “notifies” that it is changing. In the Group Control Details window, the name of the file that changed during the backup is displayed in a warning message like this:

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

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

Note: You cannot recover files from a backup that was terminated when the NetWorker daemons were killed, because the media index was not updated before the daemons died. Consequently, NetWorker does not know on which volume the requested file is located.

Monitoring Backup Devices

To monitor NetWorker backup devices, use the Media menu, as shown in Figure 8-6.



Figure 8-6 Media Menu

To find out how much space is available in the jukebox (or other autochanger), use either the Jukebox Mounting window or the *nsrjb* command. The Jukebox Mounting window displays all the media in the jukebox and the percentage used of each tape, as shown in Figure 8-7.

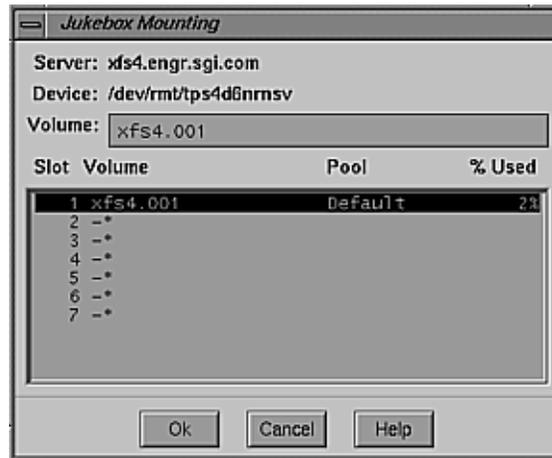


Figure 8-7 Jukebox Mounting Window

If you prefer, use the *nsrjb* command in an IRIX shell:

```
# nsrjb -jjukeboxname
```

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 full Default
2: moon.011 full Default
3: moon.012 full Default
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/tps4d6nrnsv) 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.

Monitoring and Managing Index Disk Space Usage

Every time a backup completes, NetWorker creates entries for the files saved for each client in online indexes. These indexes require disk space and must be monitored to be sure they are not getting too large. Entries can be removed from the indexes when they become too large. You can also manage your indexes automatically by selecting different index policies in the Policies window. See “Determining Browse and Retention Policies (Index Policies)” in Chapter 5 for more information.

This section explains

- viewing index information
- viewing save sets to determine resource usage
- removing the oldest cycle
- reclaiming index space

Viewing Index Information

To view the disk space usage of the indexes, click the Indexes speedbar button or select “Indexes” from the Clients pulldown menu, as shown in Figure 8-8.

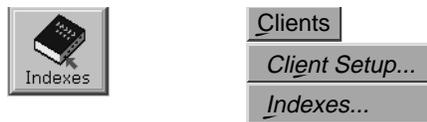


Figure 8-8 Opening the Indexes Window

The Indexes window appears, as shown in Figure 8-9.

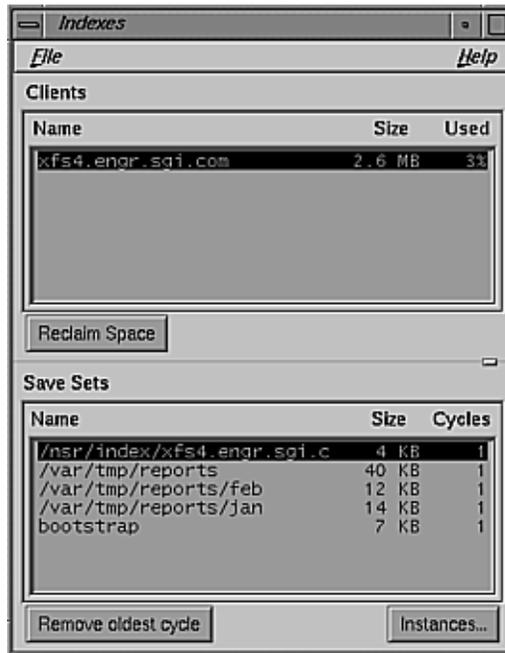


Figure 8-9 Indexes Window

Use features of this window as follows:

Clients scrolling list

View index space allotted to various clients. The Used column indicates the percentage of the index file being used. If the percent listed is 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. The allocated disk space grows automatically as the index size increases.

Reclaim Space button

Click to decrease the size of the client file index. When you use the *Reclaim Space* button, NetWorker compresses the

space in the index left by removed entries, thus freeing up disk space. A confirmation box appears once a reclaim space operation has begun.

Save Sets scrolling list

View a list of 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 name */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.

Remove oldest cycle button

Click to remove the save set's oldest cycle, if the index requires further reduction. Select the save set and 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.

For instructions on reclaiming disk space in the indexes, see "Reclaiming Index Space" later in this chapter.

Viewing Save Sets to Determine Resource Usage

To plan the amount of disk space needed for the online indexes, you can determine the size of various 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.

You can also generate a report detailing NetWorker client backups to fulfill administrative requirements.

To view the Instances window, follow these steps:

Use NetWorker's Instances window to view detailed information about save sets.

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 8-10.

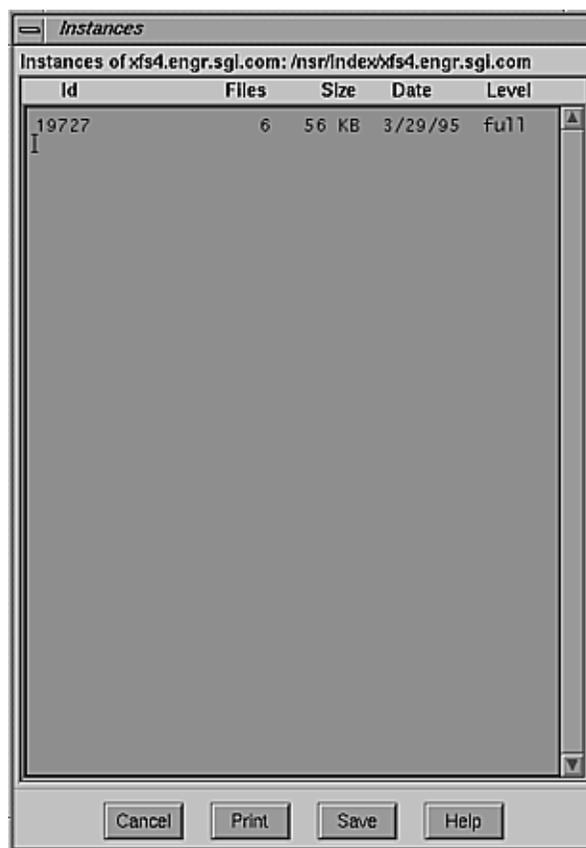


Figure 8-10 Instances Window

The features of this window have the following functions:

title bar Displays the name of the currently selected save set.

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

You can use the *Print* button to print the contents of the Instances window. To save the contents of the Instances window to a file, click the *Save* button. The information is saved to a file named *networker.rpt*, which NetWorker automatically creates, unless you choose a new filename in the dialog box that appears.

Removing the Oldest Cycle

To remove the oldest full-to-full cycle of a group of save sets from the file index for a client, follow these steps:

1. Open the Indexes window by selecting “Indexes” from the Clients pulldown 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 save set group’s oldest cycle. 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 *Cancel* to discontinue the operation.

Note: If you do not highlight any save sets 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 highlighted in the Clients scrolling list.

Figure 8-11 diagrams what happens when you remove a cycle from the index. Removing the oldest cycle opens up space so that other entries can fit into the empty space without increasing the size of the index.

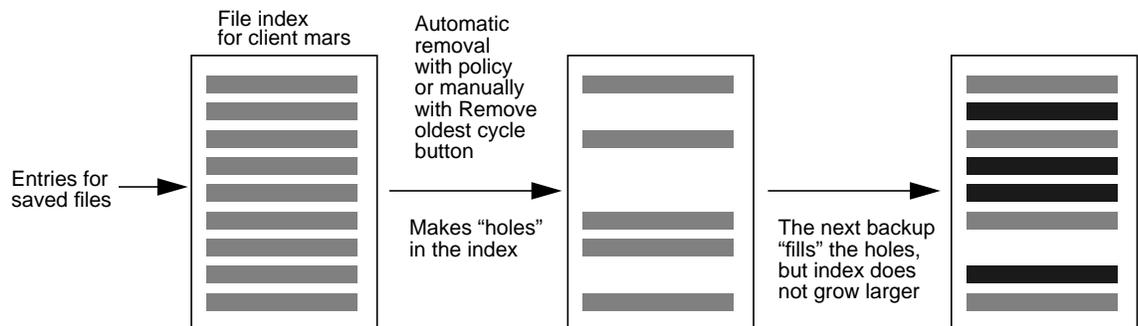


Figure 8-11 Removing a Cycle From an Index

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

After the "Remove oldest cycle" operation has finished, the statistics in the Indexes window are updated to reflect the current state of the file index.

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 these instructions to reclaim empty space in the file index created when entries are automatically or manually removed:

1. Open the Indexes window by selecting "Indexes" from the Clients pulldown 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 confirmation box appears, as shown in Figure 8-12.



Figure 8-12 Reclaim Space Confirmation Box

3. Click the *Ok* button to start the reclaim space operation. The Progress bar moves to show the progress of the index operation. When an index is busy, the *Ok* button is disabled.

The Reclaim Space confirmation box changes, as shown in Figure 8-13.



Figure 8-13 Reclaiming Space

4. Click the *Cancel* button to dismiss the dialog box. This action closes the window; however, 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; consequently, it shrinks in size and opens up space for new entries. The space is returned to the filesystem.

Figure 8-14 diagrams what happens when you use the *Reclaim Space* button.

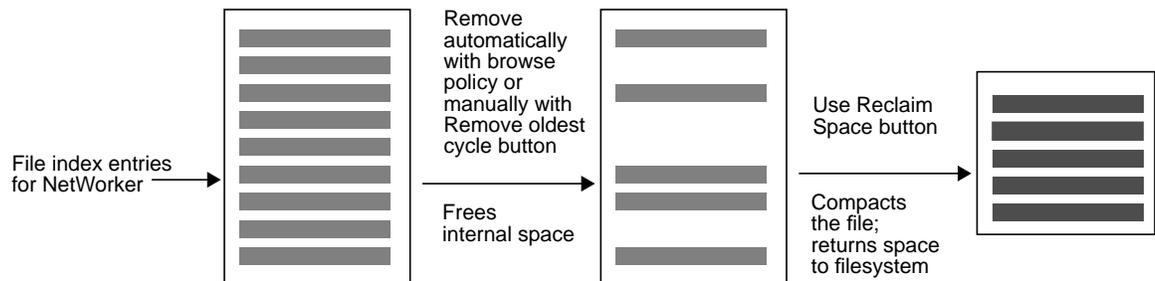


Figure 8-14 Reclaiming Index Space

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

After the reclaim space operation has finished, the statistics in the Indexes window are updated to reflect the current state of the file index.

Checking NetWorker Notices

Checking NetWorker notices is an important daily task for ensuring the smooth operation of NetWorker. NetWorker sends notices about significant NetWorker events. These notices include

- messages requesting backup volumes to be mounted
- nightly backup completion notices
- messages concerning the size of the online indexes
- jukebox request messages
- registration status

NetWorker is shipped with preconfigured notification instructions to send electronic mail to root about most events. You can check NetWorker notices three ways:

- Log in as root, and read your electronic mail.
- Look at the system console messages.
- Use the NetWorker Notifications window.

Using the Notifications Window

To open the Notifications window, select “Notifications” from the Customize pulldown menu, as shown in Figure 8-15.

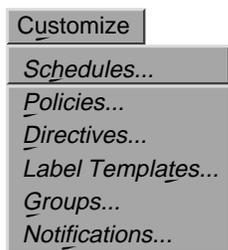


Figure 8-15 Opening the Notifications Window

The Notifications window appears, as shown in Figure 8-16.

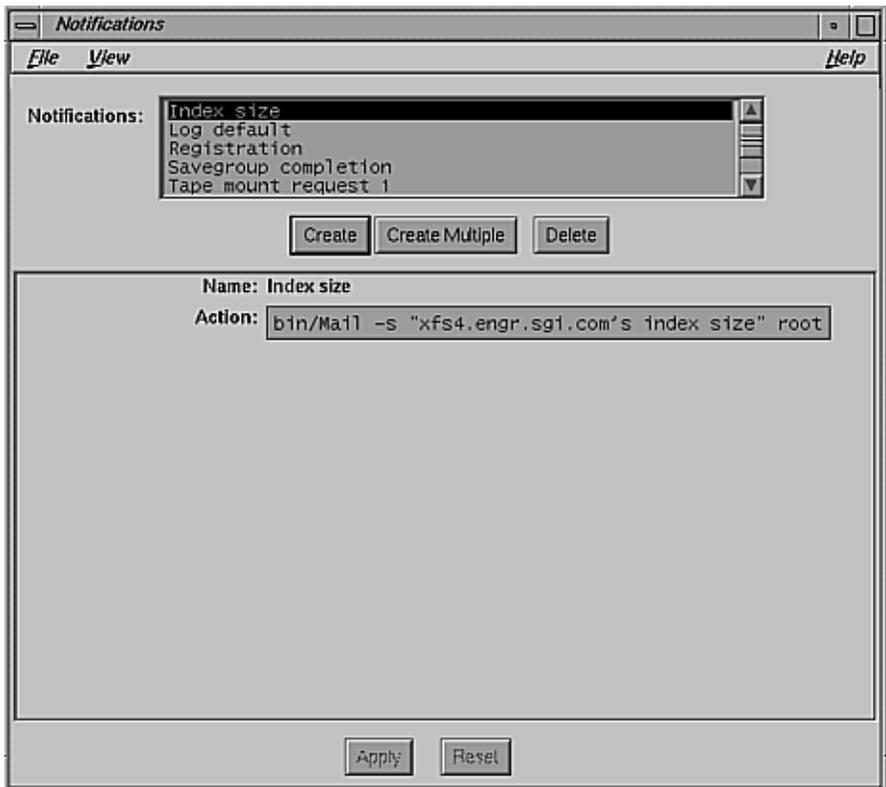


Figure 8-16 Notifications Window

The Notifications scrolling list displays the names of all the notifications known to the NetWorker server, including NetWorker preconfigured notifications. Information about the currently selected notification is displayed in the lower half of the window.

Name field Displays the name of the currently selected event notification.

Action field Displays the type of action NetWorker takes with the notice. For example, NetWorker can use electronic mail to notify individuals that a backup volume needs to be mounted in the server. You can change the Action field to customize how the notifications are posted.

Using NetWorker Preconfigured Notifications

NetWorker includes seven preconfigured types of notification: Registration, Log default, Index size, Savegroup completion, Tape mount request 1, Tape mount request 2, and Tape mount request 3. If you are using a jukebox with NetWorker, you also receive a Jukebox request notification.

Most of these notices alert you about important NetWorker events. For example, if a group of clients did not complete a nightly backup, a savegroup completion notice appears in your electronic mail.

Registration The registration notification sends a message to root that your NetWorker products are not properly registered. The registration notification appears once a day or each time you start NetWorker. The notification message includes related information about each of the NetWorker products that are not licensed correctly.

Log Default The log default notification uses an IRIX 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 in the *syslog* configuration file for logging and contacting operators; you can customize these entries. See the *syslogd(1M)* reference page for information on configuring the distribution of log information.

Index Size NetWorker checks the size of its online indexes and sends a notification if it looks as if the indexes may run out of disk space. NetWorker automatically sends the electronic mail message to root. If you want the message to be mailed to someone other than root, edit Action in the Notification

window, substituting a different user login name or a mailing list for root. To reduce the size of the indexes, see “Removing the Oldest Cycle” earlier in this chapter.

Savegroup Completion

When NetWorker finishes backing up a group of clients, it sends a completion message via electronic mail to root.

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 15 minutes. NetWorker sends a third request after another 37 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 *syslog*, and the third request sends electronic mail messages to root.

Jukebox Request

NetWorker sends a jukebox request when a backup volume is required.

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. The recover process is not nearly as convenient as when the entries are still in the online index, however.

If you do not want to rebuild the index, and the save sets you need are still in the media index, and 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*.

To rebuild the file index so that you can browse for the file you lost, follow this procedure:

1. Use the Volume Management window to find out 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 takes some time.
4. Use the NetWorker Recover window to identify the needed file(s) and initiate the recovery.

Note: Recovery is considerably easier if the file information is still in NetWorker's online index. Set a browse policy that is long enough to cover almost all recovery requests.

Recovering and Cloning Save Sets

Save sets are a set of files or a filesystem that you define for NetWorker backup. This chapter explains recovering and cloning save sets:

- recover save sets to retrieve data whose index entries have expired and to quickly restore large amounts of data
- clone save sets to make identical copies of selected data for storing off-site, send data to another location, or verify backed-up data

Note: Besides cloning individual save sets, NetWorker also clones entire backup volumes for offsite storage. See “Cloning Backup Volumes” in Chapter 7 for instructions.

Shortcuts

Use these shortcuts if you are an experienced NetWorker user, or after you have reviewed the examples in this chapter.

Recovering Save Sets

Follow these steps:

1. Select “Recover” from the Save Set menu. The Save Set Recover window appears.
2. Enter or select the client system whose save sets you wish to recover.
3. Select the versions of the save sets you wish to recover from the Instances scrolling list. You 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 window appears. If you need to relocate the files, click the *Options* button, and enter a path 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 enter nothing, the entire save set is recovered.
5. Click the *Recover* button to begin recovery. Unless you chose to relocate the recovered save set, the data is restored to its original location. You are notified if any file conflicts occur.

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

Cloning Save Sets

Follow these steps:

1. Select “Clone” from the Save Set menu. The Save Set Clone window appears.
2. Enter or select the criteria you wish to use for browsing save sets.
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 wish to clone from the Save Set scrolling list.
5. Click the *Clone* button. The Save Set Status Clone window appears.
6. Click the *Start* button in the Save Set Status Clone window.

Tip: Automatic cloning is accomplished in the Groups window. Select “Yes” to clone a group’s data, and enter or select the name of the Clone pool.

Recovering Save Sets

Index entries that have been removed from the online file index have passed their browse period. Consequently, their data is no longer available for recovery through the NetWorker Recover window. However, if the media where the files were originally backed up have not been relabeled, the data might still be available for recovery. If the file index entries have expired, you can use the Save Set Recover window to recover the data from the media.

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

Note these advantages of recovering save sets:

- reducing index entries

Files recovered with the save set recovery feature do not include the file's index entries. Consequently, the file index is not repopulated with old entries. Thus the amount of disk space required to maintain the online file index is reduced. NetWorker generates new index entries the next time it backs up the recovered save sets.

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 recovery feature. By removing index entries sooner, you free more disk space and keep your online indexes from growing too large.

- preserving access control lists

NetWorker preserves your operating system's 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.

Note: 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 unbrowsable 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 retention policy. The browse policy determines how long index entries for backed-up files will remain available in the file index. The retention policy determines how long a file will be tracked in the media index and recoverable from the backup media. For more information on browse and retention policies, see “Determining Browse and Retention Policies (Index Policies)” in Chapter 5.

Tip: The *scanner* command also recovers data from the media whose entries have been removed from the file index. The *scanner* command gives you the choice of scanning in data only, data and index entries, or just index entries. For more information on the *scanner* command, refer to the reference page.

- recovering large amounts of data

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

System administrators who are familiar with their system’s disks can easily determine which save sets they need in order 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 when it backs up data. Consequently, often only a single save set is required to restore a disk partition. Depending upon the setup of your disk, you might 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.

Note: For more information on the NetWorker Recover program, refer to Chapter 4 in the User’s Guide.

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

This section explains

- using the save set recover windows
- examples

Using the Save Set Recover Windows

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

To open the Save Set Recover window, select the “Recover” command from the Save Set menu, as shown in Figure 9-1.

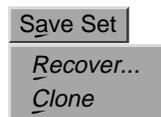


Figure 9-1 Opening the Save Set Window

The Save Set Recover window appears, as shown in Figure 9-2.

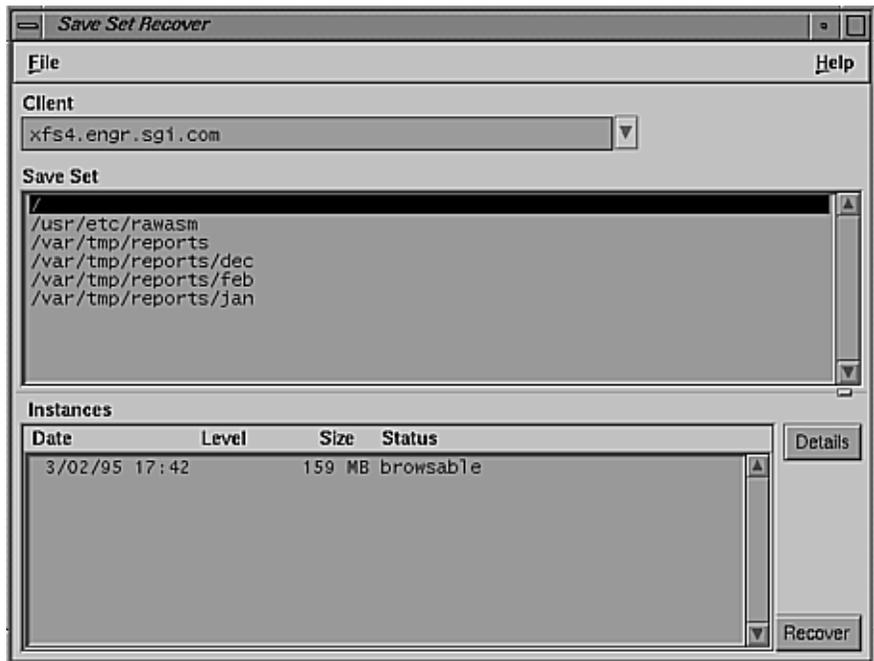


Figure 9-2 Save Set Recover Window

Use features of the Save Set Recover window as follows:

Client field Enter the name of the client system for which you wish 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 of the client's save sets.

Save Set scrolling list View the names of the selected client's save sets. Click one save set at a time to see its versions in the Instances scrolling list.

Instances scrolling list View a list of 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 once again. The seven save set status variations are as follows:

- browsable: The save set still has an entry in the online file index.
- recoverable: The entry for the save set has been removed from the client file index, but is still recoverable from the media.
- 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.
- 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 failed.

You can select an unlimited number of versions for 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.

- | | |
|-----------------------|---|
| <i>Details</i> button | Click to open the Save Set Recover Details window for more detailed information on a specific save set version. Details for the last save set selected are displayed. |
| <i>Recover</i> button | Click to open the Save Set Recover Status window to begin the save set recover process. |

Save Set Recover Details Window

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

Figure 9-3 shows the Save Set Recover Details window.

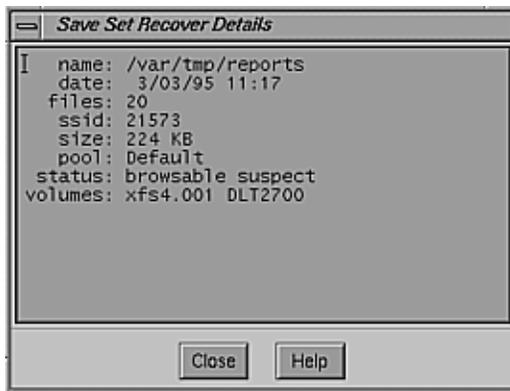


Figure 9-3 Save Set Recover Details Window

Features of the Save Set Recover Details window are display:

- name Save set name.
- date Day and time the save set was backed up.
- files Number of files in the save set. The number could be zero if the save set was an incremental backup and the files did not change since the last backup.
- ssid Internal number assigned to the save set by NetWorker for purposes of tracking 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.
- volume Backup volume name and its location, if it has one.

Save Set Recover Status Window

The Save Set Recover Status window allows you to customize save set recoveries further. Use it to select individual files rather than whole save sets and to determine where and how you will locate the recovered data. To open the Save Set Recover Status window, click the *Recover* button in the Save Set Recover window.

Figure 9-4 shows the Save Set Recover Status window.

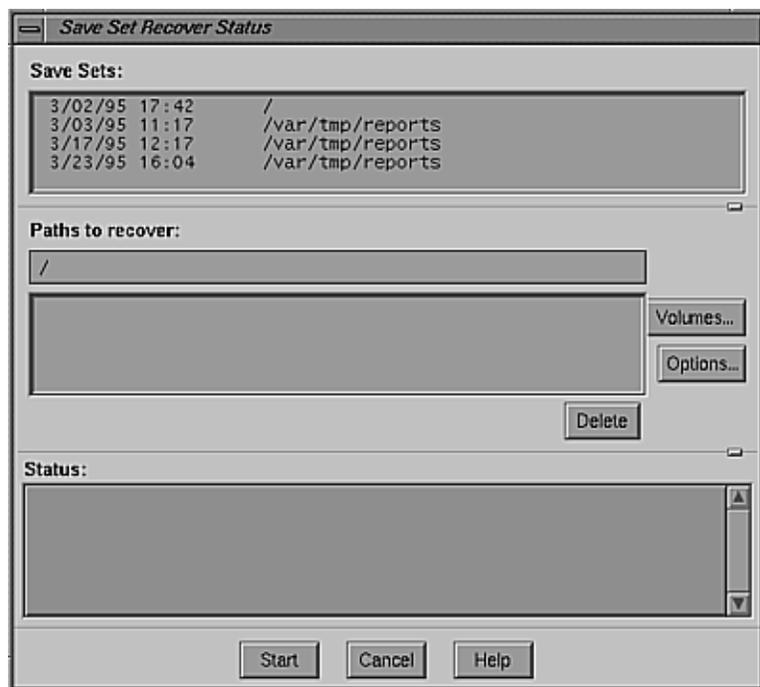


Figure 9-4 Save Set Recover Status Window

Use features of the Save Set Recover Status window as follows:

Save Sets scrolling list

View a list of save sets selected for recovery. After you have opened the Recover Save Set Status window, you can

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 pathnames. 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.

For example:

The save set pathname in the Save Sets field is */usr/src/release/nsr/4.0.4*. 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:

```
/usr/src/release/nsr/4.0.4/help
```

Press <Enter>. NetWorker recovers only the file *help*.

For example:

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>. 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.

Delete button Click to delete a file or directory you have selected from the Paths to recover scrolling list.

Status scrolling list

View messages that report on the status of the recovery. The Messages display of the NetWorker Administrator window displays the same information.

Recover button Click to start a recovery

Volumes button Click to open the Save Set Recover Volumes Required window to determine which backup volumes are needed to recover the data. All of the volumes required for all of the

save sets listed in the save sets scrolling list will be displayed. This information may be useful especially if you discover that the backup volume is located offsite, and need to make special arrangements to retrieve it. Figure 9-5 shows this window.

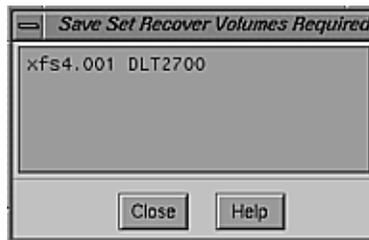


Figure 9-5 Save Set Recover Volumes Required Window

Options button Click to open the Save Set Recover Options window. Use this window to determine where and how you will recover data. Figure 9-6 shows this window.

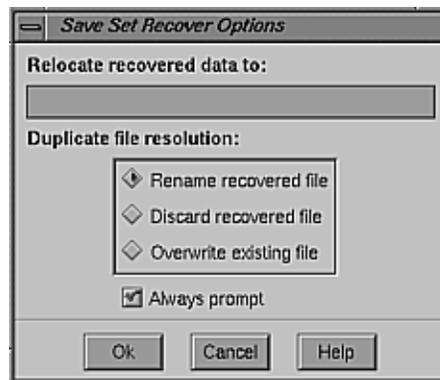


Figure 9-6 Save Set Recover Options Window

Use features of the Save Set Recover Options window as follows:

Relocate recovered data to field

Enter the full pathname of the directory where you wish 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 an item 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. 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 to be prompted when there is a filename conflict, and turn off if you do not wish to be prompted 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 contains two examples of how and why to 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.

The system administrator follows these steps to recover the needed files:

1. Open the Save Set Recover window by selecting “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.
Note: If the system administrator cannot find the client that needs to have files recovered in the Client scrolling list, 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 window.

The system administrator continues the recovery process by opening the Save Set Recover Status window.

1. Click the *Recover* button. The Save Set Recover Status window 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 <Enter> to add the file to the Paths to recover scrolling list.
3. Click the *Volumes* button to determine the backup volumes needed for recovery.
4. Click the *Options* button. Enter the full pathname of the user’s home directory 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 files still existing 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 window.

Note: 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 the recovery process by clicking the *Ok* 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.

Example 2: Recovering a Disk Partition

An engineering company has 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 has been reached over the terms of the contract. The project has been rescheduled, and the engineering team now needs to retrieve the data so it can finish developing the project software.

The MIS personnel 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 selecting “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 that contains the original project software from the Save Set scrolling list. Different versions of the save set appear in the Instances scrolling list.

Tip: Typically, a save set is comprised of data from a single disk partition because NetWorker does not automatically cross over mount points when it backs 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 window appears. Click *Ok* to close the window.

The MIS personnel continue the recovery process by opening the Save Set Recover Status window and following these steps:

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 which volume(s) is needed for recovery. The Volumes Required for Recovery window appears. Click *Ok* to close the window.
3. Click the *Options* button. The Save Set Recover Options window appears.
4. Enter the full pathname of the disk partition in the Relocate recovered data to field where the files will be recovered. Enter nothing 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 window.

8. Begin the recovery process by clicking the *Ok* 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.

Cloning Save Sets

Cloning data provides additional security for backing up and recovering data across your network by creating identical copies of your backed-up files. You might want identical copies of your data for storing off-site, or for sharing valuable engineering data with another company location. You can automatically clone save sets immediately after a backup completes, or manually on an as-needed basis. Both procedures are explained in this section.

NetWorker assures the safety of your data by copying each clone of the same save set to a different backup volume. The original backup volume is the *source volume*, and the volume to which you clone the data is the *destination volume*. NetWorker individually tracks cloned save sets and their destination volumes and marks them as cloned in the media index. More specifically, NetWorker indicates the cloned data in the windows displaying volume and save set information, which are the Volume Management, Save Set Recover, and Save Set Clone windows.

Note these features of cloning save sets:

- volume sizes for cloning save sets

Source and destination media can be of different types and capacities, and cloned save sets can span destination volumes just as they span source volumes. Save sets that fit on one backup volume might spill over to two backup volumes because the destination volume has a different capacity. However, the opposite is also possible: for the original save sets requiring two volumes, the cloned save sets might require only one volume.

- retention policies for cloned save sets

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

Volume Management window. When you apply the man mode to a volume, you must change a volume manually to the recyclable mode before NetWorker can overwrite the backup volume. This procedure is explained in this section.

- timing the cloning operation

To help minimize network traffic, NetWorker clones save sets after a backup is complete. You should take into account the time to complete the cloning process before scheduling another backup. Providing adequate time to complete either backing up or cloning of data helps reduce the amount of network traffic.

Typically, it takes the same amount of time to back up and to clone data. The cloning process might take less time 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 can 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. 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 3/4 hours.

- index entries for cloned save sets

NetWorker does not create an additional set of index entries for either online index. Both the original and cloned volumes use the same set of file and media index entries. Consequently, if you eliminate either the cloned or original volume from the media index, NetWorker does not let you remove the index entries as long as one of the volumes still exists.

- using cloning to verify data integrity

You can use data cloning to help verify the integrity of backed-up data. After the cloning process has completed, a message informs you which save sets could not be read, if any. If NetWorker performed an automatic clone, a message would appear in the savegroup completion e-mail. If NetWorker performed a manual clone, a message would appear in the Status portion of the Save Set Clone Status window.

Note: For more information on cloning entire backup volumes, see “Cloning Backup Volumes” in Chapter 7.

This section explains

- requirements for cloning save sets
- using clone pools
- using the save set clone windows
- using automatic cloning
- examples

Requirements for Cloning Save Sets

Cloning save sets requires the following:

- NetWorker release 4.1.1 or greater (this software allows you to clone save sets created prior to release 4.1.1)
- NetWorker TurboPak
- multiple backup devices
- source and destination volumes
- clone pools

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

The default NetWorker pool type for cloned volumes is the preconfigured Backup Clone pool. The Backup Clone pool is listed as one of the 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 Backup Clone pool. Use descriptive names for the clone pools you create to help identify the data they contain.

Note: For complete instructions on creating volume pools, see “Using Volume Pools” in Chapter 7.

Using the Save Set Clone Windows

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

Note: Use the Save Set Clone window to clone save sets manually. Use the Groups window to enable automatic cloning, as explained later in this chapter.

Use the Save Set Clone windows for browsing and selecting individual save sets for cloning. To open the Save Set Clone window select “Clone” from the Save Set menu, as shown in Figure 9-7.

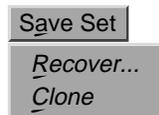


Figure 9-7 Opening the Save Set Clone Window

Figure 9-8 shows the Save Set Clone window.

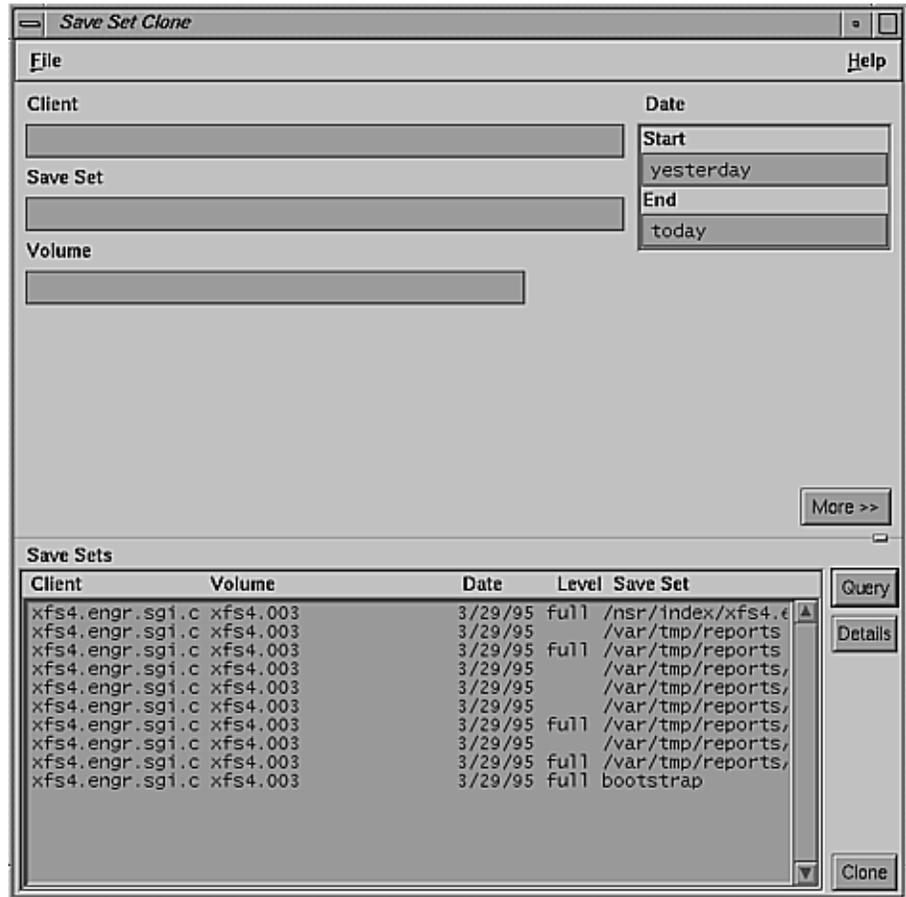


Figure 9-8 Save Set Clone Window

The upper portion of the window contains fields for selecting the save sets you wish to browse. You can fine-tune the saves sets you select by entering information in as many fields as necessary. The more fields you use for selecting save sets for browsing, the faster NetWorker locates and displays the selected save sets. If you do not enter specific values in the fields, NetWorker displays all of the client save sets.

The lower portion of the window displays the save sets selected for browsing. This portion of the window also contains buttons for initiating the browsing process, for viewing details, and for cloning save sets.

Use the features of the Save Set Clone window as follows:

Client field Enter the client name whose save sets you wish 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 wish to browse.

Volume field Enter the name of the volume whose save sets you wish to browse.

Date field Enter the range of dates on which the save set was created. By default NetWorker has yesterday's date in the Start field, and today's date in the End field. Enter the date on which you want to begin the search in the Start field and the date on which you want to end the search in the End field. You can enter the dates as follows:

- completely written out: March 1, 1994
- numerical format: 3/1/94
- phrase: for example, two weeks ago, three days ago

Tip: Open the Volume Management window to locate the 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 "Viewing Additional Browsing Choices" in this chapter for more information.

Query button Click to display the save sets chosen for browsing in the Save Sets scrolling list.

Note: The save sets selected for browsing will not appear in the Save Sets scrolling list until you click the *Query* button.

Save Sets scrolling list

View the save sets currently selected for browsing, including the Client name, Volume name, Date, Level, and Save Set name. To select a save set, click on it once. To deselect a save set version, click on it once 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 have selected in the Save Set scrolling list. The Save Set Clone Status window appears. You can clone an unlimited number of save sets at the same time.

Using the Save Set Clone Details Window

Use the Save Set Clone Details window to see 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. The information in the Save Set Clone Details window helps you better prepare for cloning selected save sets.

Figure 9-9 shows the Save Set Clone Details window.

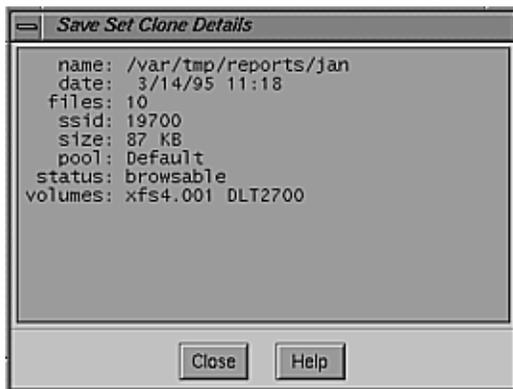


Figure 9-9 Save Set Clone Details Window

The features of the Save Set Clone Details window display:

- name** Save set name.
- date** Time and day the save set was backed up.
- files** Number of files in the save set.

Note: 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 assigned to the save set by NetWorker for purposes of tracking 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, which also indicates if a save set is a clone or has clones.
volume	Backup volume name and its location.

Using the Save Set Clone Status Window

The Save Set Clone Status window displays the save set selected for cloning. Use this window to enter the name of the clone pool to which the data will be cloned and to view the progress of the cloning operation.

Note: Volumes used for cloned data must belong to a clone pool. Use either the preconfigured Backup Clone pool, or a clone pool you have created yourself.

Figure 9-10 shows the Save Set Clone Status window.

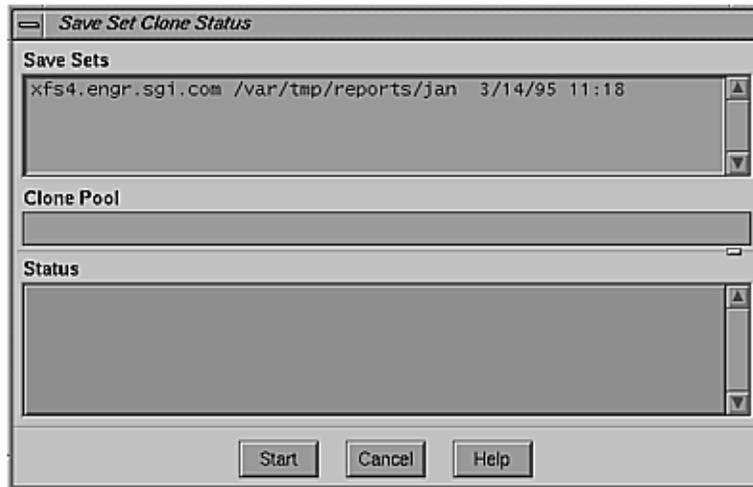


Figure 9-10 Save Set Clone Status Window

Use the features of the Save Set Clone Status window as follows:

Save Sets scrolling list¹

View save sets selected for cloning.

Clone Pool field

Enter the name of the clone pool. If a pool is not entered, NetWorker uses the default Backup Clone pool.

Status scrolling list

View the status of the cloning operation.

Start button

Click to start the cloning operation.

Viewing Additional Browsing Choices

To view additional browsing choices, click the *More* button in the Save Set Clone window. Figure 9-11 shows the additional fields.

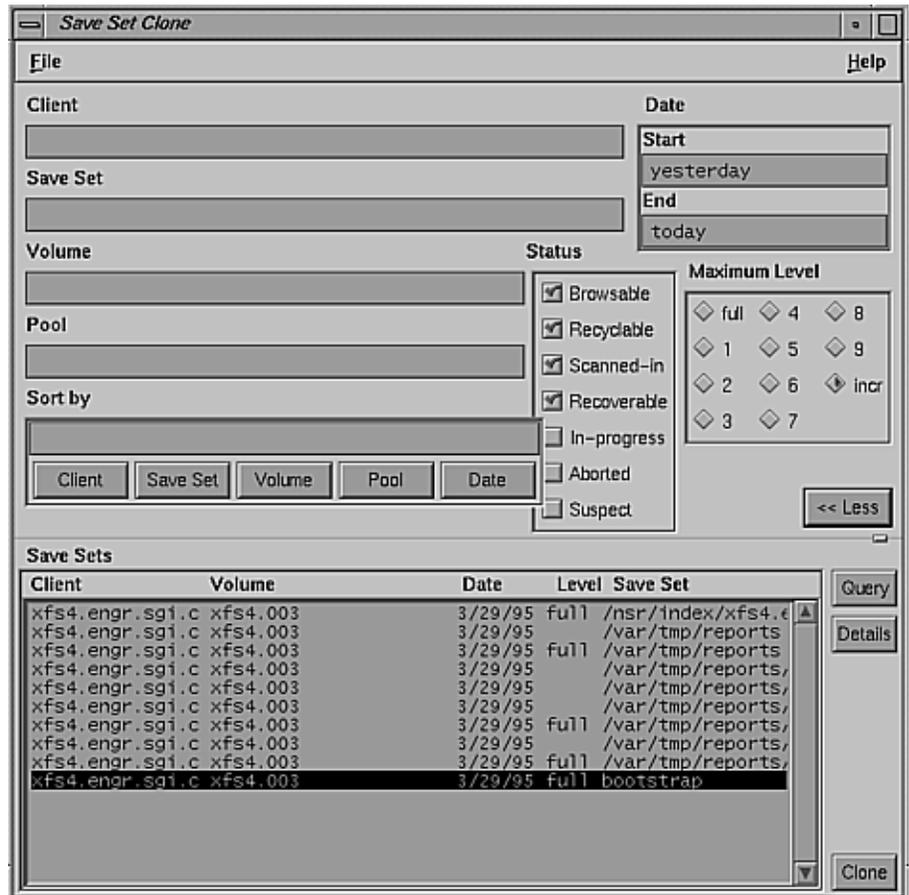


Figure 9-11 Additional Fields in the Save Set Clone Window

Use these additional choices for selecting save sets for browsing as follows:

- Pool field** Enter the name of the pool whose save sets you wish to browse.
- Status choices** Choose the status of the save sets you wish to browse. By default Browsable, Recyclable, Scanned-in, and Recoverable are already chosen. The seven save set status variations are as follows:

- **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.
- **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 still available for recovery 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 the system crashed.
- **Suspect:** A previous recovery attempt failed.

Note: 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's 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 wish 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 level 5, NetWorker browses and displays the save set levels full, one, two, three, four, and five. If you select full, NetWorker searches for and displays only level fulls. You must always have at least one level selected, but you cannot select more than one level at a time.

Sort by field

Click one or more of these buttons 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 9-12 shows the buttons.



Figure 9-12 Sort by Buttons

Try entering information in the different fields, clicking buttons, and selecting save sets to become familiar with the Save Set Clone window.

Using Automatic Cloning

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

Tip: Automatic cloning is accomplished in the Groups window; see “Configuring Backup Groups” in Chapter 6 for instructions. Select “Yes” to clone a group’s data, and enter or select the name of the Clone pool. “Example 3: Automatically Clone Several Clients” later in this chapter gives example instructions for automatic cloning.

Consider these suggestions for setting up your groups with the aid of different backup schedules for accomplishing 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 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.

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

Examples

This section contains three examples of how and why to clone save sets.

Example 1: Manually Cloning Save Sets

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

The user needs to send the data to another company location out of state that also uses NetWorker to back up their network data. The NetWorker administrator needs to clone the most recent full backup and any incrementals since the last full to make sure that the most current data is sent.

Note: The company location receiving the cloned data must use the *scanner* command to retrieve the data. See the reference page for more information about the *scanner* command.

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

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

The NetWorker administrator follows these instructions to clone the individual save sets:

1. Open the Save Set Clone window by selecting “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 for which you wish to search in the Date field.

By default **Browsable**, **Recoverable**, **Recyclable**, and **Scanned-in** are already selected from the Status choices. There is no need to change or select different status choices.

After selecting criteria for browsing, the NetWorker administrator sorts, views details, and selects the save sets for cloning.

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 window.
2. To make sure that the correct dates have been selected, sort the save sets by date by clicking the *Date* button in the Sort by field.
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 window 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 window. The Status field displays the progress of the cloning process.

The Save Set Clone window is useful for occasionally cloning individual save sets or for cloning a small number of save sets on an ongoing basis.

Example 2: Automatically Clone One Client

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

The tasks the system administrator must accomplish these 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.

The system administrator follows these steps below to automatically clone shakespeare's data:

1. Open the Groups window by choosing "Groups" from the Customize menu.
2. Create the group Documents. See the section "Configuring Backup Groups" in Chapter 6 for more information.

3. Select the appropriate choices in the Groups window. The critical choices are
 - Autostart choices: Enabled
 - Clones choices: Yes
 - Clone pool field: Backup Clone
4. Open the Clients window by selecting “Client Setup” from the Clients menu.
5. Select the appropriate choices for shakespeare. The critical choices are
 - Schedule field: Full Every Friday
 - Group choices: Documents

Following these steps, the system administrator creates a group for just one client so only that client’s data is cloned daily.

Example 3: Automatically Clone 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 using 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.

The system administrator must accomplish these tasks:

- 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 schedule named Acct Full+Clone. Table 9-1 summarizes these schedules.

Table 9-1 Acct Incr and Acct Full+Clone Schedules

Schedule	Sun	Mon	Tue	Wed	Thu	Fri	Sat
Acct Incr	Skip	incr	incr	incr	incr	incr	incr
Acct Full+Clone	Full	skip	skip	skip	skip	skip	skip

The system administrator avoids cloning all of the incremental backups by creating two schedules and selecting them for the appropriate set of clients.

- Create two groups: Create one group named Acct Incr and another group named Acct Full+Clone.
Cloning is applied to a group. Consequently, two groups are necessary: Acct Full+Clone group for performing full backups and clones and 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 set of clients uses the Acct Incr schedule, and the other identical set of clients uses the Acct Full+Clone schedule.
- Select the appropriate group and schedule for each set of clients in the Clients window.

Figure 9-13 diagrams the automatic cloning procedure for this example.

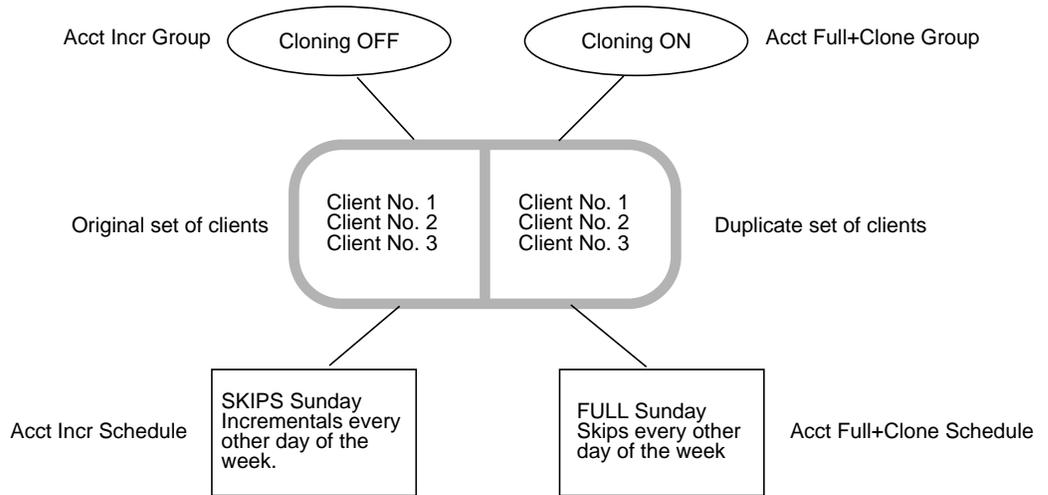


Figure 9-13 Automatic Cloning Scheme

To clone data automatically, follow these steps:

1. Open the Schedules window by clicking on "Schedules" in the Customize pulldown menu. For complete instructions on creating schedules, see "Setting Up the Backup Schedules" in Chapter 6.

The schedule Acct Incr performs incremental backups every day of the week, except Sunday when it skips the backup. Figure 9-14 shows this schedule.

Name:

Period: Week Month

August 1995							September 1995						
S	M	Tu	W	Th	F	S	S	M	Tu	W	Th	F	S
		1/i	2/i	3/i	4/i	5/i						1/i	2/i
6/s	7/i	8/i	9/i	10/i	11/i	12/i	3/s	4/i	5/i	6/i	7/i	8/i	9/i
13/s	14/i	15/i	16/i	17/i	18/i	19/i	10/s	11/i	12/i	13/i	14/i	15/i	16/i
20/s	21/i	22/i	23/i	24/i	25/i	26/i	17/s	18/i	19/i	20/i	21/i	22/i	23/i
27/s	28/i	29/i	30/i	31/i			24/s	25/i	26/i	27/i	28/i	29/i	30/i

Figure 9-14 Schedule for Acct_Incr

The second schedule, Acct Full+Clone, skips a backup every day of the week but performs a full backup and clones the data automatically on Sunday. Figure 9-15 shows this schedule.

Name:

Period: Week Month

March 1995							April 1995						
S	M	Tu	W	Th	F	S	S	M	Tu	W	Th	F	S
			1/s	2/s	3/s	4/s							1/s
5/f	6/s	7/s	8/s	9/s	10/s	11/s	2/f	3/s	4/s	5/s	6/s	7/s	8/s
12/f	13/s	14/s	15/s	16/s	17/s	18/s	9/f	10/s	11/s	12/s	13/s	14/s	15/s
19/f	20/s	21/s	22/s	23/s	24/s	25/s	16/f	17/s	18/s	19/s	20/s	21/s	22/s
26/f	27/s	28/s	29/s	30/s	31/s		23/f	24/s	25/s	26/s	27/s	28/s	29/s
							30/f						

Figure 9-15 Schedule for Acct Full+Clone

2. After the two schedules have been created, open the Groups window by selecting “Groups” from the Customize menu.
3. Create the first group, Acct Incr. For detailed instructions on creating a group see “Configuring Backup Groups” in Chapter 6. The critical choices are
 - Autostart choices: Enabled
 - Clones choices: No

The window with the Acct Incr group looks similar to Figure 9-16.

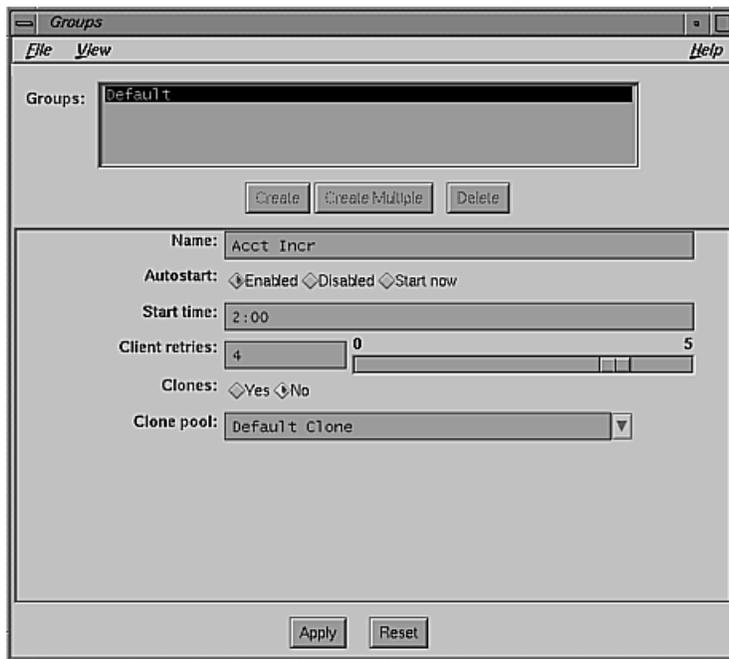


Figure 9-16 Acct Incr in the Groups Window

4. Create the second group, Acct Full+Clone. The critical choices are
 - Autostart choices: Enabled
 - Clones choices: Yes
 - Clone pool field: Backup Clone

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

The Acct Full+Clone group looks similar to Figure 9-17.

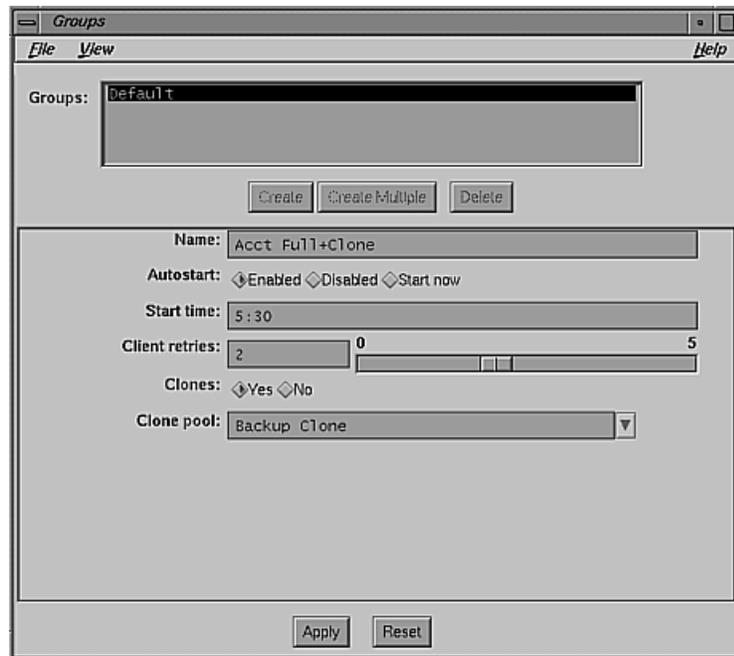


Figure 9-17 Acct Incr in the Groups Window

The system administrator now creates two sets of identical clients and selects their schedules and groups.

1. Open the Clients window by selecting Client Setup from the Clients menu. See “Adding a New Client” in Chapter 6 for complete instructions.
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. The critical choices are
 - 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. The critical choices are
 - Schedule scrolling list: Acct Full+Clone
 - Group choices: Acct Full+Clone
5. Click the *Apply* button, after making the choices for each client.

The instructions above comprise just one example of how to fulfill automatic cloning requirements.

Configuring and Scheduling Archive Requests

NetWorker includes an optional Archive product for long-term data storage. Whereas the goal of backup is to protect data against accidental loss or damage, the goal of archiving is to maintain a snapshot of data in perpetuity (or for as long as desired), and to conserve online storage space.

Note: If you are not sure if your site includes the NetWorker Archive product, read the release notes and the `/var/netls/nodelock` file.

This chapter describes how to configure, schedule, and control archive requests using the NetWorker Archive option. This chapter explains

- backup, archiving, and hierarchical storage management
- features of the NetWorker Archive product
- configuring NetWorker for archiving
- scheduling and monitoring archive requests
- example

This chapter ends with a shortcut procedure for scheduling and running archive requests.

Note: For information on installing and enabling the NetWorker Archive programs, see “Installing and Enabling NetWorker Software on Servers” and “Installing NetWorker Software on Clients” in Chapter 2. For information on using the NetWorker archive and retrieval programs, see Chapter 5, “Archiving and Retrieving,” in the *IRIX NetWorker User’s Guide*.

Backup, Archiving, and Hierarchical Storage Management

Data backup is the process of storing redundant copies of files and directories residing on a local disk onto removable media, usually tapes. Redundant copies can be recovered in case the original files are lost or damaged. The system administrator usually schedules backups daily. 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 typically initiate data archiving whenever they need, so a network-wide archiving policy is often impractical or unwise. Backups are always associated with a level—full, differential, or incremental—whereas archives are not.

When users archive project data, they sometimes want to delete those files to conserve space. As a consequence, they require that data be placed on long-lasting archive media that are durable, safe, and reliable. These media are usually removable, such as tape or optical disk.

Hierarchical Storage Management (HSM) is a policy-driven data management strategy in which data is automatically migrated from one storage medium to another, based on a set of policies. The policy often employed is *access rate*: the longer a file is inactive, the more likely it is to migrate. After a file migrates from primary to secondary storage, it can later be staged to tertiary storage based on another set of policies. 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 decrease the overall cost of storage. Migration and demigration must be automatic and reliable.

Features of the NetWorker Archive Product

The NetWorker Archive option consists of software for clients and a server:

- 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 can be enabled for backups, for archives, or for both.

The three categories of archive functions are

- data archiving
- data verification
- data retrieval

Data Archiving

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 for long-term storage on media called *archive volumes*. Archive volumes are similar to backup volumes, except that their retention policy has an unlimited expiration date.

Data archiving can be performed by end users or by the system administrator. Users archive data manually when they need to, whereas administrators can schedule an archive to take place anytime during the next 24 hours. For example, the best time to perform a large archiving operation 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 allows you to use the volume pools feature to separate backup from archive volumes, and archive volumes from archive clone volumes. See “Using Volume Pools” in Chapter 7 for information.

After archiving is complete, users and administrators are given the option of deleting archived files and directories, or leaving them in place. This option, called *grooming*, helps conserve disk space after a project is finished.

This chapter documents the Motif-based **GUI** interface for the archive and retrieval programs, although a command line-oriented interface is available as well: the commands *nsrchive* and *nsrretrieve*.

Data Verification

Because archived data might be deleted from the system immediately or eventually, NetWorker provides an extra measure of security to make sure that archived data is correct. NetWorker verifies data in two ways:

- media verification: NetWorker reads the archive volume and ensures 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 the archive volume is suspect, or if there are any problems with the data on tape, NetWorker issues a warning.

Data Retrieval

When you run the NetWorker retrieve program, NetWorker displays the names of the archive save sets on that server, listed by client name. You can retrieve a particular save set only if you have administrator or archive user privileges for that server.

You can search for specific archives or alter the sort order of archive save sets in the viewing list.

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 recovery feature.

Because archives theoretically last forever, NetWorker does not track them in the online index. Consequently, when a user retrieves archived data, the interface does not show the structure of files and directories. It shows only the archives that a server has made during a selected period, or for a certain user. Thus the best way for a user or administrator to remember archives is to keep a record of the date, or create an annotation that remains meaningful over the years.

Configuring NetWorker for Archiving

If your site has purchased the NetWorker Archive option, you can begin archive administration by invoking the *nwadmin* command:

```
# nwadmin& [ -s servername ]
```

Replace *servername* with the name of the NetWorker server where you wish to perform archive administration, or omit the *-s* option altogether.

If you have already started the NetWorker Administrator program, double-click on its icon, which is shown in Figure 10-1.



Figure 10-1 NetWorker Administrator Icon

The NetWorker Administrator window appears, as shown in Figure 10-2.

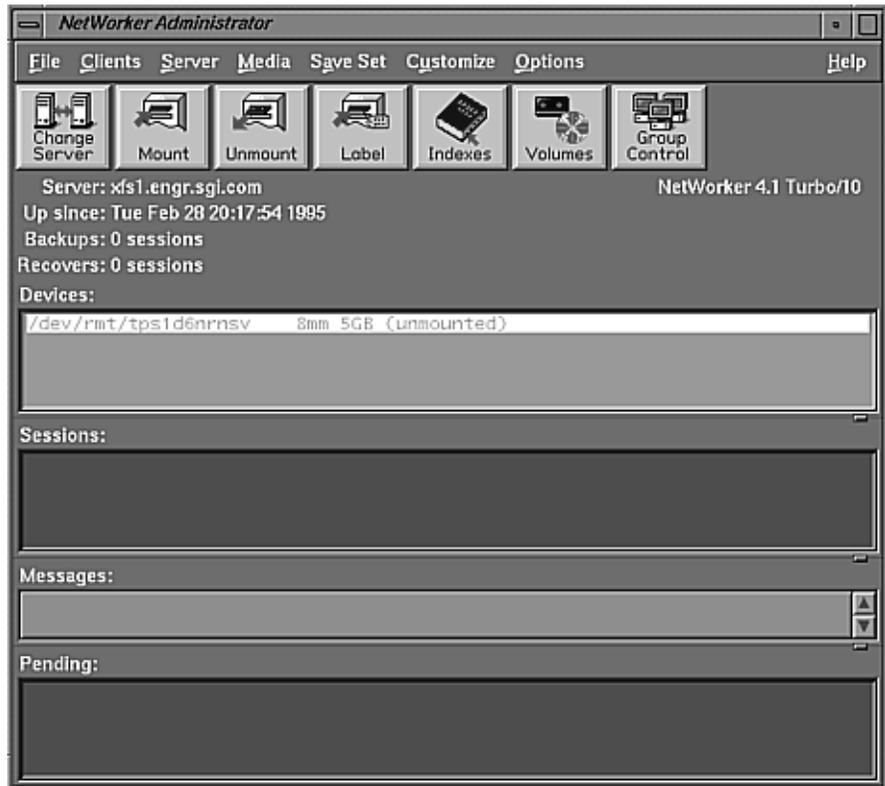


Figure 10-2 NetWorker Administrator Window

At the top level, this window looks like version 4.1.1 of NetWorker Administrator. The differences are as follows:

- in the Clients setup window, you can enable Archive services and add names to the Archive users list
- Archive Request Control menu item in the Server menu
- Archive Request menu item in the Customize menu

Use of these features is explained in the next sections.

Tip: If you have a jukebox, it is a good idea to insert an archive volume into one of the slots.

Before any archiving can occur, you must configure NetWorker to recognize one or more archive clients and set up save sets and volume pools as needed. This section explains

- configuring archive services for a client
- using archive save sets
- creating new volume pools for archiving

Configuring Archive Services for a Client

The arrow in Figure 10-3 points to the Archive services field in the Clients window. This field governs whether archives are enabled for the system selected in the Clients list. To allow archives for that client, click the *Enabled* button at Archive services. If archive services remain Disabled, that client will not be able to participate in a scheduled archive.

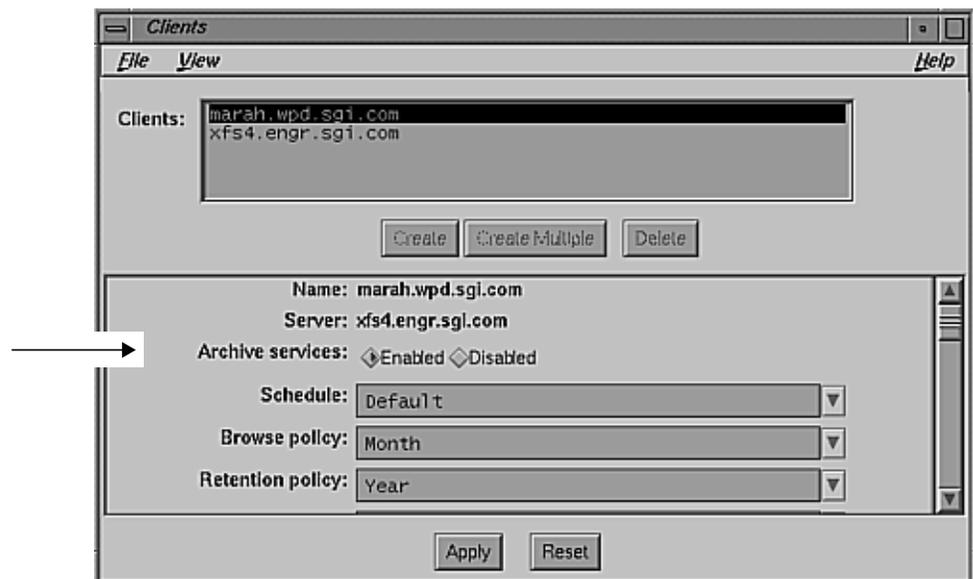


Figure 10-3 Enabling Archive Services for a Client

Note: When you enable archive services for one client, archive services are enabled for all clients of that server.

The Archive users list is at the bottom of the scrolling section of the Clients window. To allow users on the client system to perform manual archives and file grooming, enter their user names into the Archive users list. If no user names appear in the Archive users list, the client machine will be able to participate in scheduled archives, but individual users will not be able to archive files.

Using Archive Save Sets

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 complete: there are no backup levels or incremental saves.

Note: Archives are not recorded in the online file index, so they are not affected by the browse policy. The lack of associated indexes helps conserve disk space.

nsrretrieve (the retrieve program) is similar to the NetWorker *nwrecover* program, except that it works with archive save sets instead of backup save sets. Since archived files and directories are not recorded in the online index, the graphical 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 archives. These users, and users registered on the Archive users list in the Clients window, have permission to use the archiving and retrieval facilities.

Anyone can browse archive save sets—that is, look at information in the media database. However, you can retrieve only files that you own, unless you are superuser, in which case you can retrieve anyone's files.

Tip: 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; see “Labeling Backup Volumes” in Chapter 7.

Creating New Volume Pools for Archiving

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.

Note: Make sure the Store index entries field is set to No. This feature distinguishes archive pools from backup pools.

4. Click the *Apply* button to activate the new volume pool.

When you schedule a new archive request, you can use the new archive volume pool you have created. If you choose to clone archive data (as explained in “Cloning an Archive,” later in this chapter), you should also create a new archive clone pool.

Scheduling and Monitoring Archive Requests

This section explains

- scheduling an archive request
- starting or disabling an archive request
- viewing the progress of the current archive request
- verifying data on an archive volume
- cloning an archive

Scheduling an Archive Request

To schedule an archive on a one-time basis, choose “Archive Request” from the Customize menu.

The Archive Request window appears, as shown in Figure 10-4.

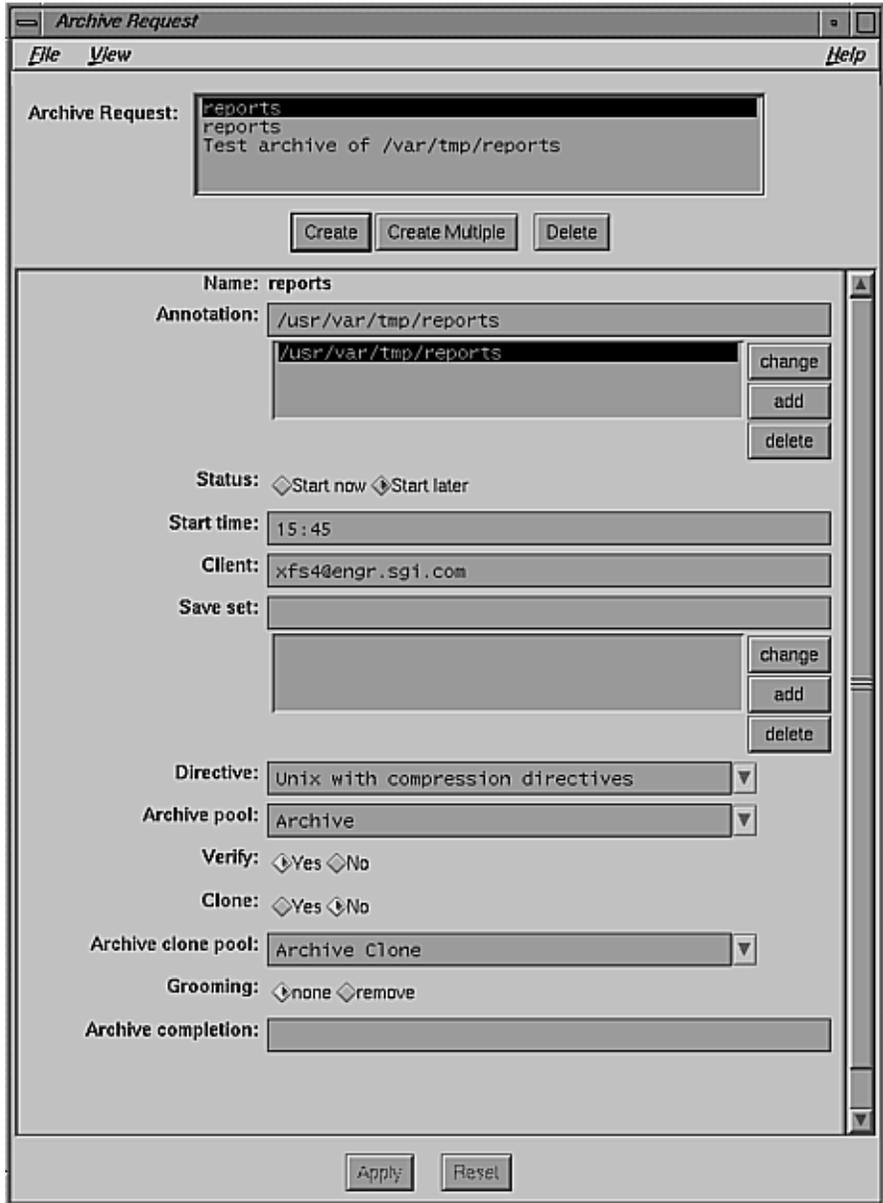


Figure 10-4 Archive Request Window, Expanded View

Use the features of the Archive Request window as follows:

Archive Request list

View a list of archives that someone has previously requested on this server. You can add new archive requests and delete old ones using the *Create* and *Delete* buttons.

scrolling panel

View the name of the selected archive. Use fields in this panel as follows:

- **Annotation:** Enter a comment of 1024 characters or less that you provide for this archive request.
- **Status:** “Start now” means archiving begins when you click the *Apply* button; “Start later” means it begins as indicated in the Start time field. After archive runs, neither is selected.
- **Start time:** Enter the time to begin, using a 24-hour clock.
- **Client field:** Enter the archive client’s hostname. For the server to archive itself, the client should be the same as the server.
- **Save set:** Enter pathnames of directories or files to archive.
- **Directive:** As in the Clients window, select the directive, usually Unix standard or Unix with compression.
- **Archive pool:** Enter the volume pool to which archives should be sorted. You should usually set this to Archive.
- **Verify buttons:** Indicate whether or not to automatically check the integrity of data archived on tape.
- **Clone buttons:** Indicate whether or not to automatically clone the archive volumes for extra security.

Note: If you enable cloning, you should also set the Archive clone pool, usually to Archive Clone.

- *Grooming* buttons: Indicate whether or not client files and directories should be removed after archiving is complete.
- Archive completion field: If desired, enter a command to execute after archiving is complete.

Starting or Disabling an Archive Request

To start an archive immediately, or disable a scheduled archive, choose “Archive Request Control” from the Server menu. The Archive Request Control window appears, as shown in Figure 10-5.

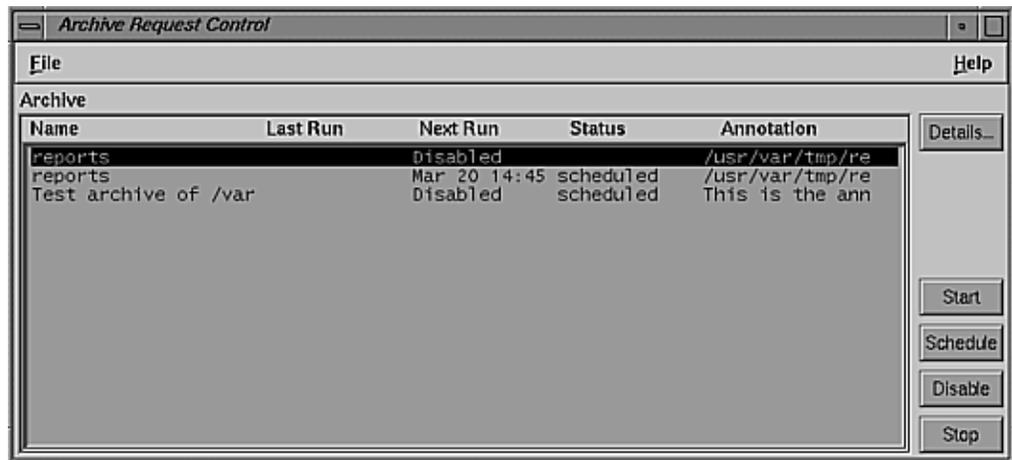


Figure 10-5 Archive Request Control Window

Use features of this window as follows:

Archive list View archive requests; the most recently run archive request is initially highlighted. 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.

- Start* button Click to initiate an archive request immediately. This action has the same effect as selecting a status of Start now for that archive request in the Archive List window.
- Schedule* button Click to reschedule an archive request; enter a new starting time. This action has the same effect as selecting Start later and specifying a new Start time in the Archive List window.
- Disable* button Click to disable an archive request that you scheduled here or in the Archive List window.
- Stop* button Click to halt an archive in progress.

Viewing the Progress of the Current Archive Request

To see the progress of a recent archive (the one currently selected), click the *Details* button in the Archive Request Control window. The Archive Request Details window appears, as shown in Figure 10-6.

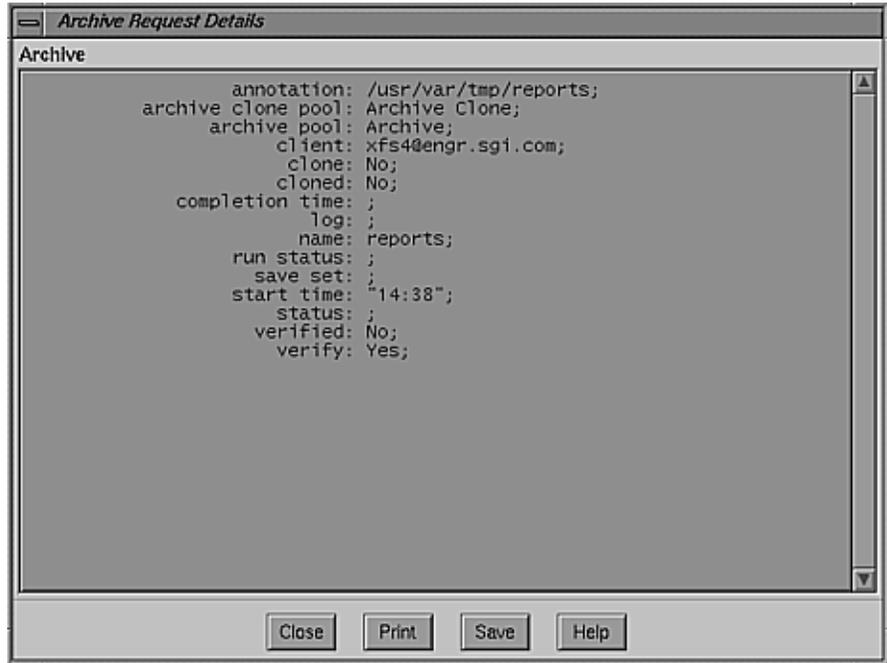


Figure 10-6 Archive Request Details Window

Note: Entries appear in the Archive Request Details window only after an archive request runs.

This window provides information about the completion of an archive request. Many fields in this window, such as annotation, archive clone pool, archive pool, client, clone, name, save set, start time, status, and verify, appear in the Archive Request window.

The features of the Archive Request Control window have the following functions:

completion time field

Indicates when the archive finished. Its duration is the difference between this and start time.

log field

Displays a message generated by the archive.

run status field Displays the success of the archive request, either completed, failed, or partial.

Note: For more information about an archive, see the log file `/usr/lib/nsrarchive`.

Verifying Data on an Archive Volume

To verify data already archived on an archive volume after the fact, you have two alternatives:

- Clone the archive data by choosing “Clone” from the Save set menu. During cloning, the original archive data is verified as the save set is copied from one volume to another. When you are done, you can reuse the cloned volume.
- Find out the save set ID of the archived data; for example, by searching the NetWorker Archive Retrieve window. Then enter the following command at the system prompt, substituting that save set ID for `ssid` (this how NetWorker verifies archived data):

```
retrieve -n -s ssid
```

Cloning an Archive

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 dialog box to activate the save set clone.

For more information on volume pools and save set cloning, see “Using Volume Pools” in Chapter 7 and “Cloning Save Sets” in Chapter 9, respectively.

Archive Example

Suppose you must decommission the workstation (and *hostname*) of someone who has left the company. It would be wise to archive the system first, in case its filesystems contain some essential files that will not be recognized as such until later.

This section gives an example of how to schedule and run an archive request:

- creating an archive client
- making an archive request
- checking the archive request
- rescheduling the 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 in Figure 10-3.
2. Click the *Create* button; the Clients window changes.
3. Enter the hostname of the decommissioned workstation into the client Name field.
4. Click *Enabled* after Archive services to allow archives for this client.
5. To permit employees on other machines to use archive and retrieval services, scroll to the bottom of the Clients window and add their user names to the Archive users field.

If you do not supply archive users, manual archives and retrieves are disabled on the client machine.

The system hostname is now a registered archive client. However, an archive does not take place until you request one.

Tip: 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 too long, especially during normal working hours when the network is busy. Therefore, it is efficient to use an archive request. Follow these steps:

6. Choose “Archive Request” from the Customize menu. The Archive Request window appears, as shown in Figure 10-4.
7. Click the *Create* button; the Archive Request window changes.
8. Type the Name you want to assign to this archive request, and a brief Annotation to remind you what the archive is for.
9. After Status, click on *Start later* to schedule this archive for tonight. Type the Start time you want, or accept the default starting time of 3:33 a.m.
10. Type the archive client machine name into the Client field, and the Save set pathname(s) you want to archive.
11. To speed the archive, select a Directive of either Unix standard or Unix with compression.
12. To separate this archive from normal backups, select the Archive pool named Archive, or be sure to use a different tape.
13. To check that archived data was accurately saved, click *Yes* in the Verify field.
14. To make a duplicate copy of this archive volume, click *Yes* in the Clone field. Accept the default Archive Clone pool.
15. To remove files and directories from disk after archiving them to tape, click *remove* in the Grooming field.

16. To be notified by e-mail when the archive completes, type the following into the Archive completion field:

```
/usr/sbin/Mail -s "archive request complete" root
```

This instruction sends it to root.

17. Click the *Apply* button to activate your changes.

You have now requested an archive of the system gallo to begin tonight at 3:33 a.m.

Checking the Archive Request

The next morning, check the outcome of the archive. If you set an Archive completion notice, you might have a 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 10-5 on page 268.
2. Click the *Details* button. The Archive Request Details window appears, as shown in Figure 10-6.

If the archive completed successfully, you are done, and you can begin dismantling the ex-employee's system. If the archive failed, you can reschedule it, as explained in the next section.

You can 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 the Archive volume had not been mounted. You decide to reschedule for the next night. Follow these steps:

1. Choose “Archive Request Control” from the Server menu. The Archive Request Control window appears, as shown on page 268. Make sure that the archive request you want is highlighted.
2. Click the *Schedule* button. The Archive Request Schedule dialog box appears, as shown in Figure 10-7.

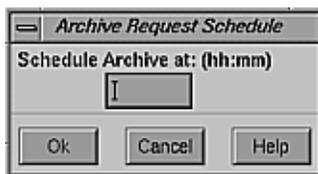


Figure 10-7 Archive Request Schedule Dialog Box

3. Type a new starting time into the Schedule Archive field, using the 24-hour clock, and click *Ok*.

The archive request is executed again that night, at the time you specified. If you change your mind and want to discontinue the archive, click the *Disable* button.

Shortcut

Use this shortcut to schedule and run an archive request if you are an experienced NetWorker user or after you have reviewed the information in this chapter. Follow this general procedure:

1. Create and enable an archive client using the Clients window.
2. Using the Archive Request 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 none, 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, use the Archive Request Schedule dialog box. To discontinue an 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.

Using NetWorker with Jukeboxes

If you are using a jukebox as one of your backup devices, this chapter provides the information you need to configure and operate NetWorker with a jukebox. Jukeboxes automate the time-consuming task of loading, mounting, and labeling backup media. Once you have installed both NetWorker and one of the optional NetWorker Support for Jukebox options, enabled the jukebox, and loaded the jukebox with backup volumes, you can run unattended backups and recovers.

Note: See “Backup Devices and Media” in Chapter 1 for a list of jukeboxes to use with Silicon Graphics systems.

By using a jukebox, users can recover files easily and efficiently without assistance from an administrator. NetWorker automatically mounts the correct volume so that the recovery proceeds as long as the volume is loaded in the jukebox.

This chapter explains

- testing the jukebox connection
- configuring a jukebox
- adding jukebox devices
- labeling volumes and using optional bar codes
- mounting jukebox devices manually
- performing routine jukebox operations

Testing the Jukebox Connection

To test the jukebox connection, run the *jbexercise(1M)* command with two pieces of “scratch” media loaded in the first and last slots of the jukebox. The tape drives should be empty with their door(s) open. See the *jbexercise(1M)* reference page for complete instructions.

When running *jbexercise*, you specify the control port and the device type. The control port for SCSI jukebox models is typically */dev/scsi/sc1d610*. For example, this command runs *jbexercise* on the EXB-10i jukebox:

```
# /usr/etc/jbexercise -c /dev/scsi/sc1d610 -m EXB-10i
```

If you are using a second jukebox, the control port is typically */dev/scsi/sc1d410*.

Note: The EXABYTE EXB-10i must be the only device on its SCSI bus.

Configuring a Jukebox

Most jukebox configuration occurs when you enable your NetWorker Support for Jukebox option (described in the section “Enabling Autochangers” in Chapter 2). The Jukeboxes window displays the configurations you selected when you installed it. After the jukebox is installed, enabled, and configured you can make changes only to Available slots, Devices, Bar code reader, and Match bar code labels in the Jukeboxes window.

This section explains

- using the Jukeboxes window
- adding devices to a jukebox
- changing available slots

Using the Jukeboxes Window

To open the Jukeboxes window, select “Jukeboxes” from the Media pull-down menu, as shown in Figure 11-1. This command is activated for those sites that have enabled one or more NetWorker Support for Jukebox options. (See “Enabling Autochangers” in Chapter 2 for information on enabling Jukebox options.)



Figure 11-1 Opening the Jukebox Window

Figure 11-2 shows the Jukeboxes window.

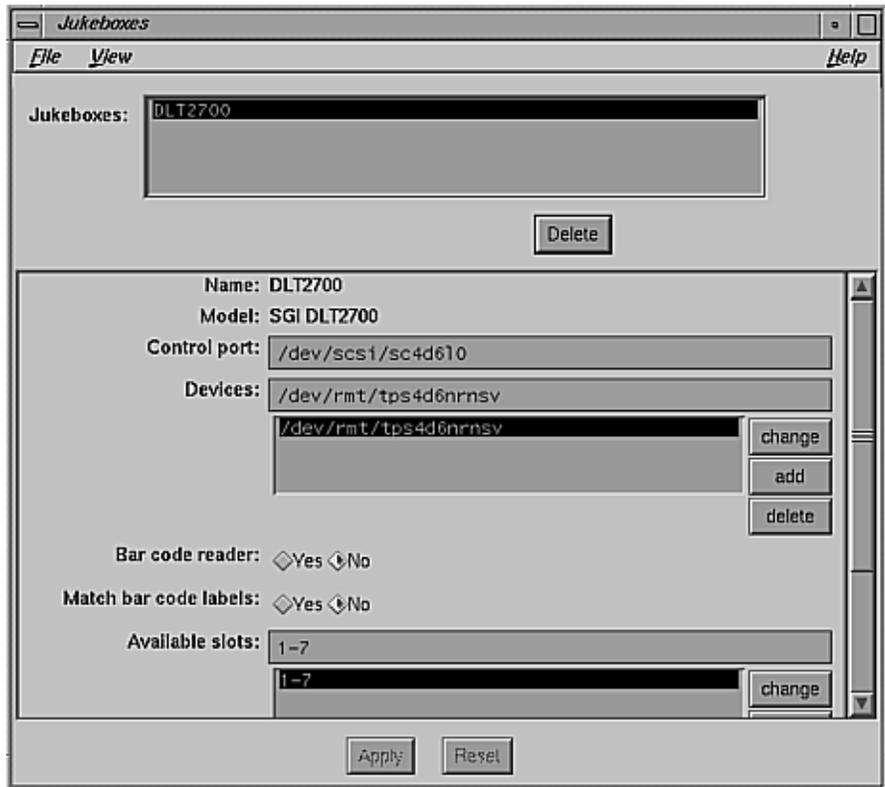


Figure 11-2 Jukeboxes Window

The Jukeboxes window lists configurations that are determined when the NetWorker Support for Jukebox option is installed and enabled. The features of this window have the following functions:

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

Display model types 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 pathnames of the device(s) 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. Before the jukebox is configured, the devices inside the jukebox are configured in the Devices window (like single-media devices). Devices inside a jukebox are designated by a “J” next to their pathname in the Devices display of the NetWorker Administrator window, as shown in Figure 11-3.

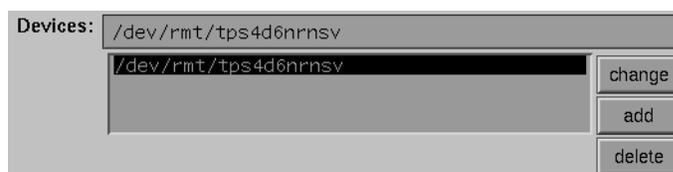


Figure 11-3 Devices Display of Administrator Window

Use the change, add, and delete buttons to change, add, or delete the highlighted device pathname in the Jukeboxes window.

- Bar code reader choices: provide choices for labeling your backup volumes with bar code labels
- Match bar code labels choices: provide choices for matching the bar code label to the volume label

Note: For more information on bar code labels, see “Labeling With Bar Codes” later in this chapter.

Available slots scrolling list

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 slots listed here correspond to the numbered slots

available for holding backup volumes inside the jukebox. Slots are identified by a range of numbers. Entering a single number, for example, 10, restricts access to just slot 10. Entering the number 10 is not the same as entering 1-10.

Adding Devices to a Jukebox

To add more devices to a jukebox, you must first use the Devices window to configure the NetWorker server to recognize the added devices. Then you must enter the pathnames for the jukebox devices in the Jukeboxes window so that NetWorker will recognize that these devices belong to the jukebox. After you have added 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.

To add a device to a jukebox, follow these steps:

1. To open the Devices window, select “Devices” from the Media menu. Figure 11-4 shows the Devices window.

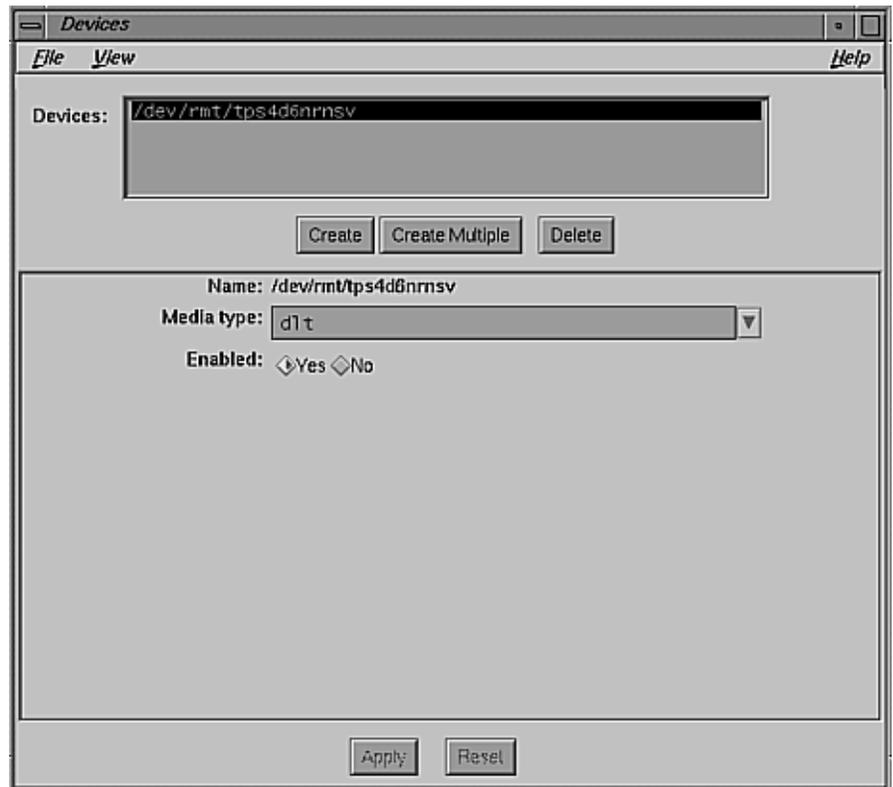


Figure 11-4 Devices Window

2. Click the *Create* button in the 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. Click the *Apply* button.
6. Select “Jukeboxes” from the Media menu to open the Jukeboxes window.
7. Enter the jukebox device pathname in the Devices field.

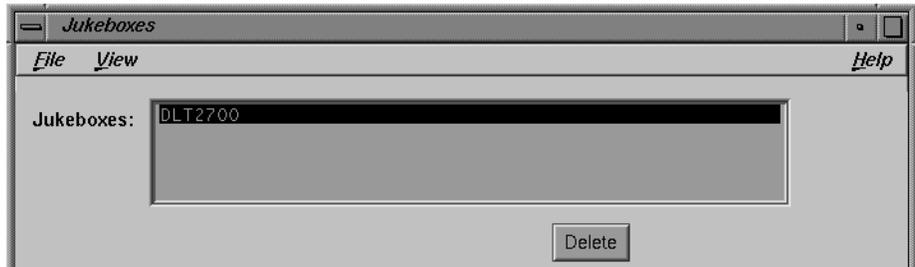


Figure 11-5 Entering the Jukebox Device Pathname

8. Click the *add* button to the right of the Devices scrolling list.
9. Click the *Apply* button.

Note: 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. Figure 11-6 diagrams an example.

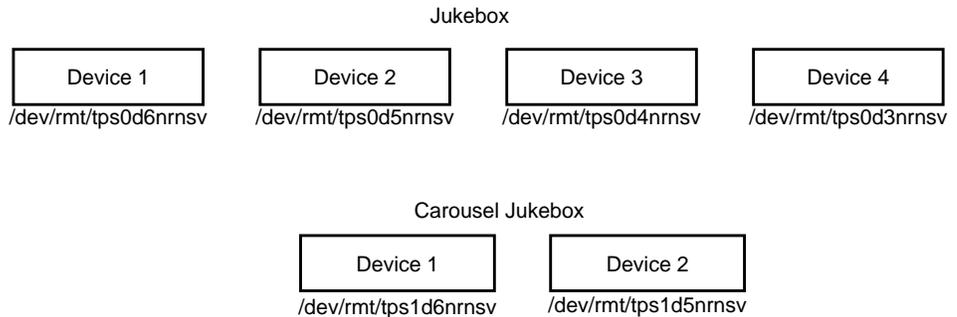


Figure 11-6 Listing Jukebox Device Pathnames

Changing Available Slots

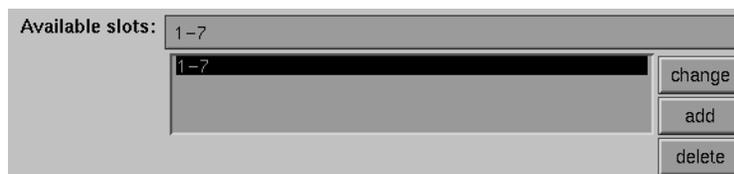
The Available slots feature allows you to control exactly which slots in the jukebox hold backup volumes available for backing up files. NetWorker uses all of the volumes in the jukebox for recovers. However, you can control which volumes NetWorker will automatically select for backups by designating a range of available slots in the jukebox.

For example, for a jukebox that contains 10 slots, you can designate slots 1 through 5 for NetWorker backups. The entries can be a range of slot numbers or a single slot number.

With two-sided media, the number of available slots will always be 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). However, the number of available slots is 32.

To change the number of available slots, follow these steps:

1. Enter the range of available slots into the “Available slots” field. Figure 11-7 shows this field.



The image shows a graphical user interface for configuring available slots. It features a text input field labeled "Available slots:" containing the text "1-7". Below this field is a scrolling list box, also containing "1-7". To the right of the scrolling list are three buttons: "change", "add", and "delete".

Figure 11-7 Changing Available Slots

2. Click the *change* button.

The slot numbers appear in the “Available slots” scrolling list.

3. Click the *Apply* button.

Note: Make sure you have placed backup volumes in all the available slots of the jukebox so NetWorker can proceed with an automatic backup.

Labeling Volumes and Using Optional Bar Codes

NetWorker uses the volume label and the optional bar code label for identifying backup volumes. This section explains how to label backup volumes in the jukebox, and how to create and apply optional bar codes to jukebox volumes:

- labeling volumes in a jukebox
- labeling with bar codes
- viewing labeling status messages

Labeling Volumes in a Jukebox

NetWorker uses the label template currently selected in the Pools window to label the backup volumes loaded in the jukebox. The Jukebox Labeling dialog box gives you control over the volume labeling. You can use a label template to generate names as needed, or you can specify a single unique name. Choose a label template from the Pools window or create your own before starting a labeling session.

Tip: Labeling volumes in the jukebox takes some time. You can perform the labeling process at a time when you do not need to back up or recover files.

To start labeling the backup volumes in the jukebox, follow these steps:

1. If you are using a label template, choose it from the Pools window. If you are creating your own label template, do so in the Pools window.
2. Select the jukebox device in the Devices display in the NetWorker Administrator window. A jukebox device is identified by a letter J in parentheses next to its pathname. (For example, `/dev/rmt/tps1d6nrnsv (J)` indicates a device inside a jukebox.)
3. Click the *Label* speedbar button. The Jukebox Labeling dialog box appears, as shown in Figure 11-8.

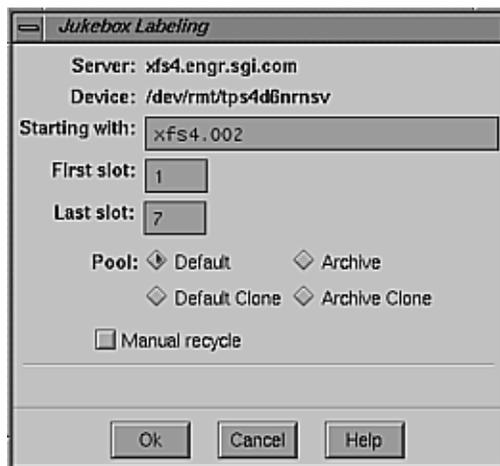


Figure 11-8 Jukebox Labeling Dialog Box

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 the First slot and Last slot fields to control the range of slots in the jukebox for labeling or updating backup volumes.

4. To override the range of slots specified by the label template, enter slot numbers into the “First slot” and “Last slot” fields in the Jukebox Labeling dialog box.

To label a single volume, enter identical numbers in the “First slot” and “Last slot” fields. To label a single backup volume, you can use any name; it does not have to match the label template. For labeling a single volume, the first and last slots should be identical, as shown in Figure 11-9.



Figure 11-9 Labeling a Single Volume

To label volumes in a series of slots, enter the range of slots in the jukebox containing the backup volumes that NetWorker will label in the First slot and Last slot fields. For example, if you previously labeled six tapes and need to label six more, your Jukebox Labeling dialog box might look like the one in Figure 11-10.



Figure 11-10 Labeling a Range of Volumes

Note: If you label a specific range of backup volumes, the “Starting with” name must match the label template format. If you label a single backup volume, you can use any name; it does not have to match the label template.

5. From the Pool choices, select a pool for the volumes you are labeling in the jukebox.
6. Click the *Manual recycle* toggle button only if you want to recycle the backup volume manually. Clicking this button overrides any previous browse and retention policies chosen for the volume and its save sets, and the volume can only be recycled manually.

Use this option if you plan on backing up important data to either cloned or non-cloned volumes, and you want to decide when the data and the volumes are ready for recycling. You can accomplish the same goal by post-labeling in the Volumes window.

Note: Only the pools that have been enabled in the Pools window appear in the Jukebox Labeling dialog box.

7. Click the *Ok* button to begin the labeling process.

A confirmation dialog box appears, as shown in Figure 11-11.

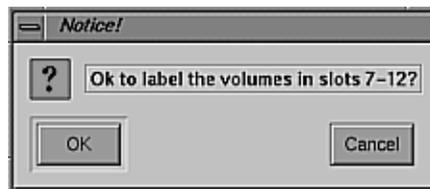


Figure 11-11 Relabel Confirmation Dialog Box

This dialog box is designed to keep you from relabeling media accidentally. 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.

Caution: If you remove backup volumes from the jukebox, remember that there might still be one or more volumes left in the jukebox devices. Unmount them by selecting the jukebox devices and selecting “Unmount”

from the Media menu in the NetWorker Administrator window. This step prevents you from accidentally relabeling these volumes when you reload the jukebox with new volumes.

Note: For more information on labeling volumes, see “Using Label Templates” in Chapter 7.

Labeling With Bar Codes

Labeling media with optional bar code labels provides two distinct advantages: it significantly speeds up volume inventory and provides improved labeling accuracy.

Using bar code labels for inventory 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 a laser beam while the volumes remain in their slots, greatly reducing the time it takes NetWorker to locate a volume or determine the contents of a carousel cartridge.

Bar code labels also provide greater labeling accuracy because the labels are affixed to the media before they are loaded and scanned in the jukebox. Once the jukebox scans the bar code label, NetWorker records and tracks it in the media index.

The volume label and the optional 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 (internal volume label). NetWorker uses bar code labels for inventorying volumes and uses the volume labels to identify the volumes needed for backup and recovery. You can make the volume and bar code label match by selecting the appropriate choices for bar code labeling in the Jukeboxes window.

Caution: 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.

This section explains

- designing a labeling scheme
- labeling volumes with bar code labels
- using bar codes that do not match volume labels

Designing a Labeling Scheme

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 you load and mount each volume individually, so NetWorker can match up the bar code label with the existing volume label.

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.

Note: When NetWorker relabels volumes automatically, it reuses the original volume label name. You can change the label name only if you relabel the volumes manually.

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, 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 to save hours of time when backup volumes are inventoried.

You can 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 order bar code labels that match your current volume labeling scheme. Labeling instructions for bar code labels should be provided with the jukebox hardware documentation. If you have questions about the bar code labels, call the hardware manufacturer.

Note: NetWorker displays both the bar code label and volume label in the Pending display of the NetWorker Administrator window when a volume is requested. You can also view both the volume label and its associated bar code label in the Volume Management window.

Labeling Volumes With Bar Code Labels

Use the Jukeboxes window and the Jukebox Labeling dialog box to label your volumes with bar codes. 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. Open the Jukeboxes window by selecting “Jukeboxes” from the Media menu. Figure 11-12 shows the Jukeboxes window.

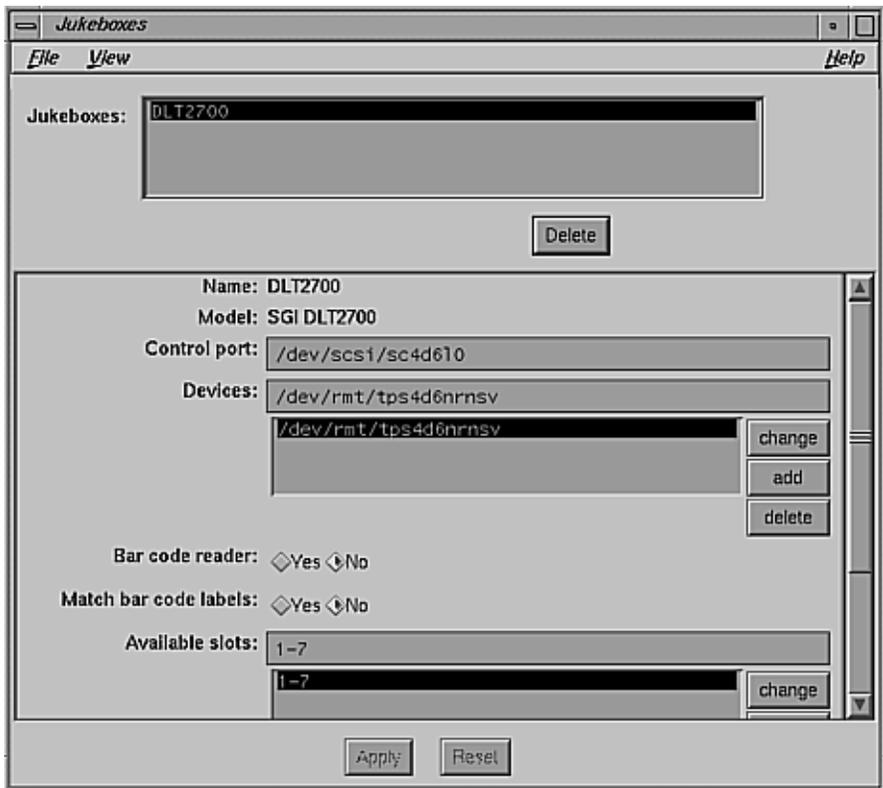


Figure 11-12 Jukeboxes Window

4. Select “Yes” from the Bar code reader choices.

If you want the volume label (normally created using a label template) to match the bar code label, select “Yes” from the Match bar code labels choices. If you do not want to override the label created by the label template with the bar code label, select “No.”

In either case, you will have two labels: the volume label (which may or may not match the bar code label depending upon your bar code choices in the Jukeboxes window) and the bar code label.

5. Open the Jukebox Labeling dialog box by clicking the *Label* speedbar button or by selecting “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 11-13.

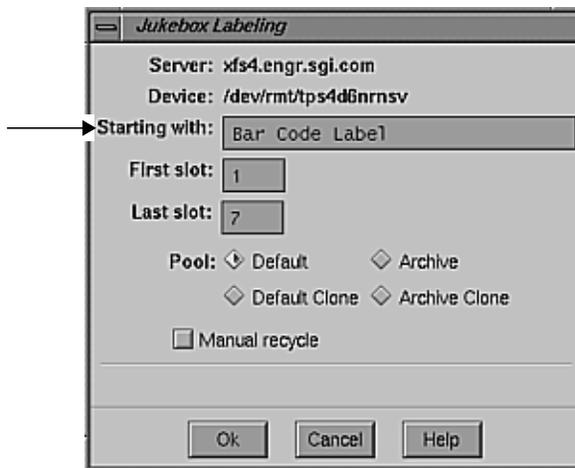


Figure 11-13 Bar Code Label in Jukebox Labeling Dialog Box

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 in Figure 11-14.

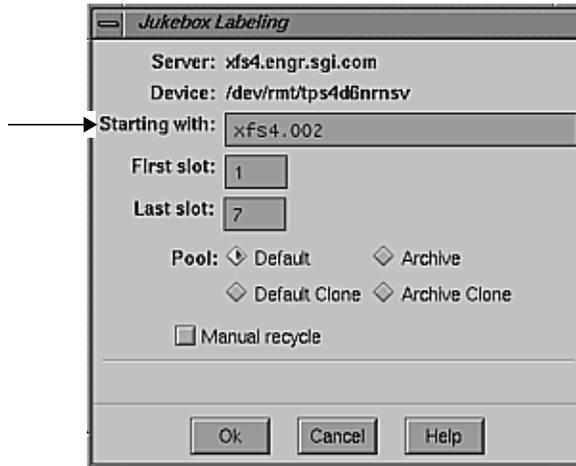


Figure 11-14 Template Name in Jukebox Labeling Dialog Box

Once you have completed selecting your bar code label choices and labeling the backup volume, both label names for the volume appear in the Volume Management window.

The Volume Management window looks similar to the one shown in Figure 11-15.

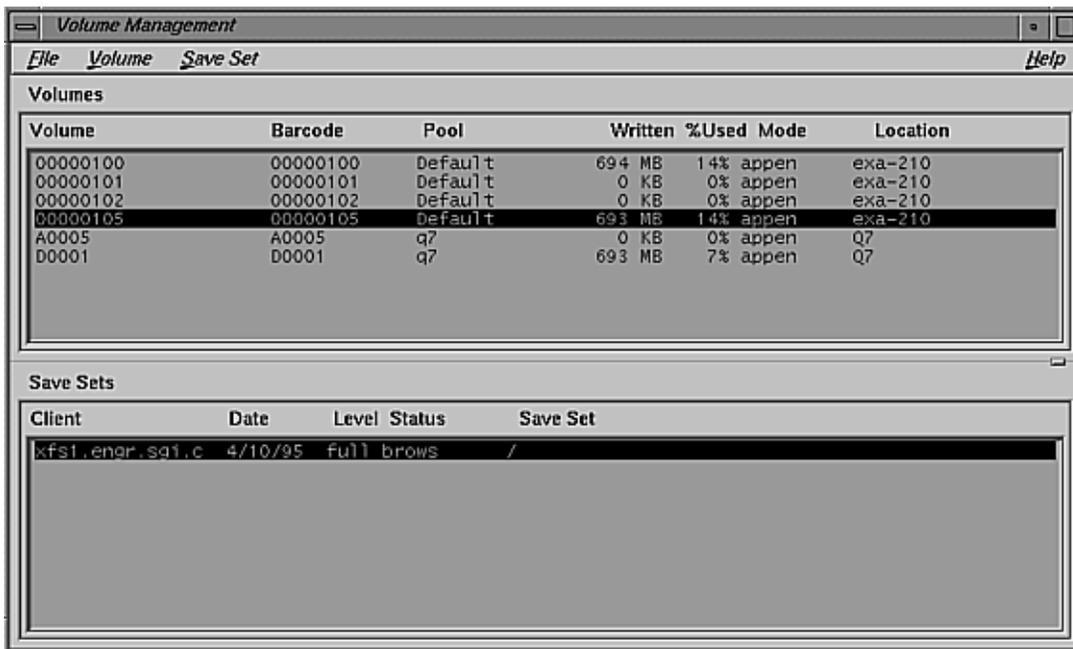


Figure 11-15 Volume Management Window with Complete Bar Code Selections

Caution: 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, apply a different label and begin the labeling process again, or disable Match bar code labels in the Jukeboxes window when you label the second volume.

Using Bar Codes That Do Not Match Volume Labels

If you choose not to match the volume label to the bar code label, create and attach volume labels to the outside of your media:

1. Affix all of the bar code labels.
2. Load the volumes in the jukebox.
3. To begin the labeling process, select “Yes” for Bar code reader choices, and select “No” for Match bar code labels choices in the Jukeboxes window. NetWorker uses the next available label from the label template for the volume name.

NetWorker proceeds labeling the volumes and records both labels in the media index.

4. After NetWorker completes the labeling process, open the Volume Management window to determine the volume label and bar code label for each volume.
5. Create your own volume labels to attach to the backup volumes. Use the information in the Volume Management window to match the correct volume labels to the bar code labels.

Viewing Jukebox Labeling Messages

Whenever you start a labeling operation for a jukebox, NetWorker displays a status message in the Jukebox Labeling dialog box to show the progress of the labeling process.

A message appears in the dialog box, as shown in Figure 11-16.



Figure 11-16 Jukebox Labeling Dialog Box: Status Message

You can cancel the label operation by clicking the *Cancel* button. When labeling is complete, another message informs you that the labeling operation is finished. Click the *Ok* button to close the dialog box.

Mounting Jukebox Backup Volumes Manually

To mount a backup volume in a jukebox, you must first make sure the jukebox device is selected in the NetWorker main window. This is especially important if the NetWorker server is configured to use more than one device.

Tip: Although jukebox devices are identified by a “J” next to the pathname in the main window, it is a good idea to place an adhesive label on the outside of a jukebox identifying its device pathname(s), to remind you of which device pathname(s) belong to the jukebox. This practice is especially helpful when you are using more than one jukebox.

To mount a jukebox volume manually, follow these steps:

1. Select a jukebox from the Devices display.
2. Click the *Mount* speedbar button. The Jukebox Mounting dialog box appears, as shown in Figure 11-17.

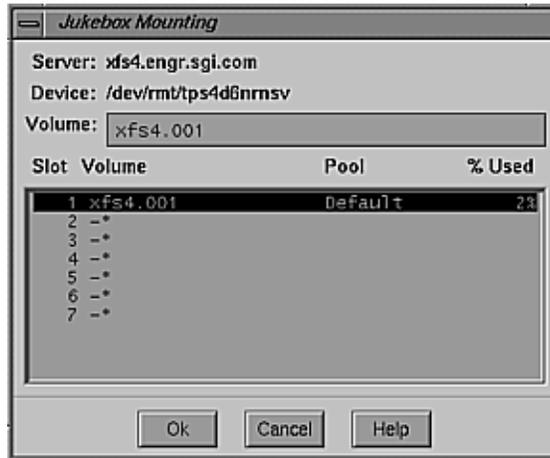


Figure 11-17 Jukebox Mounting Window

The Jukebox Mounting window displays the name of the NetWorker server and the system pathname of the selected jukebox device. The Volume scrolling list displays the slot numbers (Slot), the label of volume stored in the slot (Volume), and the percent capacity filled of the volume (% Used).

The Jukebox Mounting window also displays the pool (Pool) to which the volumes belongs. This information can be used to determine which volumes to mount in the jukebox. You might not want to mount volumes that are full or near capacity. The name of the currently selected backup volume appears in the Volume field.

3. To mount the selected backup volume in the jukebox, select it from the list, and click the *Ok* button.

Performing Routine Jukebox Operations

This section contains information and suggestions for operating the jukebox on a day-to-day basis. It explains

- viewing jukebox notices
- checking jukebox space
- recycling versus adding more backup volumes
- taking inventory of the jukebox contents
- using pools with a jukebox
- cleaning the devices
- calibrating the devices

Viewing Jukebox Notices

NetWorker uses electronic mail to send notices about NetWorker events. The jukebox software uses the three Tape mount request notices to inform the system administrator that the jukebox needs attention. The notification Tape mount request 1 comes undefined, so that you may determine your own message for this particular notification.

The following situations require attention:

- the backup volumes in the jukebox are 90% full
- the jukebox is out of space for the backups and needs more backup volumes to continue
- the jukebox has a mechanical problem

To open the Notifications window, select “Notifications” from the Customize pull-down menu. The Notifications window with the three tape mount request notifications is shown in Figure 11-18.

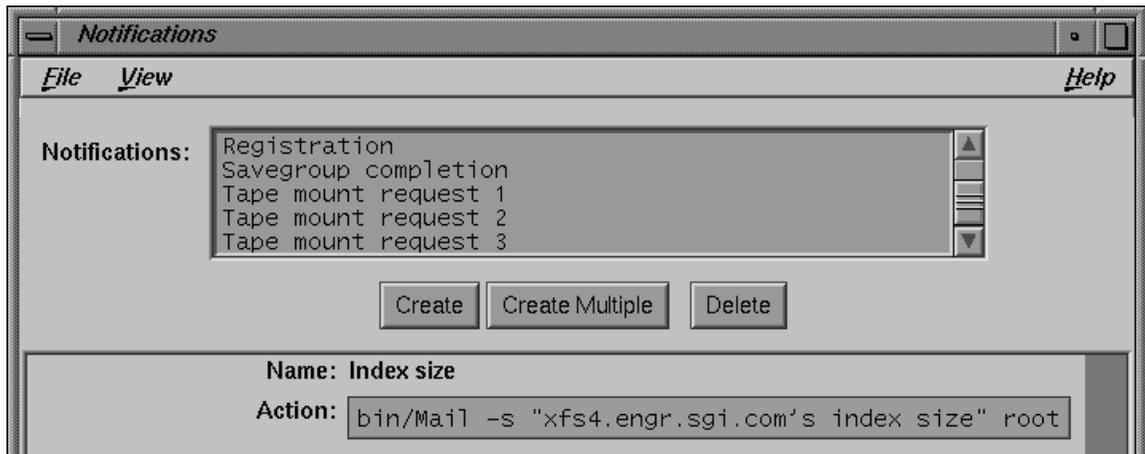


Figure 11-18 Notifications Window With Three Tape Mount Request Notifications

See “Checking NetWorker Notices” in Chapter 8 for more information on the Notifications window.

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, a notice appears in the Pending display of the NetWorker Administrator window.

Note: After you correct a jukebox problem, you must mount a backup volume (using the “Mount” command from the Operation pull-down menu in the main window) before continuing to back up or recover files.

Checking Jukebox Space

To find out how much space is available in the jukebox, you can use either the Jukebox Mounting window (shown in Figure 11-17) or the `nsrjb(1M)` command. The Jukebox Mounting window displays all the media in the jukebox and the percentage used of each tape.

If you prefer to use the *nsrjb* command, follow these steps:

1. Open a shell.
2. Enter the *nsrjb -v* command:

```
mars# /usr/etc/nsrjb -v -jjukeboxname
```

NetWorker displays information about the backup volumes in the jukebox that looks similar to this:

```
Jukebox 8LT:
  slot  volume                used  pool
    1:  bitbucket.008          full  Default
    2:  bitbucket.009          full  Default
    3:  bitbucket.010          full  Default
    4:  bitbucket.011          full  Default
    5:  bitbucket.012          full  Default
    6:  bitbucket.013           20%  Default
    7:  -*                      -

    * not registered in the NetWorker media database
    6 volumes, 1 less than 80% full.
    56 GB total capacity, 10 GB remaining (85% full)
```

```
drive 0 (/dev/rmt/tps131d6nrnsv) slot 1: bitbucket.008
```

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.

Recycling Versus Adding More Backup Volumes

NetWorker can save files on volumes in either appendable or recyclable mode as indicated in the Volume Management window. If the volumes inside of the jukebox are in full mode, they cannot receive backups.

If your volumes are all in full mode, you can do one of the following:

- If you need to keep the volumes for long-term storage, remove the full volumes and replace them with new media. Label and mount the new volumes.
- If you do not need the data on the full volumes, you can manually change the mode to recyclable in the Volume Management window.

NetWorker overwrites the data on the volumes with new backups and maintains the existing labels.

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 the retention policy of the client that generated them.

There are advantages to either recycling or adding more media to the 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 pool instead of recycling it. For example, if the jukebox needs new backup volumes every three months, and the company policy is to maintain the backups for a year, you have to keep adding new media to the pool until you can recycle the backup volumes that contain expired or old backups.

Taking Inventory of the Jukebox Contents

Whether or not you use bar code labels in conjunction with your volume labels, when NetWorker assigns a volume label, it registers the location of the volumes in the jukebox slots. If 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 tell NetWorker that the jukebox now holds a different set of labeled volumes or that the volumes are in a different order. This procedure is called *taking an inventory*.

When you take an inventory of the volumes in the jukebox, NetWorker reads the label of each volume and remembers its slot number. For example, if you have more than one job pack (removable rack of tapes) for a jukebox, you must take an inventory each time you remove one job pack and load another one into the jukebox, unless you perform the label process for the new job pack.

To tell NetWorker you have changed the contents of the jukebox, follow these steps:

1. Load the new backup volumes into the jukebox.
2. Select the jukebox device in the Devices panel of the main window. Remember, jukebox devices are marked with a “J.”
3. Open the Jukebox Inventory dialog box by selecting “Inventory” from the Media menu. Figure 11-19 shows the dialog box.



Figure 11-19 Jukebox Inventory Dialog Box

4. Make sure the slot numbers in the First and Last Slot fields are the slots you want to inventory. If they are not, enter the correct range of slot numbers.

The Jukebox Inventory dialog box displays a message that monitors the progress of the inventory operation.

You can cancel the inventory operation after NetWorker completes inventorying the current volume by clicking *Cancel* in the Jukebox Inventory dialog box.

5. Click *Cancel* to close the window when the inventory is complete.

When an inventory is completed, NetWorker has registered the contents of the jukebox and can therefore proceed with its network-wide backup and recover services.

Note: For information on determining jukebox capacity, see “Monitoring Backup Devices” in Chapter 8.

Using Pools With a Jukebox

If you plan to have full and nonfull backups, it is a good idea to estimate the number of volumes needed for your full backups and assign them to the Full pool in the Label window. This technique ensures that your full backups are located in a consecutive range of slots in the jukebox so the volumes can be easily removed.

Cleaning the Devices

Check the jukebox manufacturer's documentation for information about the frequency and cleaning method for the jukebox's devices. It is important to keep the jukebox in good working order with regular maintenance. Refer to the jukebox hardware manuals for information about preventive maintenance.

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.

NetWorker Performance

The performance of a server system is affected by the speed of the backup device, the network speed, the amount of main memory, the disk speed, the CPU speed, and the number of CPUs. The factors affecting the performance of a client system are filesystem traversing, generation of data, data on multiple disks, CPU speed, and number of CPUs.

This chapter explains

- guidelines for choosing a configuration
- measuring server performance
- measuring client performance

Guidelines for Choosing a Configuration

Several factors determine which NetWorker server configuration best suits your backup and recover needs. The configuration consists of the hardware and software, which includes tape drives, jukeboxes, client systems, and network connection.

This section provides a few simple rules that you can use to guide your choices, and focuses on backup, since backup requires far more server capacity than recover and occurs so much more often.

Note: Please keep in mind that these are guidelines, and actual performance 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, network connection, jukeboxes, clients, 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. Table 12-1 shows the data transfer rate for several tape drives.

Table 12-1 Tape Drives and Data Transfer Rates

Drive	Data Transfer Rate
DAT	200 KB/s
EXABYTE 8200	250 KB/s
EXABYTE 8500	500 KB/s
Digital Linear Tape	1.25 MB/s

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. The numbers listed below are considered average transfer rates for each client system. However, it is a good idea to run several clients simultaneously to help smooth out fluctuations in each client's data transfer rate. Table 12-2 shows the data transfer rates for several types of clients.

Table 12-2 Client Data Transfer Rates

Client: 1 Backup	Data Transfer Rate
Silicon Graphics	600 KB/s
PC/DOS/SPX	80 KB/s
PC/DOS/TCP	150 KB/s

Table 12-2 Client Data Transfer Rates

Client: 1 Backup	Data Transfer Rate
Sun™ SS2	200 KB/s
IBM® RS/6000	300 KB/s

Network

Ethernet has an upper limit on bandwidth of about 1 MB per second, but in practice, most Ethernet networks are limited to about 800 KB per second between a set of clients and a single server. For higher performance, use FDDI. Table 12-3 shows the data transfer rates for several networks.

Table 12-3 Data Transfer Rates for Networks

Network	Data Transfer Rate
Ethernet	800 KB/s
FDDI	8 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 number of CPUs increases.

Jukeboxes

Jukeboxes provide automatic loading and unloading of tapes or optical disks. This capability assists the administrator in two different ways. During nightly backups, NetWorker uses the jukebox to automatically switch to the next tape when a tape fills up. During restores, NetWorker uses the jukebox to load all of the tapes needed for the recovery without operator

intervention. Table 12-4 shows data transfer rates and capacities for several jukeboxes.

Table 12-4 Jukebox Data Transfer Rates and Capacities

Jukebox	Maximum Data Transfer Rate	Capacity	Average Data Transfer Rate
EXB-10i	480 KB/sec	50 GB	480 KB/sec
DLT2700	1.25 MB/sec	70 GB	1.25 MB/sec

To determine the capacity requirements of a jukebox for a scheduled unattended backup, select 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 a DLT2700 jukebox or two EXB-10i autochangers.

To determine how much disk space you need for the online indexes for quick recovers, first do a rough calculation of the amount of data backed up in a single schedule period (for example, week, month, or quarter). Use the guidelines in Table 12-5 showing how much data is backed up with different levels of backup.

Table 12-5 Percent of Data Backed Up for Each Type of Backup

Level	Percent of Data Backed Up
Full	100%
Level 1-9	25%
Incremental	10%

For example, a monthly schedule that has one full backup on the first Sunday, a level 5 backup on other Sundays, and incrementals every other day looks like Table 12-6.

Table 12-6 Example of the Percent of Data Backed Up

Level	% of data backed up
1 Full	100%
26 Incremental	260%

Table 12-6 Example of the Percent of Data Backed Up

Level	% of data backed up
4 Level 5	100%
Total	460%

This table illustrates that 460% of the total amount of data is backed up over the course of a month. 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.

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 for them files that they had any time during the past six months. So you need six months times 46 GB, which yields 276 GB of capacity.

In practice, you need a little extra jukebox capacity, since there are a small number of “unavailable” volumes as 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 that you should purchase a jukebox, or set of jukeboxes, with about 50% more capacity than your current requirement.

NetWorker Configuration: Example 1

Site A has approximately 70 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). This equation calculates the required data transfer rate to achieve this goal.

$$70000 \text{ MB} / 12 \text{ hours} = 6700 \text{ MB} / \text{hour} = 1660 \text{ KB} / \text{second}$$

To back up the data with a single NetWorker server, this configuration is suggested:

- IRIX NetWorker with the TurboPak option
- two DLT2700 jukeboxes containing DLT tape drives connected to the NetWorker server, with licenses for each
- two Ethernet network interfaces
- at least 1600 MB of free disk space on the server for the client index files
- 50 Client upgrade license

NetWorker Configuration Example 2

Site B has 50 GB of data on a single network with 80 clients; the administrators want to be able to schedule backups in a single night (eight hours). The full backups for the clients must be staggered due to the limit of 800 KB/sec data transfer rate per Ethernet network. Calculate the backup capacity required to complete the backups in one night:

$$800 \text{ KB / second} * 8 \text{ hours} = 23 \text{ GB / night}$$

By using three different backup schedules to stagger their full backups into three nights instead of one, the administrators reduce the load of the nightly backup data from 50 GB per night to about 20 GB per night.

$$\text{Full: } 50 \text{ GB} / 3 = 16.7 \text{ GB}$$

$$\text{Incr: } (50 \text{ GB} - 16.7 \text{ GB}) * .1 = 3.3 \text{ GB}$$

$$\text{Total: } (16.7 \text{ GB} + 3.3 \text{ GB}) / 8 \text{ hours} = 694 \text{ KB / second}$$

To back up the data with a single NetWorker server, this configuration is suggested:

- IRIX NetWorker with the TurboPak option
- one EXB-440 jukebox with two EXABYTE 8505 tape drives connected to the server, with license
- one network interface
- approximately 1 GB of free disk space on the server for client index files
- 50 client upgrade licenses

Measuring Server Performance

This section provides examples of how to measure the performance of a server.

Backup Device Speed

Most tapes have a step function in their data rate. 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 and list its size. For example:

```
# cat /unix > big
# ls -l big
-rw-rw-r-- 1 root sys 20675420 Mar 5 16:11 big
```

2. Use the `dd(1M)` command to write the large file to tape four times and measure the time results:

```
# time dd if=big of=/dev/rmt/tps1d6nrnsv bs=32k conv=sync
95.2 real 13.0 user 11.9 sys
# time dd if=big of=/dev/rmt/tps1d6nrnsv bs=32k conv=sync
78.2 real 12.9 user 12.7 sys
# time dd if=big of=/dev/rmt/tps1d6nrnsv bs=32k conv=sync
78.0 real 12.8 user 12.5 sys
# time dd if=big of=/dev/rmt/tps1d6nrnsv bs=32k conv=sync
76.8 real 13.0 user 12.4 sys
```

3. Divide the file's size by the average of the last three real times. For example:

```
Rate: 20675 KB / 77.66 seconds = 266 KB / second
```

This number gives you the rate of the tape speed.

Network Speed

NetWorker 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.

```
# cat /unix > big
# ls -l big
-rw-rw-r-- 1 root sys 20675420 Mar 5 16:11 big
```

2. Use the `rcp(1C)` command to copy the file from the client to the server and time the result:

```
# time rcp big server:/dev/null
38.2 real 0.2 user 30.7 sys
```

3. To find the network speed, divide the number of bytes in the file by the real time. For example:

```
Rate: 20675 KB / 38.2 seconds = 541 KB / second
```

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.

Server CPU Speed

The speed and the number of the CPU(s) of a server limits the following:

- total data throughput to tape
- interrupts per second for network data
- context switches per second between processes

The best measure is the number of CPUs for the server. More CPUs means a faster system.

Memory

The memory on the server limits the amount of data buffered between the NetWorker `save(1M)` command, agent daemon, and media management daemon. In general, the more memory the server has, the better performance is.

Measuring Client Performance

This section provides examples of how to measure the performance of a NetWorker client.

Filesystem Traversing

To measure the filesystem traversing speed, follow the steps below:

1. Time the `uasm(1M)` command with the `-bi` option. For example:

```
# time /usr/etc/uasm -bi /usr
33848 records 6961176 header bytes 644814472 data bytes

real 51.86
user 7.73
sys 27.95
7.7u/27.9s (68% of 0:51) 0k+0k+0k 0pf+0sw 2993i+16o
```

2. Divide the number of records by real time for rate per file. For example:

```
33848 records / 51.86 seconds = 652.7 files / second
```

Data Generation Rate

To measure the rate at which a client generates data for a backup, follow the steps below:

1. Time the `uasm` command with the `-si` option and redirect the output to `/dev/null`. For example:

```
# time /usr/etc/uasm -si /usr > /dev/null

real 6:39.49
```

```
user 12.62
ys 1:27.19
12.6u/87.2s (24% of 6:39) 0k+0k+0k 0pf+0sw 45302i+27o
```

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:

```
629701 KB / 399.49 = 1576 KB / second
```

Data on Multiple Disks

NetWorker automatically backs up multiple disks in parallel.

To measure parallel disk speeds, follow these steps:

1. Use the *df(1)* or *du(1M)* command to find two directories of approximately the same size 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.

Client CPU Speed

The CPU of a client limits the following:

- the total data throughput to tape
- the interrupts per second for network data
- the context switches per second between processes

The best measure is the **MIPS**[®] rating for the client. A larger **MIPS** rating means a faster system.

Troubleshooting

This appendix contains troubleshooting information that answers common questions about operating and configuring NetWorker:

- maximum filename length
- checking the NetWorker daemon
- */nsr* not configured
- solving licensing problems
- renaming a client
- using recover access
- using savegroup completion messages
- using savegroup completion warning messages
- solving errors in executing external ASMs
- solving X errors
- moving indexes
- recovering files from an interrupted backup
- determining the NetWorker server
- using *nsrexecd*

Maximum Filename Length

NetWorker supports a maximum filename size of 1024 characters. This limitation is the same as the IRIX *svid* limitation.

Checking the NetWorker Daemons

If you have trouble starting NetWorker or see error messages in */var/adm/SYSLOG* from NetWorker daemons, the daemons may not be running properly. To check the daemons, enter this command:

```
# ps -ef | grep nsrd
```

If you discover that you need to start the NetWorker daemons, enter the commands below:

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

/nsr Not Configured

If you have not created the directory */nsr* during installation and enabling of the NetWorker server software (see the section “Enabling Server Software” in Chapter 2), this message appears on the system console and in */var/adm/SYSLOG* when the NetWorker daemons are started:

```
nsrd: error: /nsr must be a directory or a symlink to a directory  
nsrd: info: Please create a symbolic link called /nsr that points  
nsrd: info: to a directory in a filesystem with sufficient space  
nsrd: info: to store NetWorker's index files
```

Follow step 6 or 7 in the section “Installing Server Software” in Chapter 2 to solve this problem.

Solving Licensing Problems

The following messages, and others, indicate that NetLS licenses for NetWorker products have not been properly enabled:

```
NetWorker Registration: (notice) server disabled; No valid
Base or Advanced licenses found,

nsrd: either install new licenses or recover a valid
resource file. (15017)

NetWorker registration: (info) Server is disabled (Install
base enabler).

nsrccap: RPC error, Program not registered (24015)
nsrccap: nsrd did not respond. nsrd must be started before
running nsrccap. If nsrd is running, you may have exceeded a
resource limit.
```

To correct licensing problems, follow the procedures in the section “Enabling Server Software” in Chapter 2 to enable your NetLS licenses. To verify that licenses have been enabled, display them in the Registration window, as described in the section “Displaying NetWorker Licenses” in Chapter 4. Licenses that don’t appear in this window aren’t enabled.

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 NetWorker cannot recover any files it backed up under the client’s 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 which contains the corresponding index.

Follow these steps:

1. Create and configure the new client.
2. Open the Clients window, and delete the old client name.
3. In an IRIX shell, become root on the NetWorker server.

4. Change directories to */nsr/index*.
5. Use the *mv* command to rename the client's index. For example:

```
# mv old_client_name new_client_name
```
6. Restart the NetWorker daemons, using the *nsrd* command. The daemon *nsrmmdbd* renames all the instances of the save set under the old client name to the new client name.

These steps cause the database daemon *nsrmmdbd(1M)* to rename the client inside its database:

1. Make the old client name an alias of the new name.
2. Kill the NetWorker daemons. To kill the daemons, use this command:

```
# /etc/init.d/networker stop
```
3. Restart the NetWorker daemons. To restart the daemons, use this command:

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

As soon as possible, complete a full backup of the renamed client's files.

Using Recover Access

System administrators control client recover access by configuring the client. The Recover access list in the Clients window lists the systems that can recover the client's files.

These users can recover any files on any client:

- root
- operator
- member of the operator group

Other users can only recover files for which they have read permission, relative to the file mode and ownership at the time the file was backed up. Files recovered by a user other than root, operator, or the operator group, are owned by that user.

Using 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@server_name`. This section contains descriptions of error messages that can appear in the savegroup completion mail. Possible solutions are suggested.

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 IRIX 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 IRIX, the *portmapper* program 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 (for example, *save*, *recover*, *mminfo*, and so on), these messages may appear in your savegroup completion mail:

```
Server not available
RPC error, remote program is not registered (24015)
```

These messages indicate that the NetWorker daemons *nsrd*, *nsrindexd*, *nsrmmmd*, and *nsrmmdbd* are not running.

To restart the *nsr* daemons, enter *nsrd* at the system prompt:

```
mars# /usr/etc/nsrd
```

Saving Remote Filesystems

The following error message might appear in your savegroup completion mail when you are backing up a remote filesystem:

```
All: host hostname cannot request command execution
```

You may also see this message:

```
All: sh: permission denied
```

This means the *nsrexecd* on the client was not configured to allow the server *hostname* to back up its files. See the *nsrexecd(1M)* reference page for information on how to reconfigure *nsrexecd*.

When saving remote filesystems, files, or remote clients, NetWorker first tries to use *nsrexecd*, and if that fails, invokes an *rsh(1M)* to that client as the *rcmd(3N)* user (the default *rcmd(3N)* user is *root*). If the client is supposed to run *nsrexecd*, restart it. If not, make sure the NetWorker server can *rsh* as the *rcmd* user to the remote clients by adding this entry in the clients' */.rhosts* file(s):

```
NetWorker_server rcmd_user
```

Two areas where you may encounter problems with NetWorker and *rsh* are as follows:

- remote shell

The remote shell is determined by the *rcmd* user's login shell (the login shell is determined by */etc/passwd* file). One of the actions of *rshd(1M)* is to read commands from a file. This file is usually named *.<type_of_shell>.rc*. For example, in the C shell this file is called *.cshrc*.

One of the commands *rshd* may find is the PATH variable. The PATH has to have the directory in which the NetWorker executables are located. If the PATH command is not in the remote shell's *rc* file or does not have the correct directory where the NetWorker executables are located, then NetWorker is unable to find client executables.

The following error messages may be sent if the PATH is not correct:

```
clientname probe unexpectedly exited
...: savefs not found
```

If the remote shell's *rc* file does not have the PATH variable set, then the hidden attribute *executable path* in the NSR client resource has to have the full pathname of the NetWorker executables as its value. If the PATH does not have the directory of the executables, then add the directory and give a command to read the *.<remote_shell>rc* file.

Also, the remote shell has to allow an interactive session.

- *rcmd* user and *.rhosts*

The *rcmd* user is the user login name the NetWorker server uses to remotely log in and run commands on the client. The default *rcmd* user is *root*. If you are using *root* as the *rcmd* user then *root@NetWorker_server_name* should be in the *.rhosts* file. Otherwise, *NetWorker_server rcmd_user* has to be in */etc/hosts.equiv* or *.rhosts* file in the home directory of the local user on whom access is being attempted.

The hidden attribute *rcmd* user is displayed in the Clients window when NetWorker is invoked with the *-x* option. If you are using an *rcmd* user other than *root*, enter a valid *user_name* in the *rcmd* user hidden attribute field.

If the above conditions are not met, then these error messages may appear:

```
.... : client rcmd(3) problem for command....
.... : Permission denied
```

Make sure that *rcmd* user is either *operator* or in *operator's* group.

File Changed During Backup

NetWorker backs up the image that is in the filesystem at the time it comes across the file. NetWorker notifies 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 is through being used, 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 these steps:

1. Open the Groups window and select “Details” from the View menu.
2. In the Printer field, enter the name of the printer you are using to print the bootstrap.
3. Click *Apply* to save your changes.

Using Savegroup Completion Warning Messages

Occasionally the savegroup completion message includes one or more additional messages to help the administrator understand why NetWorker performs certain tasks. For example:

```
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 preselected for that client's schedule, NetWorker performs a full backup. This feature is important because it allows you to perform disaster recoveries for that client.

This message can also appear if the server and client's clocks are not synchronized. To avoid this situation, make sure the NetWorker server and client:

- are in the same time zone
- have their clocks synchronized

Also make sure that a non-full save both begins and ends after midnight.

This *savegroup* message can 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 saves the server's bootstrap information along with this group to avoid a long recovery process. As soon as possible you should enable the group your

NetWorker server belongs to, or add the server to another group that is enabled.

Errors in Executing External ASMs

This error message may appear during a backup:

```
client:/ save: external ASM '/xxxxxasm' not found or not executable
```

All external ASMs (see *uasm(1M)* for a complete list of external ASMs) are separate commands, and invoked as needed. For security reasons, external ASM names must end in “asm” and be located in the same directory as the originally invoked command (typically *save* or *recover*). Rename or move the ASM following these conventions to correct this error.

Solving X Error Messages

The following error message can 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 package)
```

This message indicates that the client is not authorized to display NetWorker.

To correct this situation, follow these steps:

1. Enter this command at the client system:

```
client% /usr/bin/X11/xhost Networker_server
```

2. Remotely log in to the NetWorker server and enter this command at the server prompt:

```
Networker_server% setenv DISPLAY client:0.0
```

Note: For shells other than *csh* use the following commands instead:

```
# DISPLAY=client:0.0
# export DISPLAY
```

Moving Indexes

Because the index databases are holey files, the *cp* command creates a file that consumes more disk space than the original file. To move indexes, enter the following command in the */nsr/index* directory:

```
mars# uasm -s -i <client index directory name>|(cd target_dir; uasm -r)
```

Note: On an EFS filesystem, do not use this command; EFS filesystems have no holey files.

Recovering Files From an Interrupted Backup

You cannot recover files from a backup that was terminated by killing the NetWorker daemons, because the media index was not updated before the daemons died. 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, the following message might appear:

```
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 systems. 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 an IRIX client, *nsrexecd* is started, and statements are added to */etc/rc2.d/S95networker* to restart *nsrexecd* each time the client reboots.

Security is increased by a challenge and response scheme, which ensures that only the NetWorker server is initiating connections, and not another command. Commands are restricted to those needed for saving data (the *save* and *savefs* commands are described in Appendix C, “Theory of Operations”).

If you must reconfigure *nsrexecd*, for example, to allow a different NetWorker server to back up the client, follow these steps:

1. In the file */etc/rc2.d/S95networker* on the client, make the changes to the *nsrexecd* startup command.

Note: See the *nsrexecd(1M)* reference page for a description of the command line configuration options.

2. Restart *nsrexecd* by entering these commands:

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


Recovering From a Disk Crash

This chapter contains procedures to follow to prepare for and recover from a major disk crash. Four different types of disk crash situations are covered:

- The primary disk, which contains the operating system and NetWorker binaries, is damaged. This can apply to a client system or a NetWorker server.
- A secondary disk, which contains other filesystems, is damaged. This can apply to a client system or a NetWorker server.
- The NetWorker server's disk, which contains the online indexes (the */nsr* filesystem), is damaged. You have to recover the indexes before using NetWorker to recover any filesystems.
- Your NetWorker fileserver was destroyed. You must recover everything to a new NetWorker server.

If a primary disk suffers a head crash, you may need to replace the disk, boot from miniroot, format and partition the disk, reinstall the operating system and NetWorker binaries, and then recover the affected filesystems. In this case, you must consult the system administration manuals you used to set up your fileserver for the first time before using NetWorker to recover the data on the disk.

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.

This appendix gives a shortcut recovery procedure and explains

- preparing for a crash
- using `recover` or `save set recover`
- recovering a secondary disk
- recovering a primary disk on a client

- recovering */nsr* on a NetWorker server
- recovering to a new server
- disaster recovery with jukeboxes

Note: It is impossible to provide step-by-step instructions on how to recover your system from a disaster, because every site and every disaster is unique. The discussions in this chapter are designed to give you general principles on how to recover a primary or secondary disk, and to help you understand the procedure. They are meant to be examples only, not instructions.

Shortcut

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. If the operating system is lost, reload and boot the system using the same hostname and disk partitioning.
2. Replace the damaged disk, if necessary. Format it, partition it, and make new filesystems. Mount the filesystems in the same locations and with the same names as before.
3. If the NetWorker software was lost, reinstall it from the NetWorker software distribution CD.
4. If the */nsr* filesystem is destroyed on a NetWorker server, use *recoverindex* to recover the NetWorker server's indexes.
5. Recover the lost filesystems by using the save set recover feature or the normal recover process. Recover the client indexes using the normal recover process.

Preparing for a Crash

The ultimate disaster for a system is to lose all the files on its disk. Most sites back up their fileservers daily as insurance against this event. If a system's primary disk suffers a crash, you can rebuild its filesystems with NetWorker, after you reinstall the operating system (if necessary).

If the NetWorker server's filesystem or disk that contains */nsr* is destroyed, the recovery procedure involves an extra step: you must recover the server's online indexes as well as the server's filesystems. The server's */nsr* filesystem 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 system. You may do this as long as you perform the following procedures:

- Give the replacement server the same hostname as the original NetWorker server.
- Get new NetLS licenses for the new system. The licenses for the old system will not work.
- Reinstall NetWorker using the same directory locations for the online indexes as in the original installation.

Tip: Once you understand the procedure for a disaster recovery, make sure you have carefully thought out a disaster recovery plan for your site. If possible, you should test your ability to recover from a disaster at your site.

If you have set up your network and enabled NetWorker to execute automatic, 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, including the indexes for the server itself. The media database, NetWorker configuration files, and part of the server's index is 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.

Take two additional precautionary steps to help you recover from a future crash:

- Keep a file containing hard copies of the daily bootstrap records. Place these 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 the recovery procedure much smoother for you in the future.

Filing the Bootstrap Information

NetWorker sends a record of the index backup to your default printer, so you have a piece of paper with the date, name of the backup media, and save set ID number for bootstrap. If you ever need to recover the server's online indexes, you need the information on this piece of paper. Save this information in a safe place.

The information sent to the printer looks similar to this:

```
March 11 16:42 1994 NetWorker bootstrap information Page 1

date    time    level    ssid    file    record    volume
3/10/94 2:46:02  9 1148868949 56      0 bitbucket.002
3/11/94 2:53:34  9 1148868985 77      0 bitbucket.003
```

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's indexes by using the *saveindex* command. Using this command also sends the bootstrap information to a printer. For example:

```
# /usr/etc/saveindex -c server_name
March 7 03:30 1994 NetWorker bootstrap information Page 1

date    time    level    ssid    file    record    volume
3/01/94 7:44:38 full 1148869706 55      0 bitbucket.008
3/02/94 6:12:09  9 1148869754 48      0 bitbucket.008
3/03/94 6:14:23  9 1148869808 63      0 bitbucket.009
3/04/94 6:29:58  9 1148869870 88      0 bitbucket.009
```

To use the *saveindex* command, you must be root on the NetWorker server.

Filing the Disk Information

Use the disk information commands to find out how the NetWorker server's disks are partitioned and mounted, and make a hard copy of this information. Do the same for any NetWorker clients that have local hard disks. The information you need is

- the output of the *df(1M)* command
- the output of the *prtvtoc(1M)* command
- a copy of the file */etc/lvtab*

Note: For XLV, there is no */etc/lvtab*. To list all XLV volumes on the system and their contents, do one of these actions:

- run *xlvmake* and enter the *show* command to display the current set of XLV objects on your system

or

- run *xlvassemble -n* which will display the same information

Because of a bug in XLV, you'll see the following output from a *prtvtoc* command when you are running on an XLV root volume:

```

xfs-mp9 1# prtvtoc
Printing label for root disk

prtvtoc: /dev/rdisk/xlv_root: Can't get volume header
xfs-mp9 2#
```

For example, the *df* information looks similar to this:

```

mars% df
Filesystem                Type blocks   use  avail %use Mounted on
/dev/root                  xfs  67800    42135 25665 62% /
/dev/dsk/xlv/usr           xfs 6732736 5992080 740656 89% /usr
/dev/dsk/xlv/f             xfs 86673816 58296424 28377392 67% /f
/dev/dsk/xlv/e            xfs 31405776 27886176 3519600 89% /e
/dev/dsk/xlv/d            xfs 3907872 2477400 1430472 63% /d
/dev/dsk/xlv/b            xfs 31316992 27869408 3447584 89% /b
```

The *prtvtoc(1M)* command gives you information about how each disk is partitioned. It takes device volume header names as arguments; you can construct these from the output of *hinv(1M)*. For example:

```

mars% hinv -c disk
Integral SCSI controller 131: Version WD33C95A
Integral SCSI controller 130: Version WD33C95A
Integral SCSI controller 4: Version WD33C95A
Integral SCSI controller 3: Version WD33C95A
Integral SCSI controller 2: Version WD33C95A
Disk drive: unit 4 on SCSI controller 2
Disk drive: unit 3 on SCSI controller 2
Disk drive: unit 2 on SCSI controller 2
Disk drive: unit 1 on SCSI controller 2
Integral SCSI controller 1: Version WD33C95A
Disk drive: unit 4 on SCSI controller 1
Disk drive: unit 3 on SCSI controller 1
Disk drive: unit 2 on SCSI controller 1
Disk drive: unit 1 on SCSI controller 1
Integral SCSI controller 0: Version WD33C95A

```

One device volume header name is constructed from each of the disk drive lines. The device name for the first disk drive line in the example is */dev/rdisk/dks2d4vh*.

Give the *prtvtoc* command with the device names as arguments (the */dev/rdisk/* prefix can be omitted). For the example above, the command is:

```

mars# prtvtoc dks2d4vh dks2d3vh dks2d2vh dks2d1vh \
dks1d4vh dks1d3vh dks1d2vh dks1d1vh
* /dev/rdisk/dks2d4vh (bootfile "/unix")
*      512 bytes/sector
*      94 sectors/track
*      15 tracks/cylinder
*      34 spare blocks/cylinder
*      2858 cylinders
*      3 cylinders occupied by header
*      2855 accessible cylinders
*
* No space unallocated to partitions

Partition Type Fs Start: sec (cyl) Size: sec
(cyl) Mount Directory
0 efs 4128 ( 3) 49152 ( 35.7)
1 raw 55040 ( 40) 524288 ( 381.0)

```

6	efs	579328 (421.0)	3353280 (2437.0)
7	lvol	4128 (3)	3928480 (2855)
8	volhdr	0 (0)	4128 (3)
10	volume	0 (0)	3932608 (2858)
...			

The file `/etc/lvtab` describes the logical volumes used for LV. This file, along with the information from the `df` command, tells how the disk partitions are combined and used and how to map names onto partitions. `/etc/lvtab` looks similar to this:

```

mars% cat /etc/lvtab
lv0:/b volume:stripes=4:devs=/dev/dsk/dks1d2s7,/dev/dsk/dks1d5s7, \
/dev/dsk/dks2d1s7, /dev/dsk/dks2d2s7
lv4:/a volume:stripes=4:devs=/dev/dsk/dks1d3s7,/dev/dsk/dks1d4s7, \
/dev/dsk/dks2d3s7, /dev/dsk/dks2d4s7

```

Print and file this information in case you ever have to recover from a disk crash. If a disk is destroyed in a head crash, you can rebuild it exactly as it was and recover the filesystems to their original state using the hard copy information from the disk information commands.

Using Recover or Save Set Recover

You can use 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 can demultiplex the filesystem save sets from the same backup volume in parallel, thus reading each backup volume only once. Thus you should try to 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).

Note: With release 4.1.1, 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.

The big advantage to using the save set recover feature is that it takes less time to browse and mark 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.

Note: Any time you have to recover the primary disk (for example, root), do so in single-user mode, from the system console, not from the X window system. Before starting this procedure, make sure all your filesystems are mounted.

The two ways to use the save set recover feature are as follows:

- Run the *nwadmin* program and choose “Recover” from the Save Set menu to open the Save Set Recover window.
- If you are not using the X window system, run *recover* with the *-S ssid* option from the system prompt. Refer to the *mminfo(8)* reference page for instructions on how to find the save set(s) you want to recover. Refer to the *recover(8)* reference page for detailed instructions on using the *recover -S* command.

In the Save Set Recover window, note 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 in order to bring the system back to the state it was in before the system crash.

Remember, a level 0 backup is a full backup, and 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. Figure B-1 diagrams several backups over time.

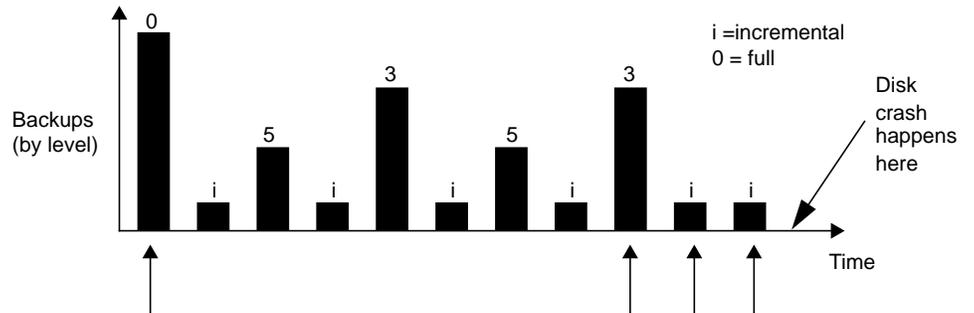


Figure B-1 Backups Over Time

In this diagram, the bars represent the level backups since the last full backup. The arrows point to the save sets you would need in order to recover your filesystems from the disk crash. For example, suppose you wish to recover the save set `/nsr/index/xfs4.engr.sgi.com`, shown in the Recover Save Set window shown in Figure B-2.

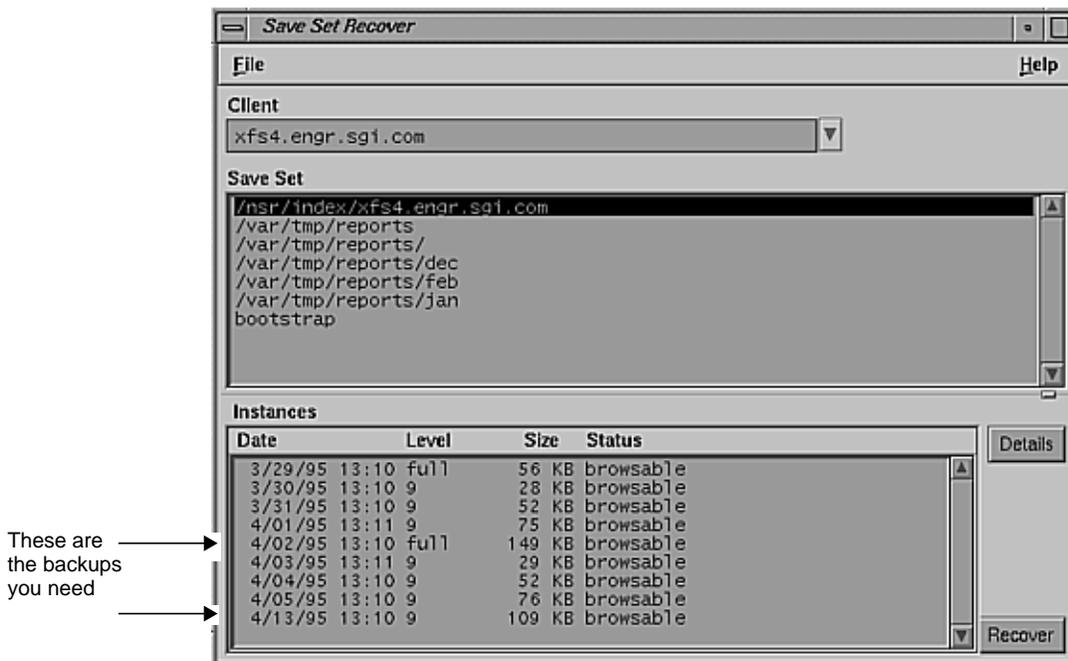


Figure B-2 Recover Save Set Window

When you select `/nsr/index/xfs4.engr.sgi.com` in the Save Set scrolling list, the Instances scrolling list displays its backup history. You must mark (for a recovery) the most recent full backup of `/nsr/index/xfs4.engr.sgi.com` and the most recent level 9 in order to restore `/nsr/index/xfs4.engr.sgi.com` to the state it was in before the system crash. In other words, you must mark the appropriate save set levels for each filesystem you are recovering.

Normally during disaster recovery, you force the recover program to overwrite existing files. When you use save set recover, marking the save set level is even more important, because the same file might be recovered multiple times with each successive version coming from a later save set. And since each file in each save set is recovered with save set recover, files or directories that have been previously deleted or renamed between backups (with the `mv` command) are still recovered and therefore must be manually deleted. You might 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, you might be better off using the normal recover instead of save set recover. If, however, you only need the last full backup of a save set, save set recover should be your method of choice.

By contrast, if you decide to use the Recover command from the *nwrecover* program, you recover your filesystems exactly as they were as of the last backup before the crash, and NetWorker reads the minimum amount of tape to recover the files. The disadvantages are as follows:

- It takes more time to browse and mark 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 might run out of swap space if you add too many filesystems at once to your list of data to recover.

In summary, we recommend the save set recover feature for faster disaster recovery if

- you can determine the correct save sets to recover
- there are only a few save sets to recover for each filesystem
- you are not recovering the NetWorker indexes
- recovering extra files is acceptable (not an issue if you are recovering from full backups only)

Use the normal recovery procedure if

- you cannot determine which save sets to recover
- you must recover files from many backups to restore the filesystem to an acceptable level
- you are recovering the NetWorker indexes
- recovering extra files is not acceptable, and you are recovering from several backups, not just a full one

Recovering a Secondary Disk

This section gives you an outline of how to recover a secondary disk using NetWorker. The outline applies to either a NetWorker server or a client.

The example in Figure B-3 assumes the primary disk is still operational, so the system has an operating system and can run NetWorker. However, a secondary disk is lost due to a head crash.

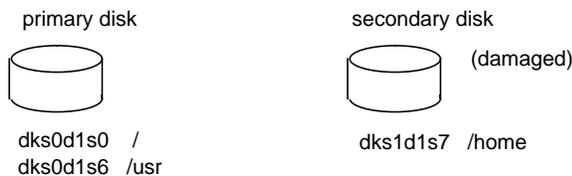


Figure B-3 Damaged Secondary Disk

If the disk is damaged, replace it with a new disk of the same type. Try to get a disk of the exact same size as the old one. You need a disk large enough to hold all the filesystems to be recovered.

1. Install the replacement disk. Make sure that IRIX recognizes the new disk.
2. Label the new disk, partition it, and make filesystems. This procedure is outlined in the section “Repertitioning a Hard Disk” in Chapter 8 of the *IRIX Advanced Site and Server Administration Guide*. Use the hard copy of the disk information to remember how large each partition was. (See the section “Filing the Disk Information” in this appendix.)
3. Recreate logical volumes on this disk, if used. Use the procedure in the section “Logical Volumes and Disk Striping” in Chapter 8 of the *IRIX Advanced Site and Server Administration Guide* and again use the hard copy of the disk information.
4. Invoke the `recover` command:

```
venus# /usr/etc/recover
```
5. Select one filesystem to recover. You should recover one filesystem at a time, because NetWorker adds all the files in one filesystem but stops at a mount point, and you may run out of swap space if you add too many filesystems at once to your list of data to recover.

6. Add the directories and files under the filesystem to be recovered:

```
NetWorker> add /home
```

7. Recover the filesystem:

```
NetWorker> recover
```

```
/home is being recovered into its original location
Requesting 33023 files, this may take a while...
```

8. Repeat steps 5 through 7, adding and recovering one filesystem at a time. Remember to open a main window on your screen (or use the *nsrwatch* command) so you can monitor NetWorker requests for backup media during the recovery.

9. Exit the *recover* program:

```
NetWorker> quit
```

Recovering a Primary Disk on a Client

In the example in Figure B-4, a disk with the operating system or the NetWorker binaries is damaged.

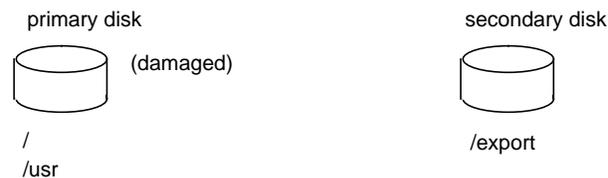


Figure B-4 Damaged Primary Disk

After replacing the damaged disk, format it and reinstall the operating system, using the original software distribution. Next, reinstall NetWorker from the original software distribution, reinstall the NetWorker licenses, and activate the licenses as described in the section “Installing and Enabling NetWorker Software on Servers” in Chapter 2.

With the operating system and NetWorker back in place, you are ready to start recovering the remainder of the data lost from the disk.

1. 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 recreate them.
2. Invoke NetWorker to recover each filesystem on the disk being recovered, one at a time. For example:

```
venus# /usr/etc/recover
NetWorker> add /
```

3. Use the `list` command to check the list of directories and files you are about to recover. Use the `force` command so that you are not prompted every time a file naming conflict occurs. For example:

```
NetWorker> list
NetWorker> force
NetWorker> recover
```

4. Recover the lost filesystems, one at a time. Open the main window or use the `nsrwatch` command to monitor NetWorker requests for backup media during the recovery.

Note: Always reboot a system after recovering a primary disk.

Recovering `/nsr` on a NetWorker Server

This section addresses the case in which 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 commands, these have to be reinstalled first. See “Recovering a Primary Disk on a Client” earlier in this appendix.

The next important step is to recover the server’s indexes from the backup media, using the `recoverindex` command. The `recoverindex` command asks you for the bootstrap save set ID (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 save set ID.

For example, save set ID 1148869870 below is the most recent bootstrap backup:

March 11 16:42 1994 NetWorker bootstrap information Page 1

date	time	level	ssid	file	record	volume
3/10/94	2:46:02	9	1148868949	56	0	bitbucket.002
3/11/94	2:53:34	9	1148868985	77	0	bitbucket.003

If you do not have this piece of paper, you can still recover the indexes by finding the save set ID using the *scanner* command. (See the section “Finding the Bootstrap Save Set ID” in this appendix.)

You may need more than one backup medium to recover the server’s indexes. During the recovery, you can use the *nsrwatch* command or open the main window to watch for pending messages requesting backup media.

With the operating system and NetWorker in place, recover the indexes from the backup medium:

1. Obtain the printout of the bootstrap save set ID information.
2. Retrieve the backup medium that contains the most recent backup named bootstrap, and load it into the server’s device.
3. Use the *recoverindex* command to extract the contents of the bootstrap backup. For example:

```
mars# /usr/etc/recoverindex
recoverindex: Using mars as server
```

NOTICE: recoverindex 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. recoverindex is not used to recover NetWorker clients’ on-line indexes; normal recover procedures may be used for this purpose. See the recoverindex(1M) and nsr_crash(1M) man pages for more details.

```
What is the name of the tape drive you plan on using
[dev/rmt/tps1d6nrnsv]?
Enter the latest bootstrap save set id []: 1148869870
```

```
Enter starting file number (if known) [0]: 88
Enter starting record number (if known) [0]: 0
```

```
Please insert the volume on which save set id 1148869870
started into dev/rmt/tps1d6nrnsv. When you have done
this, press <RETURN>:
```

```
Scanning dev/rmt/tps1d6nrnsv for save set 1148869870;
this may take a while...
scanner: scanning 8mm 5GB tape space.006 on
dev/rmt/tps1d6nrnsv
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/space/
nsrindexasm -r nsr/index/space/db
scanner: ssid 449955156: scan complete
scanner: ssid 449955156: 31 KB, 10 files
nsr/index/space/db: file exists, overwriting
uasm -r nsr/index/
uasm -r nsr/mm/
uasm -r nsr/
uasm -r
space: 31 records recovered, 0 discarded.
nsrindexasm: Building indexes for mars...
nsrindexasm: Caching save times for mars...
8mm 5GB tape space.006 mounted on dev/rmt/tps1d6nrnsv,
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.
# nsrindexasm: Pursuing index pieces of
nsr/index/space/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
.db: file exists, overwriting
: 25711 records recovered, 0 discarded.
nsrindexasm -r .db
```

```
.db: file exists, overwriting
nsrindexasm: waiting for lock on ../db.SCAVENGE
nsrindexasm: lock on ../db.SCAVENGE acquired
Received 2 files from NSR server 'mars'
: 733 records recovered, 0 discarded.
nsrindexasm: Building indexes for mars...
nsrindexasm: Caching save times for mars...
nsrindexasm: Suppressing duplicate entries in mars - 50
duplicates discarded.
```

The on-line index for space is now fully recovered.

Notice that the shell prompt appears once bootstrap is recovered. You can use NetWorker commands such as *nsrwatch* to watch the progress of the server, or *nwadmin* to bring up the NetWorker Administrator window during the recovery of the index. Open a new shell to monitor the recovery so that the *recoverindex* output does not display on top of the *nsrwatch* output.

Replacing the /nsr/res Directory

The *recoverindex* command also recovers the */nsr/res* directory, which is used by NetWorker to store configuration information such as the list of NetWorker's clients and registration information. However, this directory, unlike the indexes, cannot be overwritten or relocated; instead, the recovered */nsr/res* directory is renamed */nsr/res.R*.

After *recoverindex* has finished, this final message appears:

```
nsrindexasm: The on-line index is now fully recovered.
```

To complete the recovery of the */nsr/res* directory, you need to shut down NetWorker, move the recovered */nsr/res* directory into its original location, and then restart NetWorker:

1. Shut down the NetWorker server using the *nsr_shutdown* command:
2. Save the original */nsr/res* directory, and move the recovered version into the correct location.

```
# /usr/etc/nsr_shutdown
# cd /nsr
# mv res res.orig
# mv res.R res
```

3. Restart the NetWorker server

```
# cd /
# nsrd
```

When NetWorker restarts, it uses the recovered configuration data.

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 the section “Recovering a Secondary Disk” earlier in this appendix.

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` command.

For example:

1. Place the most recent medium used for backups in the server device.
2. Read the contents of the backup medium with the `scanner` command:

```
mars# /usr/etc/scanner /dev/rmt/tps1d6nrnsv
```

Substitute the pathname for your server device.

The `scanner` command displays the contents of the backup media, for example:

```
# scanner /dev/rmt/tps1d1nrnsv
scanner: scanning 8mm 5GB tape atlas.005 on /dev/rmt/tps1d1nrnsv
client name   save set      save time    level  size  files  ssid S
atlas.engr.s  /             3/10/94  3:37  f    6072660  1053  3388 E
atlas.engr.s  /usr         3/10/94  3:37  f    57315808  2985  3387 E
atlas.engr.s  /usr/people  3/10/94  3:33  f    504417356  19737  3372 S
atlas.engr.s  /usr/people  3/10/94  3:33  f    1340638812  41067  3372 E
scanner: done with 8mm 5GB tape atlas.005
```

In this example, the bootstrap save set ID is 1340638812. Once you find the most recent bootstrap save set ID, you can use the `recoverindex` command to

recover the server's index. 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.

Note: The new NetWorker server must have the same hostname as the old NetWorker server.

To recover NetWorker to a new server, follow these steps:

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. See Chapter 11, "Using NetWorker with Jukeboxes," and "Enabling Autochangers" in Chapter 2 earlier in this guide for jukebox device driver installation and testing.

2. Obtain new NetWorker licenses for all affected systems; see Table 2-2 in Chapter 2, "Installing NetWorker."
3. Find the printout of the bootstrap save set ID information from the old server. You will need it for the next two steps.
4. Retrieve the backup media that contains the most recent backup named bootstrap, and load it into the new server's device.
5. Use the *recoverindex* command to extract the contents of *the bootstrap backup*.

The NetWorker daemons should start up on the new server. Messages like the following should appear on the new server:

```
new_server syslog: NetWorker Server: (notice) started
new_server syslog: NetWorker Registration: (notice)
invalid auth codes detected.
new_server syslog:
new_server syslog: The auth codes for the following
licenses enablers are now invalid.
```

```
new_server syslog: The cause may be that you moved the
NetWorker server to a new computer.

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

new_server syslog:

new_server syslog: License enabler #xxxxxx-xxxxxx-xxxxxx
(NetWorker TurboPak)
```

6. Contact Silicon Graphics to obtain the new licenses you need.

Finally, 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 entering `saveindex -1 full` at the system prompt. (Or, if you prefer, perform a full backup of the new server as soon as possible.)

Disaster Recovery With Jukeboxes

To use jukeboxes during disaster recovery, follow these steps:

1. Read the disaster recovery procedures listed in the reference page *nsr_crash(1M)*. Perform all steps up to giving the `recoverindex` command. If only one volume is needed to recover your NetWorker file indexes, follow the instructions in *nsr_crash(1M)*.
2. Run the `jbm_enabler` command to add the jukebox.
3. Give the command `nsrjb -H` to reset 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(1M)*, determine which volumes are needed to retrieve the NetWorker file indexes. Load these volumes into the jukebox.

5. Give the command `nsrjb -I` to reinventory the jukebox. To speed up this process, use the command with the `-S` flag and list the slots where you placed the backup volumes. You must list the slots continuously; for example, `nsrjb -I -S 1-3`. To inventory slots out of order (for example, 2, 1, and 4), you must use the `nsrjb -I -S` command separately for each slot.

All the volumes currently loaded in the jukebox are marked with an asterisk because there is no media database.

6. Load the first volume that `recoverindex` requests into the first drive in the jukebox. Give this command:

```
# /usr/etc/nsrjb -l -n -s slot -f device_name
```

`slot` is the slot where the first volume is located and `device_name` is the pathname of the first drive.

7. Run `recoverindex`. Provide the same device name as in step 6 above and the last save set ID requested by `recoverindex`. At this point, NetWorker recovers the file indexes.
8. After the indexes have been recovered, give the command `nsrjb -u`.
9. To recover whole filesystems, use the save set recover feature as explained earlier in this chapter to determine which volumes will be needed. If any of the volumes are not located in the jukebox, load them into the jukebox and inventory those slots. Continue with this process until all your filesystems have been restored.

Theory of Operations

This appendix is for the system administrator or operator who is familiar with the *IRIX NetWorker Administrator's Guide* and is interested in more information about NetWorker and how it works. You should be familiar with IRIX administrative concepts, the X Window System interface, and the *nsradmin* command.

The NetWorker reference pages provide complete and detailed information to help you administer NetWorker. The reference pages list the commands in alphabetical order, so that you can verify them easily.

This appendix gives a pictorial overview of the major components of NetWorker and how they are controlled. The pictures illustrate basic points; they are not meant to be detailed or precise. This appendix lists the various NetWorker reference pages that contain details about each section.

You should read this appendix while sitting at your workstation so that you can read the IRIX and NetWorker reference pages, execute NetWorker commands (especially *nsradmin*), and verify the pictures (using IRIX administrative commands like *ps*).

The basic picture elements used in this appendix are shown in Figure C-1.

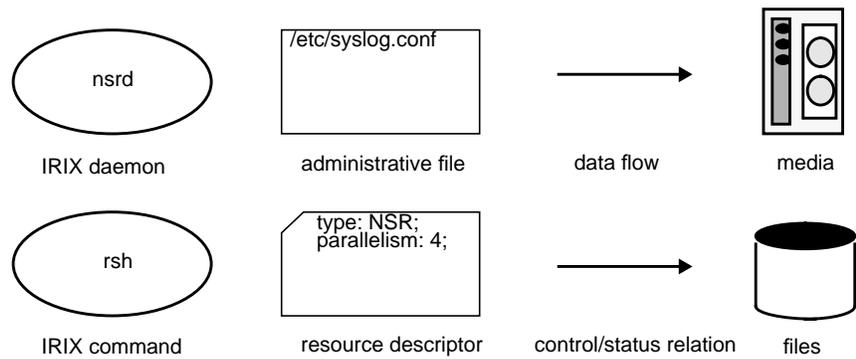


Figure C-1 Major Components of NetWorker

This appendix explains

- NetWorker media format
- the NetWorker system at rest
- browsing and modifying resources
- recovering files
- backing up the entire system
- automatic network-wide backups
- event processing
- media management
- index management
- disaster recovery
- NetWorker security

NetWorker Media Format

The NetWorker media format is designed to meet these requirements:

- the ability to multiplex data from several clients simultaneously to maximize performance
- machine architecture independence (such as byte order differences) through the eXternal Data Representation (XDR) format
- filesystem independence, allowing NetWorker to back up heterogeneous clients with different operating systems and filesystems
- the ability to fully utilize media capacity by appending to the media until it is full, and then continuing to additional media
- support for quick seeking on media, by writing occasional file marks
- special handling for certain types of files via the Application Specific Module (ASM) architecture
- the ability to track media through media labels

No other existing media format meets these requirements.

The NetWorker media format is fully documented in the *mm_data(5)* reference page. Third parties are welcome to adopt the format.

The NetWorker System at Rest

Four daemons continue running on the NetWorker server system even when it is at rest. The two internal server daemons, *nsrmmmd* and *nsrmmdbd*, are not meant to be *networker* client-accessible; their only clients are forked copies of *nsrd*. As these internal daemons execute, they may send status messages (or “events”) to the controlling *nsrd*.

The controlling *nsrd* is the root of all NetWorker functionality. It receives status information from the internal daemons and reports the status to commands such as *nsrwatch* and *networker*. The *nsrd* daemon allows the browsing and editing of NetWorker resource descriptors by clients using the *nsradmin* command. Finally, *nsrd* accesses save and recover sessions from network clients and forks “agent” copies of itself to handle each session.

From a monitoring point of view, the NetWorker main window and *nsrwatch* are functionally equivalent. They monitor the dynamic state of the NetWorker server. A very important field in these monitors is the Pending display, which shows you what the server needs to make progress.

The *networker* daemon provides an X Window System graphical interface for all the system configuration tasks. The *nsradmin* daemon provides a character-based administrative interface to the NetWorker server that can be used from any ASCII terminal.

The four daemons that are always running on the NetWorker server and three network-based monitors are shown in Figure C-2.

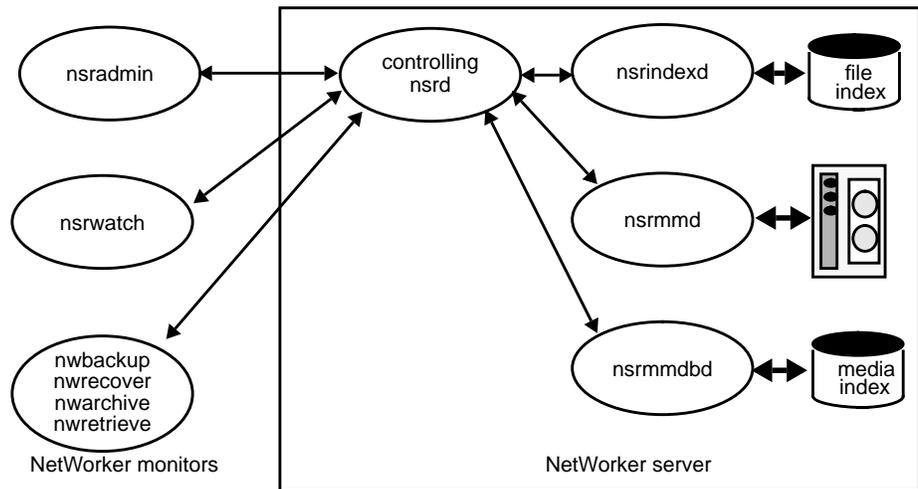


Figure C-2 NetWorker Monitors and Daemons

Table C-1 lists the relevant reference pages.

Table C-1 Reference Pages for NetWorker at Rest

Reference Page	Comments
<i>nsradmin</i> (1M)	Controls all aspects of NetWorker server administration
<i>nsrwatch</i> (1M)	Character-based monitor

Table C-1 (continued) Reference Pages for NetWorker at Rest

Reference Page	Comments
<i>nwadmin</i> (1M)	X Window System interface for NetWorker system administration
<i>nsrd</i> (1M)	NetWorker server
<i>nsrmmmd</i> (1M)	Internal daemon interfaces to all backup devices
<i>nsrmmdbd</i> (1M)	Internal daemon that manages the media index

Browsing and Modifying Resources

You can configure NetWorker by using the X Window System interface or by editing the configuration information (called *resources*) from an ASCII terminal.

If you are not using the X Window System, you can browse and edit the NetWorker server's resources with the *nsradmin* command. You can either use your IRIX editor to browse and modify the NetWorker server's resource descriptors, or you can use the visual mode of *nsradmin*. See the *nsradmin*(1M) reference page for more information.

If you are using your editor, *nsradmin* copies the resource descriptors from the server into a temporary file on your system. When you are finished making changes, *nsradmin* compares any edits you may have made with the original copies and prompts you for confirmation for any additions, deletions, or changes you may have made to the resource descriptors.

Changes are passed back to the server and are immediately applied to the NetWorker server subsystems that are affected. This way, administrators do not have to kill and restart daemons or reboot systems in order to have changes applied. In addition, the changes can be applied from a client workstation using the administrator's favorite editor. All administrative aspects of a NetWorker server are handled by resource modifications.

The resource types that describe the NetWorker server configuration are shown in Figure C-3.

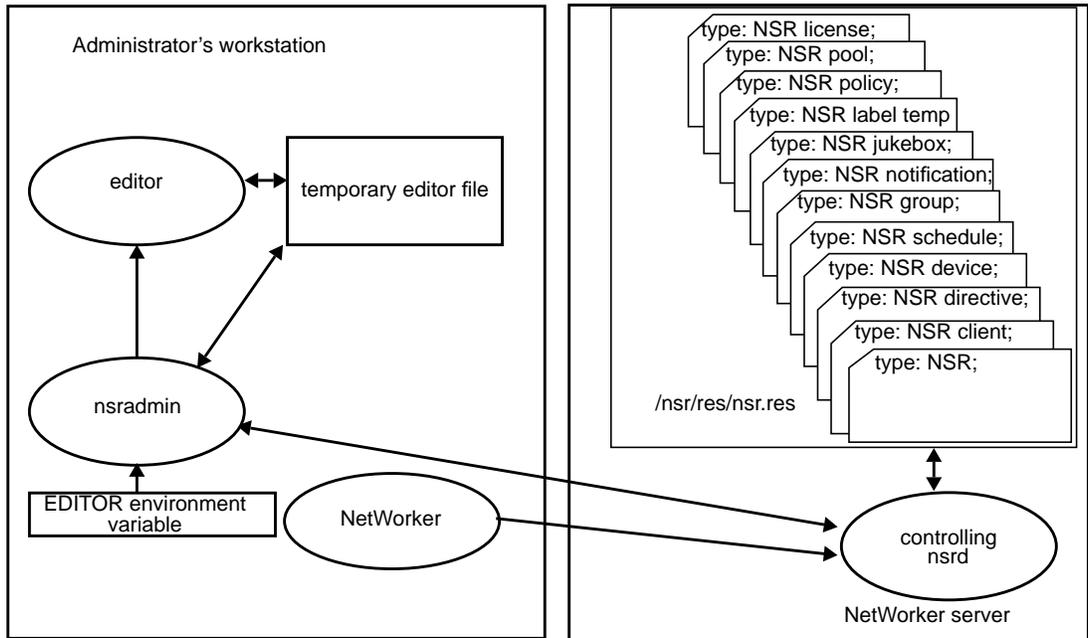


Figure C-3 NetWorker Resource Types

Table C-2 lists the relevant reference pages.

Table C-2 Reference Pages for Browsing and Modifying Resources

Reference Pages	Comments
<i>nwadmin</i> (1M)	X Window System interface
<i>nsradmin</i> (1M)	Invokes your editor on resource descriptor copies
<i>nsr_resource</i> (5)	Describes resource file format and types
<i>nsrd</i> (1M)	NetWorker server checks and applies changes

Backing Up Files

This section describes the operation of backups, called “saves.”

The *save* command traverses a client’s filesystems and backs up a client’s files, subject to directives. See “Using Directives” in Chapter 6 for information.

The *save* daemon first initiates a session with the server’s controlling *nsrd*. The server accepts a connection if it originates from a privileged port on a client listed in the server’s “NSR client” resource descriptor. Once the network connection is accepted, a forked agent *nsrd* handles all subsequent requests from the client.

After session establishment with the server, *save* reads files and sends a save set to its agent *nsrd*. Upon completion, the client calls upon the server to commit the save set. To commit the save set, the agent *nsrd* instructs the *nsrmmmd* to commit the save set data to the backup media. At this point, the agent returns from the commit call. *save* then terminates the session and both *save* and the agent *nsrd* exit.

Under normal circumstances, in order to acquire a privileged port, *save* must be a “set uid root” command. Upon acquiring the privileged port, *save* reverts back to the invoking the user ID, unless the user is *operator* or the user is in the group *operator*.

The controlling *nsrd* does the original access control, then forks an agent to handle the actual saved data. By forking one agent per save session, the server can handle an arbitrary number of sessions simultaneously.

Figure C-4 diagrams a typical *save* session between client and server.

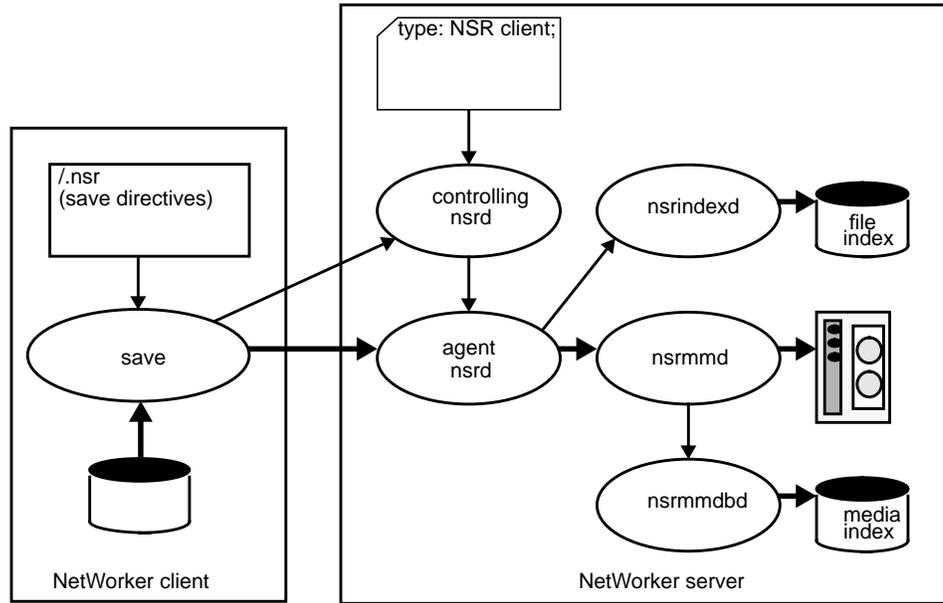


Figure C-4 Typical save Session

Table C-3 lists the relevant reference pages.

Table C-3 Reference Pages for Backing Up Files

Reference Pages	Comments
<i>nsr</i> (5)	Documents <i>save</i> directives
<i>nsr_client</i> (5)	Documents the resource descriptor type "NSR client"
<i>save</i> (1M)	Command that saves files to the server
<i>nsrd</i> (1M)	NetWorker server daemon checks access control
<i>nsrexecd</i> (1M)	NetWorker client execution daemon
<i>rcmd</i> (3N)	Discusses privileged ports
<i>nsr</i> (1M)	Security section discusses policies and issues
<i>nsr_data</i> (5)	Describes basic data types passed from client to server

Recovering Files

The *recover* command is the counterpart of *save* and subject to similar security and session establishment policies. Once a session is established, two major functions of *recover* are used: file browsing and recovery.

Saved files are browsed by using commands that are familiar to the IRIX user. Since file attributes are kept online, the browsing is possible even when no backup volume is mounted on a device. During browsing, media information may be needed if the user requests information associated with the “Versions” option in the Recover window’s View menu.

As users browse, they may build a recover list. Eventually the user may give the *recover* command and the client submits the recovery list to the agent *nsrd* on the server. At this point, the file and media indexes are used to determine the backup volume(s) and position(s) of the desired file(s).

The *nsrmmmd* requests backup volumes and reads them until the entire file recovery list is processed.

Clients browse the online indexes via remote procedure calls to their agent daemons. The device and *nsrmmmd* are involved only when files are actually requested for recovery. Thus, NetWorker supports multiple concurrent browsers.

Figure C-5 diagrams a typical *recover* session between client and server.

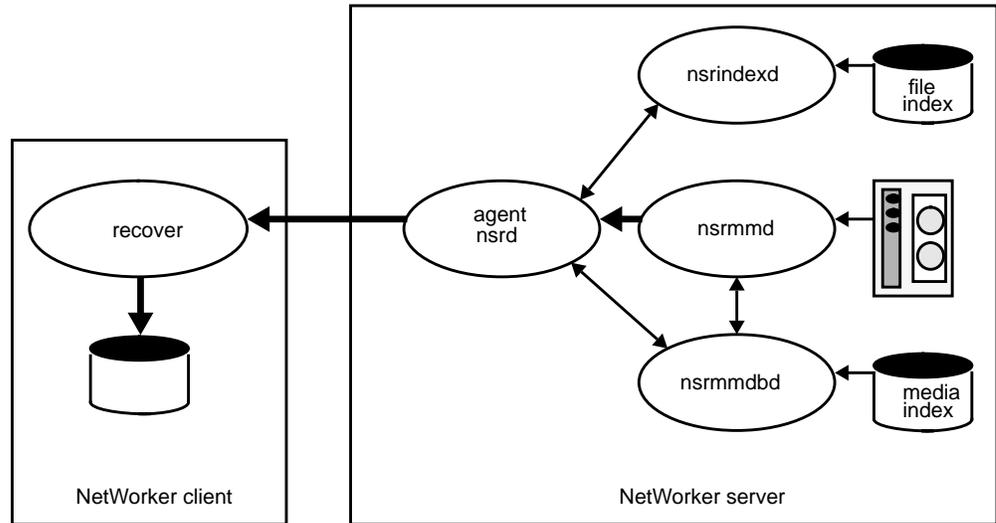


Figure C-5 Typical *recover* Session

The relevant reference page is *recover(1M)*, which documents browsing and recovery subcommands.

Backing Up the Entire System

While it is useful that the *save* command allows individuals to back up their directories, an unattended backup system must back up whole systems according to predefined backup levels. The *savefs* command accomplishes this task.

The *savefs* daemon determines which backup schedule to use for a client by looking for the schedule name in the "NSR client" resource. It then consults the appropriate "NSR schedule" resource to determine which backup level to use (full, incremental, or a level 1–9). Given the save level and a filesystem, *savefs* saves only files within the filesystem that have been modified since the most recently recorded lesser-level save. Upon completion, the current save level and the time at which *savefs* began are recorded in the media database.

If the client has been configured to back up all its filesystems, *savefs* uses all local filesystems.

After compiling information about the local filesystems and when they were last saved, and determining the desired save level, the *savefs* command invokes *save* against each of the client's filesystems.

Figure C-6 diagrams the *savefs* command saving two filesystems in parallel.

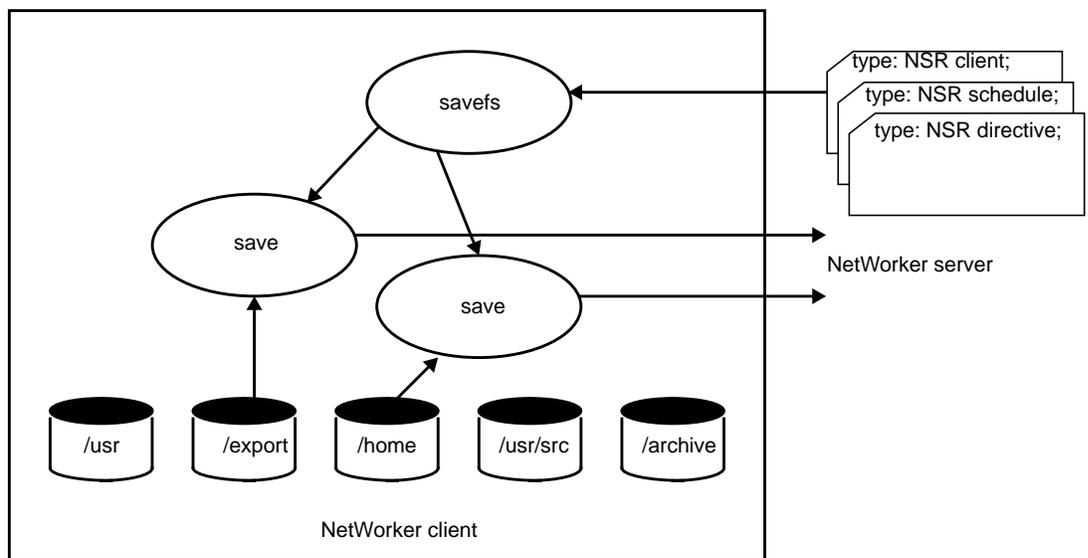


Figure C-6 Saving Two Filesystems in Parallel

Table C-4 lists the relevant reference pages.

Table C-4 Reference Pages for Backing Up the Whole System

Reference Page	Comments
<i>savefs</i> (1M)	Saves the client's filesystems
<i>nsr_client</i> (5)	Determines the client configuration and which filesystems are backed up
<i>nsr_directive</i> (5)	Determines which save directives apply to the client

Table C-4 (continued) Reference Pages for Backing Up the Whole System

Reference Page	Comments
<i>nsr_schedule(5)</i>	Determines save level
<i>save(1M)</i>	Does the actual save, given many arguments

Automatic Network-Wide Backups

NetWorker uses the “NSR group” resource to find the start time of a network-wide backup. The pre-configured time is 3:33 each morning. Each “NSR client” resource describes the groups to which it belongs.

The controlling *nsrd* starts the *savegroup* command at the appropriate time for each group that is enabled. The *savegroup* daemon uses the “parallelism” attribute of the “NSR” resource to determine how many client sessions to save concurrently.

Upon completing all client saves, *savegroup* invokes the *saveindex* command to insure that the NetWorker server index is safely backed up.

The NetWorker server *nsrd* starts the nightly saves by invoking *savegroup*. Each client’s *savefs* is initiated by *savegroup* in an orderly manner. After the clients are backed up, *saveindex* backs up the server’s index.

The daemon *savegroup* also uses *nsrexecd* to run *savefs* on clients. This daemon uses *rshd* if *nsrexecd* is not present and running on a client.

The NetWorker server *nsrd* starts the nightly save by invoking *savegroup*, as illustrated in Figure C-7.

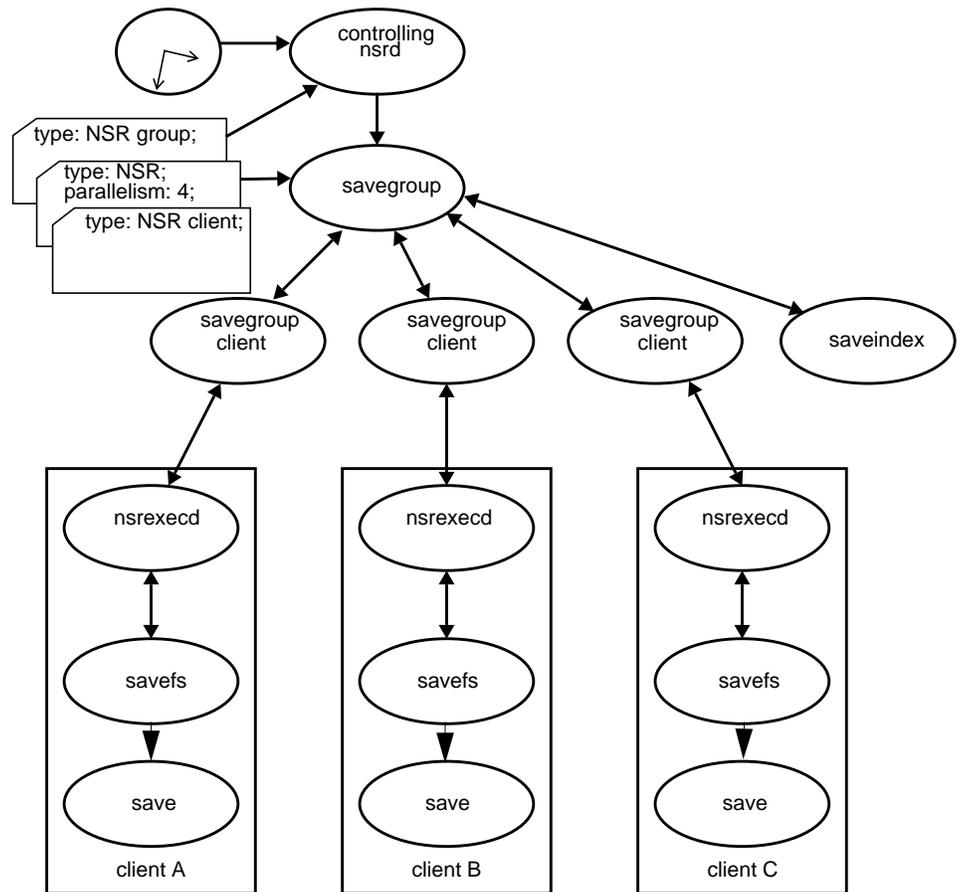


Figure C-7 Starting the Nightly Save

Hint: Make liberal use of incremental save levels, which are very efficient, since they take minimal backup media space and run quickly.

Note: The daemon *saveindex* does a level 9 backup of the indexes to promote faster recovery after a disaster.

Hint: For unattended backups, a NetWorker server with two backup devices is worth more than twice as much as a NetWorker server with only one backup device. Often the NetWorker server with two backup devices is more productive than two NetWorker servers with only one device each. Using a jukebox is the most efficient way to complete unattended backups.

A nightly network backup is shown in Figure C-8. (Not shown are the server's controlling *nsrd* and *nsrmmdbd*, and the client's *nsrexecd*.)

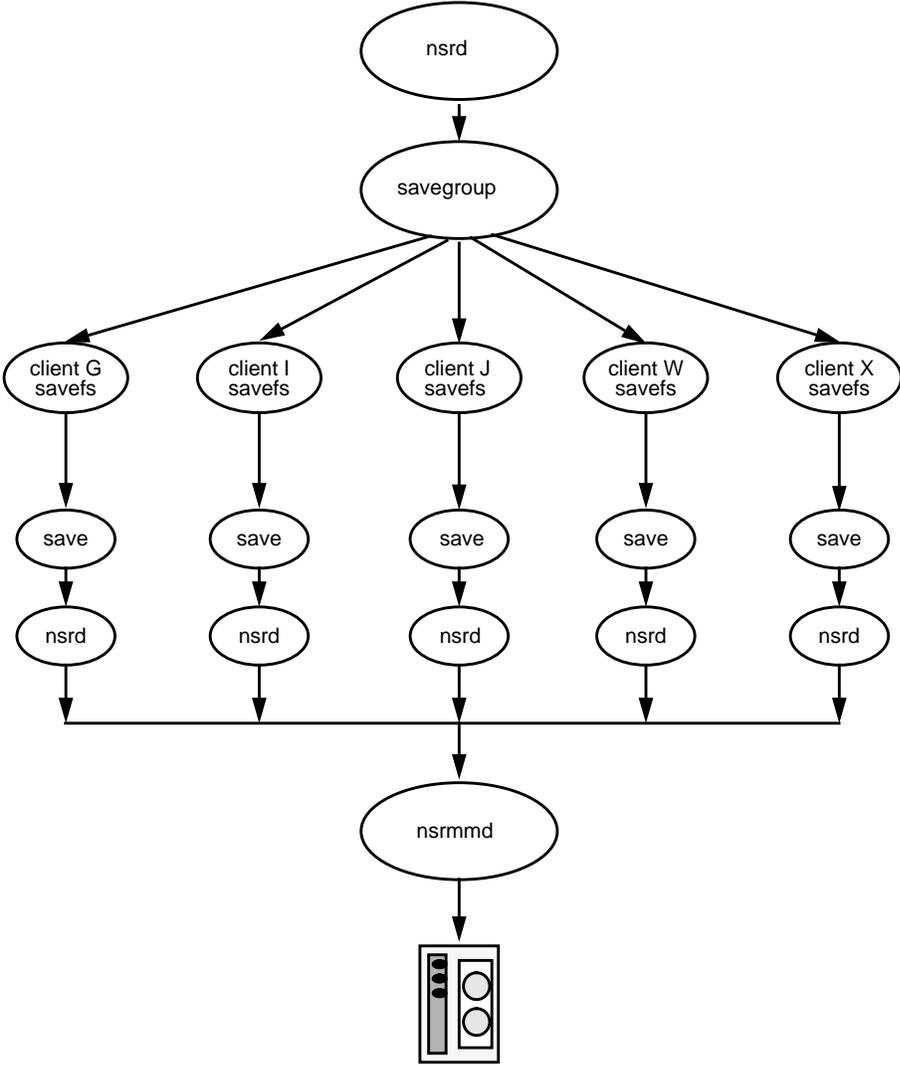


Figure C-8 Nightly Network Backup

Table C-5 lists the relevant reference pages.

Table C-5 Reference Pages for Automatic Nightly Backups

Reference Page	Comments
<i>savegroup</i> (1M)	Conducts the nightly network saves
<i>nsr_service</i> (5)	Parallelism attribute controls how many clients save at once
<i>nsr_group</i> (5)	Selects the time of nightly saves
<i>nsrexecd</i> (1M)	Remote execution system used to start <i>savefs</i>
<i>rshd</i> (1M)	Remote execution system used to start <i>savefs</i> if <i>nsrexecd</i> is not present
<i>savefs</i> (1M)	Saves each client's filesystems
<i>saveindex</i> (1M)	Saves the server's index; detailed in a later section

Event Processing

Besides reporting routine status to the controlling *nsrd*, the NetWorker subsystems are set to report major events. The administrator can cause any IRIX command to be invoked when an event occurs. He or she does this simply by modifying an "NSR notification" resource on the NetWorker server.

When an event occurs, the controlling *nsrd* inspects all "NSR notification" resources. An action is taken if the event type matches any event type listed in the hidden "event" attribute list, and if the event's priority matches any priority listed in the hidden "priority" attribute. Therefore, one event may trigger any number of actions, and one "NSR notification" may match any number of events.

All NetWorker daemons post events to the controlling *nsrd*. As each event is posted, the controlling *nsrd* matches the event against all NSR notification resources and acts on every match by executing the corresponding IRIX command. NetWorker defaults provide interfaces to the *syslog* system and electronic mail.

Figure C-9 diagrams the notification subsystem.

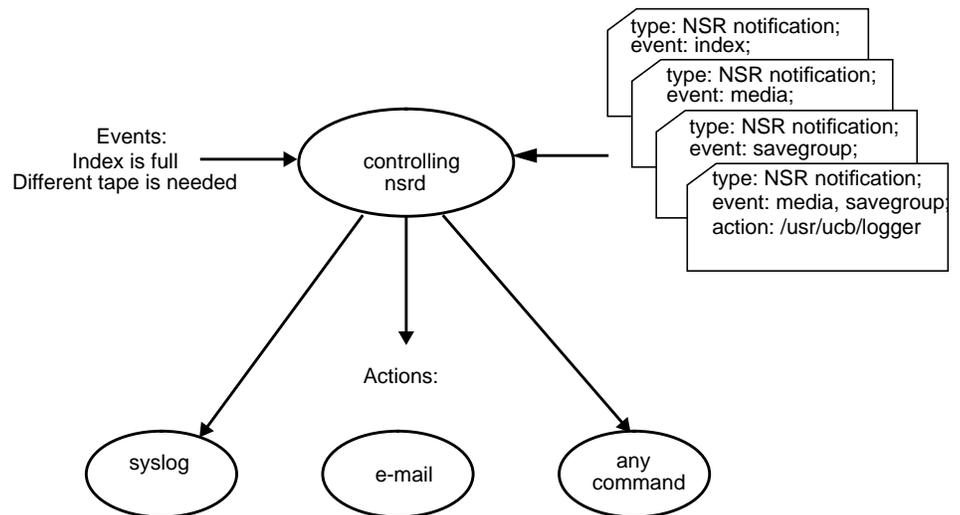


Figure C-9 Notification Subsystem

The relevant reference page is *nsr_notification(5)*, which describes the notification process.

Media Management

There is a working relationship between *nsrmmd*, the other daemons, the device resources, and the media manipulation commands. As shown, the *nsrmmd* daemon may be writing to one device while the operator brings a second device online. When *nsrmmd* needs a new backup volume, it polls the controlling *nsrd* for the backup volume and the device on which it is located. The *nsrmmd* daemon also records save set information to the *nsrmmdbd* daemon each time a backup volume filemark is written.

During recovery, *nsrmmd* queries *nsrmmdbd* to discover the backup volume and volume-file location of the desired data. The administrator uses the *mminfo* command to display information about the backup volume library.

Figure C-10 diagrams *nsrmmd* and its relation to other processes.

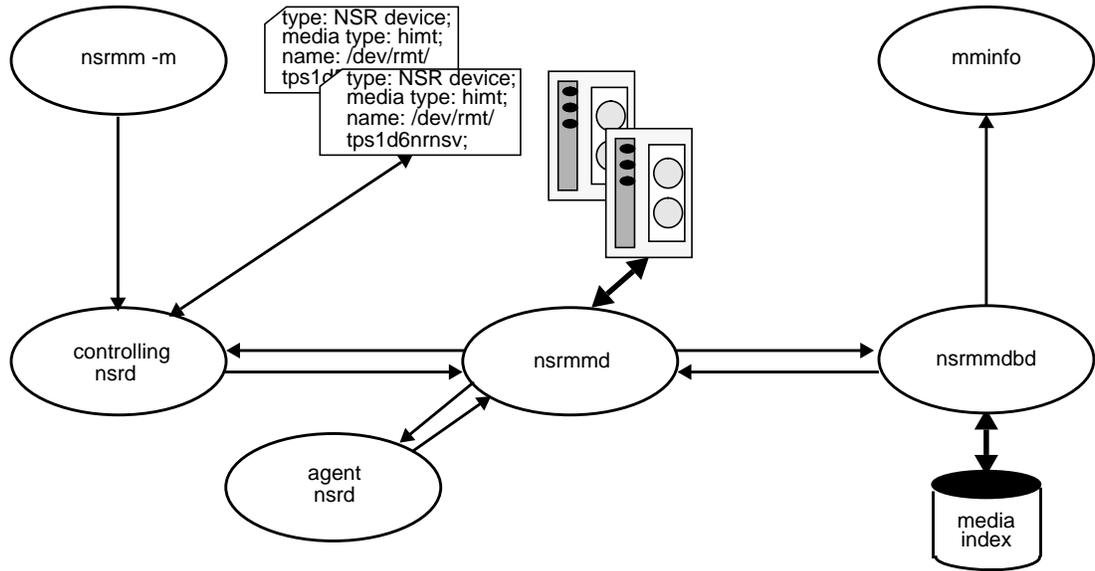


Figure C-10 nsrmm and Other Processes

Table C-6 lists the relevant reference pages.

Table C-6 Reference Pages for Media Management

Reference Page	Comments
<i>nsrmm</i> (1M)	Media manager mounts, unmounts, deletes backup volumes
<i>nsr_device</i> (5)	Describes and names backup devices
<i>nsrmm</i> d(1M)	Writes and reads data to and from the backup devices
<i>nsrmm</i> dbd(1M)	Manages the backup media library
<i>mminfo</i> (1M)	Displays information about the backup media library

Index Management

The online index is built during saves, and is queried during browsing and recovering. The *nsrls* command allows the administrator to gather information about the index sizes and record counts. Index entries are purged by the *nsrmm* command after it deletes a backup volume from the media index. The *nsrck* command is automatically invoked after failure (for example, an operating system crash) to guarantee the index database consistency before the NetWorker service is enabled.

Figure C-11 diagrams the file index and its relationship to processes.

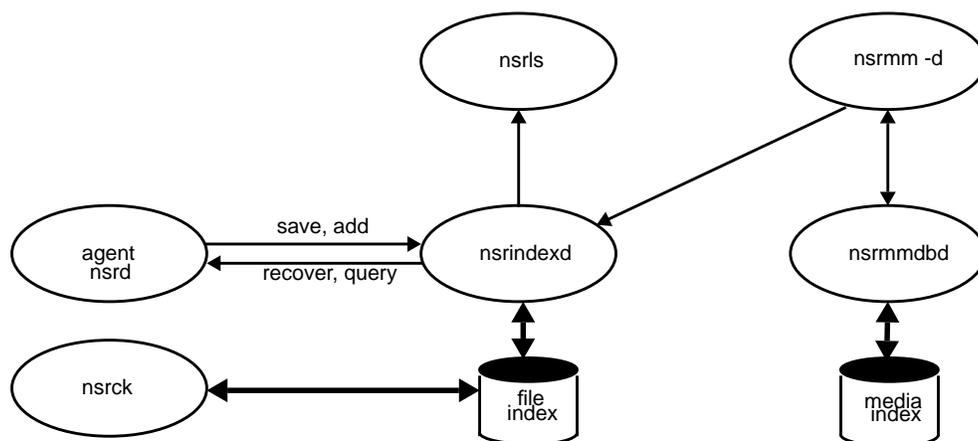


Figure C-11 The File Index and Processes

Table C-7 lists the relevant reference pages.

Table C-7 Reference Pages for Index Management

Reference Page	Comments
<i>nsrmm</i> (1M)	Media manager deletes backup volumes
<i>nsrindexd</i> (1M)	Manages the online index
<i>nsrls</i> (1M)	Displays index usage statistics

Table C-7 (continued) Reference Pages for Index Management

Reference Page	Comments
<i>nsrck</i> (1M)	Rebuilds indexes after a hard crash
<i>nsr_layout</i> (5)	Explains where all the NetWorker server's files are located

Disaster Recovery

The last act that *savegroup* performs each morning is backing up the NetWorker server's own index so that the index can be recovered without using itself. Besides saving the index, *saveindex* sends bootstrap recovery information to the printer. Should the server's index ever need to be recovered, this bootstrap information is sufficient to find the saved index on a NetWorker backup volume.

The *saveindex* daemon is invoked by *savegroup*; its job is to save index information vital to the server itself and to print recovery parameters to the line printer. Should the server's indexes ever need to be recovered, these parameters will be fed to the *recoverindex* command. Once the server's indexes are recovered, all other files may be recovered in the normal manner.

Figure C-12 shows the progression of index information from saving to recovering.

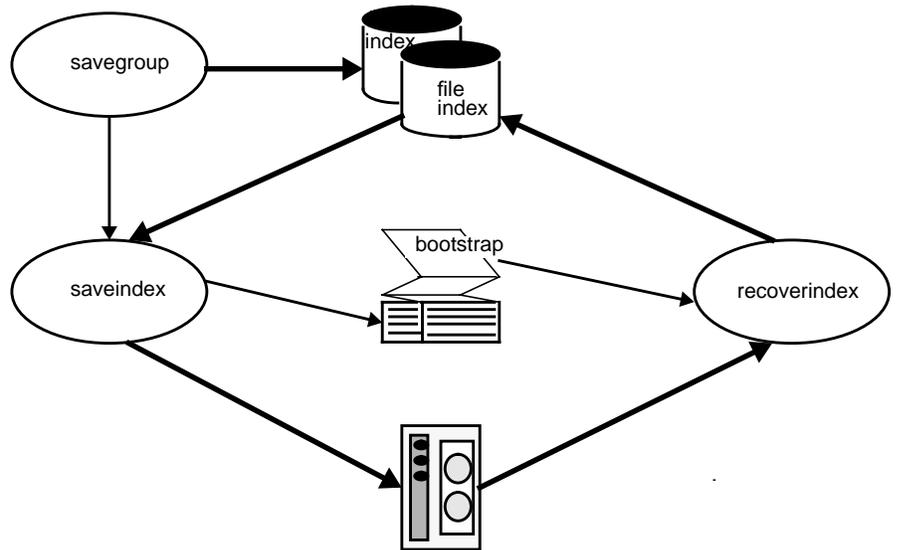


Figure C-12 Index Information From Saving to Recovering

Table C-8 lists the relevant reference pages.

Table C-8 Reference Pages for Disaster Recovery

Reference Page	Comments
<i>savegroup</i> (1M)	Typically invokes <i>saveindex</i>
<i>saveindex</i> (1M)	Also covers <i>recoverindex</i>
<i>scanner</i> (1M)	Used by <i>recoverindex</i> ; reads raw NetWorker backup volumes
<i>mm_data</i> (5)	Describes backup media format
<i>nsr_crash</i> (1M)	More general information on recovering after a disk crash

NetWorker Security

When a client attempts to save files or filesystems or recover them, NetWorker starts the *save*, *savefs*, or *recover* program, respectively. However, before the server accepts the command to start any of these programs, the client must establish a connection with the server. Before the server accepts the connection, it must verify that the user who initiated the save or recover request has the necessary permissions:

- The request for the connection must be made from a secure port on the user's machine. 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 workstation has permission to save or recover a client's files. For a user to have permission, one of the following criteria must be met:
 - the user's machine name must be 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

This access control is similar to that used by the *rsh* (remote shell) command, except that instead of using the *.rhosts* file, NetWorker uses the Remote access list in the Clients window.

Once a connection has been established, the client commands *save* and *recover* set their effective UID to the UID of the user who initiated the command so that all local filesystem and system call access is done as that user. Thus users cannot recover or back 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. Therefore, the administrator can 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 systems.

Note: Since any user named *operator* or any user belonging to a group named *operator* can access all data on all clients, care should be taken in assigning these names.

NetWorker's preconfigured selections allow the clients to browse and recover only their own files. To give other clients recover access to a client's

files, the administrator must explicitly add the access to the Remote access list in the Clients window, as explained earlier in this chapter.

You can further tighten access control for the client commands by turning off the set-uid bit. Doing this restricts client system use of the *save*, *savefs*, and *recover* commands to root. To allow access by root and *operator*, but not by other users, change the group ownership of these commands to *operator*, and set the mode bits to allow execution by owner and group.

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

The *nsrexecd* command runs on NetWorker client systems. This command provides a secure and restrictive way for NetWorker to start automatic backups on clients. The *nsrexecd* command allows you to restrict access to a select set of NetWorker servers. When you install a client, *nsrexecd* is started, and statements are added to the appropriate boot files to restart *nsrexecd* each time the client reboots.

Command Summary

This appendix contains a list of the most commonly used NetWorker commands entered at the system prompt, and lists of the NetWorker reference pages for NetWorker commands and resources.

Table D-1 summarizes the commands most often used from an ASCII terminal.

Table D-1 Commonly Used NetWorker Commands

Command	Purpose
<i>nsrmm -l volume_name</i>	Label a backup volume with the name <i>volume_name</i>
<i>nsrmm -m</i>	Mount the backup volume in the server drive
<i>nsrmm -u</i>	Unmount the backup volume in the server drive
<i>nsrmm -d -P volume_name</i>	Purge the backup volume named <i>volume_name</i>
<i>nsrmm -d volume_name</i>	Delete the backup volume named <i>volume_name</i>
<i>nsrmm -m -R</i>	Recycle the backup volume
<i>mminfo</i>	Display information about the backup volumes
<i>nsradmin</i>	Configure networker resources
<i>nsrck -F</i>	Compact the index size after purging or deleting backup volumes
<i>nsrexecd</i>	NetWorker client execution daemons
<i>nsrjb</i>	List the backup volumes in the jukebox
<i>nsrls</i>	Display information about the server's online indexes

Table D-1 (continued) Commonly Used NetWorker Commands

Command	Purpose
<i>nsrwatch</i>	Display the character-based NetWorker status monitor
<i>savegroup -p group_name</i>	See a preview of the backup of the group named <i>group_name</i>
<i>savegroup group_name</i>	Start the backup of the clients in the group <i>group_name</i>

Table D-2 lists the NetWorker maintenance commands.

Table D-2 NetWorker Maintenance Commands

Command	Reference Page	Description
<i>jbexercise</i>	<i>jbexercise</i> (8)	Jukebox diagnostic program
<i>nsrmm</i>	<i>nsrmm</i> (8)	NetWorker media interface command
<i>mminfo</i>	<i>mminfo</i> (8)	NetWorker media information reporting command
<i>mmlocate</i>	<i>mmlocate</i> (8)	Reports the backup volume location information contained in the media database
<i>nwadmin</i>	<i>nwadmin</i> (8)	Window-based display of the NetWorker server status
<i>nsr</i>	<i>nsr</i> (8)	Basic information on NetWorker
<i>nsradmin</i>	<i>nsradmin</i> (8)	NetWorker character-based program for system administration
<i>nsrck</i>	<i>nsrck</i> (8)	NetWorker check and repair program for the server indexes
<i>nsr_crash</i>	<i>nsr_crash</i> (8)	How to use NetWorker to recover from a disk crash
<i>nsrd</i>	<i>nsrd</i> (8)	NetWorker server daemon
<i>nsrexecd</i>	<i>nsrexecd</i> (8)	NetWorker client execution daemon
<i>nsrim</i>	<i>nsrim</i> (8)	NetWorker index management program, usually invoked from the <i>savegroup</i> command
<i>nsrindexasm</i>	<i>nsrindexasm</i> (8)	Module for saving and recovering NetWorker indexes

Table D-2 (continued) NetWorker Maintenance Commands

Command	Reference Page	Description
<i>nsrindexd</i>	<i>nsrindexd(8)</i>	NetWorker index daemon
<i>nsrjb</i>	<i>nsrjb(8)</i>	NetWorker optional jukebox software module
<i>nsrls</i>	<i>nsrls(8)</i>	Lists statistics of NetWorker index files
<i>nsrmmmd</i>	<i>nsrmmmd(8)</i>	NetWorker media management daemon
<i>nsrmmdbasm</i>	<i>nsrmmdbasm(8)</i>	Module for saving and recovering NetWorker media databases
<i>nsrmmdbd</i>	<i>nsrmmdbd(8)</i>	NetWorker media index daemon
<i>nsrwatch</i>	<i>nsrwatch(8)</i>	Displays NetWorker server status from an ASCII terminal
<i>recover</i>	<i>recover(8)</i>	Browses online indexes and recovers files from the NetWorker server
<i>recoverindex</i>	<i>recoverindex(8)</i>	Recovers online indexes of NetWorker server
<i>save</i>	<i>save(8)</i>	Saves files to the NetWorker server
<i>savefs</i>	<i>savefs(8)</i>	Saves filesystems to the NetWorker server (see also <i>savegroup</i>)
<i>savegroup</i>	<i>savegroup(8)</i>	Starts saving the files of a group of NetWorker clients
<i>saveindex</i>	<i>saveindex(8)</i>	Saves the online indexes of a NetWorker server
<i>scanner</i>	<i>scanner(8)</i>	Reads the contents of a volume to recover from NetWorker server crashes
<i>tapeexercise</i>	<i>tapeexercise(8)</i>	Exercises a tape drive in order to uncover problems
<i>uasm</i>	<i>uasm(8)</i>	NetWorker module for saving and recovering generic IRIX files

NetWorker Reference Pages

Table D-3 lists the reference pages for NetWorker client commands.

Table D-3 Reference Pages for NetWorker Client Commands

Reference Page	Description
<i>compressasm</i> (1M)	NetWorker program for compressing and decompressing files
<i>logasm</i> (1M)	NetWorker program for saving log files
<i>mailasm</i> (1M)	NetWorker program for saving and recovering mail files
<i>mminfo</i> (1M)	NetWorker media information reporting command
<i>mmlocate</i> (1M)	Reports the backup volume location information contained in the media database
<i>mmpool</i> (1M)	Reports on pool information stored in the media database
<i>networker</i> (1M)	A window-based display of the NetWorker server status
<i>nsr_shutdown</i> (1M)	Kills NetWorker processes on a NetWorker server
<i>nsrexecd</i> (1M)	NetWorker client execution daemon
<i>nsrmm</i> (1M)	NetWorker media interface command
<i>nsrwatch</i> (1M)	Displays the NetWorker server status from an ASCII terminal
<i>recover</i> (1M)	Command to browse the online indexes and recover files from the NetWorker server
<i>save</i> (1M)	Command to save files to the NetWorker server
<i>savefs</i> (1M)	Command to save filesystems to the NetWorker server
<i>swapasm</i> (1M)	NetWorker program for saving and recovering swap files
<i>uasm</i> (1M)	NetWorker program for saving and recovering generic IRIX files
<i>xlateasm</i> (1M)	NetWorker program for encoding and decoding files

Table D-4 lists other reference pages for NetWorker clients.

Table D-4 Other NetWorker Client Reference Pages

Reference Page	Description
<i>mm_data</i> (5)	Data format that <i>nsrmmd</i> uses
<i>nsr</i> (8)	Format of <i>.nsr</i> directive files
<i>nsr_client</i> (5)	NSR client resource type
<i>nsr_data</i> (5)	Data formats for <i>save</i> and <i>recover</i>
<i>nsr_device</i> (5)	NSR device resource type
<i>nsr_directive</i> (5)	NSR directive resource type
<i>nsr_getdate</i> (3)	Explains formats of dates that can be entered in date fields
<i>nsr_group</i> (5)	NSR group resource type
<i>nsr_jukebox</i> (5)	NSR jukebox resource type
<i>nsr_label</i> (5)	NSR label resource type
<i>nsr_layout</i> (5)	NetWorker file layout
<i>nsr_license</i> (5)	NSR license resource type
<i>nsr_notification</i> (5)	NSR notification resource type
<i>nsr_policy</i> (5)	NSR policy resource type
<i>nsr_pool</i> (5)	NSR pool resource type
<i>nsr_resource</i> (5)	NetWorker resource format
<i>nsr_schedule</i> (5)	NSR schedule resource type
<i>nsr_service</i> (5)	NSR resource type
<i>rap</i> (1M)	Command-line based resource administration program (RAP)
<i>rapcheck</i> (1M)	RAP resource directory consistency checking program
<i>rapd</i> (1M)	RAP server daemon
<i>rapxfer</i> (1M)	RAP resource directory manual area transfer
<i>resource</i> (5)	RAP resource file format

Table D-5 lists the reference pages for NetWorker server commands.

Table D-5 Reference Pages for NetWorker Server Commands

Reference Page	Description
<i>jbexercise</i> (1M)	Jukebox diagnostic command
<i>jbm_enabler</i> (1M)	Jukebox resource activator
<i>nsr</i> (8)	Guide to using NetWorker
<i>nsr_crash</i> (1M)	How to use NetWorker to recover from a disk crash
<i>nsradmin</i> (1M)	NetWorker character-based command for system administration
<i>nsrcap</i> (1M)	Update NetWorker information about NetLS licenses
<i>nsrck</i> (1M)	NetWorker check and repair command for the server indexes
<i>nsrd</i> (1M)	NetWorker server daemon
<i>nsrim</i> (1M)	NetWorker index management command, usually invoked from the <i>savegroup</i> command
<i>nsrindexasm</i> (1M)	Program for saving and recovering NetWorker indexes
<i>nsrindexd</i> (1M)	NetWorker index daemon
<i>nsrjb</i> (1M)	NetWorker optional jukebox program
<i>nsrls</i> (1M)	Lists the statistics of NetWorker index files
<i>nsrmmd</i> (1M)	NetWorker media management daemon
<i>nsrmmdbasm</i> (1M)	Program for saving and recovering NetWorker media databases
<i>nsrmmdbd</i> (1M)	NetWorker media index daemon
<i>recoverindex</i> (1M)	Command to recover the online indexes of a NetWorker server
<i>savegroup</i> (1M)	Command to start saving the files of a group of NetWorker clients
<i>saveindex</i> (1M)	Command to save the online indexes of a NetWorker server

Table D-5 (continued) Reference Pages for NetWorker Server Commands

Reference Page	Description
<i>scanner</i> (1M)	Command to read the contents of a volume to recover from NetWorker server crashes
<i>tapeexercise</i> (1M)	Command to exercise a tape drive in order to uncover problems

Glossary

1-9

Values for *override* in “NSR schedule” resource. Each number represents a backup level. Lower levels back up more files.

8 mm

Choice in the Devices window; represents eight-millimeter magnetic tape.

archive

The process by which NetWorker backs up directories or files that have not been accessed for a period of time to an archive volume, 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 that describes a NetWorker resource. It has a name and a list of values.

autochanger

A device that can move media among various components located in the device, including slots, media drives, media access ports, and transports; examples are jukeboxes (for optical disks) and tape libraries (for tapes).

bootstrap

In NetWorker, the bootstrap consists of three files: the server’s file index (*server.db*), the media index (*mmvolume*), the server configuration information (*nsr.res*) and, if the TurboPak option has been purchased, the jukebox configuration file (*jb.res*).

browse policy

The browse policy determines how long entries for your files remain in the online file index.

client

A system that accesses the NetWorker server to back up or recover files. Clients may be workstations, PCs, or file servers with gigabytes of data.

command line

The shell prompt, where you enter commands.

compressasm

The NetWorker directive used for compressing and decompressing files.

concurrent device support

Concurrent device support is used when you have more than one device attached to the NetWorker server and want to use all of the devices for backups simultaneously. Concurrent device support spreads the backups evenly across the devices.

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 that tracks every file or filesystem backed up.

fileserver

A system with disks that provides services to other systems on the network.

filesystem

1. A hierarchical data structure that holds and organizes files in directories.
2. The entire set of all IRIX files.
3. A method for 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 multiple backup volumes.

heterogenous

Heterogenous networks are networks with different platforms that interact across the network.

himt

A choice of media type in the Devices window; stands for half-inch magnetic tape.

holey

A directive used to efficiently back up files that do not have all of their data blocks allocated. This directive is not available on EFS systems.

incremental (i)

A backup level in which only files that have changed since the last backup are backed up.

interactive

A program or script that asks questions requiring input from the user.

interoperability

The ability of software and hardware on multiple systems from multiple vendors to communicate.

job pack

A tray or carousel that holds several backup volumes.

jukebox

A device that has the ability to move media among various components located in the device including slots, media drives, media access ports, and transports. Jukeboxes automate the media loading, labeling, and mounting functions during backups and recovers.

level 1-9

A backup level that backs up files that have changed since the last backup of any lower level.

mailasm

The 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 still tell if new mail arrived before 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, which tracks 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.

middle (m)

A command that requests a middle portion of a save set that spans multiple backup volumes.

mminfo

The NetWorker command that displays information about the backup volumes.

NetWorker

The network-based software product that backs up and recovers filesystems.

NetWorker client

A system that can access the backup and recover services from a NetWorker server.

NetWorker daemons

Daemons specific to the NetWorker environment.

NetWorker server

The system on a network running NetWorker software, containing online indexes, and providing backup and recover services to clients on a network.

NetWorker resources

Components of NetWorker software configuration information, described by a list of attributes and values.

notice

A response to a NetWorker event.

nsrhost

The logical hostname of the system that is the NetWorker server.

nsrjb

Command line interface to a jukebox device.

nsrls

The NetWorker command to display information about the online indexes.

nsrwatch

A command that gives you a character-based display of NetWorker status information.

online indexes

The databases located on the server that contain all the information pertaining to the client backups and backup volumes.

operator

The person, group, or account name that 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.

pre-configured

Existing selections or configurations for different NetWorker features.

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; stands for quarter-inch cartridge tape.

recover

The NetWorker command used to browse the server index and recover files from a backup volume to a client's disk.

retention policy

A policy that determines how long entries are retained in the media index and thus how long they are recoverable.

recycle

Relabel a backup volume with the same name, thus removing all entries associated with the backup volume from the server's index.

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.

scanner

The NetWorker command used to read a backup volume.

server

The system running the NetWorker software, containing the online index, and providing backup and recover services to clients on a network.

skip (s)

A backup level in which files are skipped and not backed up.

skip

The directive to skip files during a backup. Useful for skipping files that do not require a backup.

swapasm

The directive used to back up NFS swap files.

system administrator

The person normally responsible for installing, configuring, and maintaining NetWorker.

tail (t)

Represents the end of a save set that spans multiple backup volumes.

user

The person who can use 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 level backup, 63, 66-68
 - and full and incremental, 68
 - setting, 104
- 3480 media type in Devices window, 195
- 4mm media type in Devices window, 195
- 8 mm
 - 5GB media type in Devices window, 195
 - media type in Devices window, 195
- tape
 - capacity, 62
 - estimating time needed for backup, 125
 - number required, 69

A

- abort save set status, 148
- access control list, 223
- add* button, 42
- adding
 - administrator, 133
 - backup device, 196
 - client, 121-124
 - device to jukebox, 282-284
- adhesive label, 138, 168, 297
- administrator
 - adding or changing, 133-134
 - deleting, 134
 - list, 23
 - permissions, 130, 133-134

- Administrator program, 23-24
 - archive options in window, 262
 - icon, 24
 - starting, 24-25
 - window, 25-27, 198
 - and archiving, 262
 - Devices display, 27
 - menu bar, 26
 - Messages display, 27
 - Pending display, 27
 - pull-down menus, 28-40
 - server status display, 26
 - Sessions display, 27
 - speedbar, 26
- alias, configuring for client on server, 121
- appen mode, 146, 147
- Append* button, 42
- Application Specific Modules (ASMs), 112, 323, 351
- Apply* button, 42
- archive
 - and backup, 258
 - cloning, 267, 271
 - configuring client on server for, 118
 - creating client, 272
 - creating volume pool for, 265
 - details, 268
 - differences in NetWorker Administration
 - window, 262
 - directive, 267
 - grooming, *see* grooming
 - request
 - disabling, 268-269

- making, 273-274
 - rescheduling, 269
 - shortcut, 275
 - starting, 268-269
 - save set, 264
 - users, configuring client on server, 121
 - verifying, 271
 - volume, 259
- Archive Clone label template, 165
- Archive label template, 165
- Archive Request window, 266
- autochanger, 1, 57
 - license required, 13
 - policies, 88-91
 - see also* jukebox
- Autochanger Software Module, 13
- automatic cloning, 247
- Autostart, for group, 98
- B**
- backup, 4
 - 1-9, 63, 66-68
 - active device, maximum, 129, 132
 - and archive, 258
 - automatic network-wide, 98, 360-364
 - automatic start, 98
 - clone data for group, 98
 - data compression, 113
 - device
 - active, 132
 - adding and deleting, 193-196
 - required information, 196
 - calibrating, 304
 - concurrent, 131-133
 - deleting, 196
 - displaying, 193
 - enabled, 195
 - media type, 195
 - monitoring, 206-207
 - multiple, 193
 - no-rewind, 5, 195
 - nonfunctional, disabling, 195
 - speed, 311
 - supported, 5
 - with *compressasm* directive, 113
 - disabling network-wide, 98
 - distributing, 124
 - concurrent devices, 131, 132
 - entire system, 358-360
 - estimating time needed, 125
 - failed, 202
 - file changed during, error message, 321
 - file details, 55-56
 - first time, 50
 - full, 63
 - and incremental, 65
 - and incremental and level, 68
 - setting, 104
 - group, 95-101
 - changing time, 100
 - configuring, 95-101
 - client in, 119
 - creating, 96-100
 - permission, 99
 - deleting, 100-101
 - dividing filesystems, 124
 - in Pools window, 183
 - preconfigured, 95
 - restarting, 204-205
 - status, 200
 - stopping, 204
 - viewing details, 201-202
 - incremental, 63
 - and full, 65
 - and full and level, 68
 - setting, 104
 - interrupted, recover from, 205, 324
 - large filesystem, 124-126
 - level, 63-68

1-9, setting, 104
and disk crash recovery, 334-336
and full and incremental, 68
disk space required, 308
in Pools window, 184
of save set, 246
override, 103
manual, 129
media
 capacity, 153
 request notice, 219
monitoring, 55, 199-205
parallel, 4, 69, 314
previewing, 202-203
program
 icon, 27
 starting, 27
reducing time to complete, 69
retries, 98-99
 setting, 98
save session, 132
schedule, 61-80
 and save set cloning, 247
 changing, 104
 convenience versus security, 72
 creating, 104
 Default, 74-75
 for jukebox, 88-89
 Full Every Friday, 75
 Full on First Friday of Month, 76-77
 Full on First of Month, 77-78
 override, setting, 104
 preconfigured, 73-80, 102
 Quarterly, 78-80
 renaming, 104
 setting up, 101-108
 skip, setting, 103
 staggering, 71-72
 time considerations, 69
 viewing, 49-50
scheduled
 enabling, 52-54
 starting, 54-55
skipping, 63, 104
speeding up, 131
start time, setting, 98
starting immediately, 203
status, 200
theory of operations, 355-356
unattended, 4
unscheduled, 48
volume
 adding and recycling, 301-302
 and label template, 138
 and retention policies, 85
 appended not overwritten, 138
 capacity, 153
 cloned, recovering from, 158
 cloning, 156-164
 deleting, 155-156
 displaying information, 144, 145-148
 finding name, 143
 jukebox, mounting, 297-298
 label
 adhesive, 138, 168, 297
 for jukebox, 286-297
 name and physical location, 171
 labeling, 47-48, 137-142
 bar codes, 290-296
 location, setting, 150-151
 managing automatically, 144-148
 managing manually, 148-156
 marked "full" when not, 153
 media format, 351
 mode, 146
 changing, 151-153
 Full, 152
 mounting, 47-48, 142-143
 number required, 62
 overriding retention policy, 154
 physical location, 171

- pool, 178-192
 - changing, 190
 - creating, 179-180, 187-191
 - deleting, 179, 182
 - preconfigured, 185-186
 - premounting, 143
 - Read Only mode, 152
 - recycling, 154
 - recycling manually, 289
 - relabeling, 142
 - removing, 155-156
 - requested for recovery, 143
 - unmounting, 142-143
 - writable, 143
- bar code, 281, 290-296
- label, identical, 290, 295
 - matching volume label, 293
 - on existing volumes, 291
 - physical label, 291, 292
- binding to server errors, 319
- bootstrap
- file, 329
 - information
 - filing, 330-331
 - not printed, 322
 - save set ID
 - and recovering indexes, 340-342
 - finding, 344-345
- brows save set status, 147
- browse and retention policies, 80-87
- configuring client on server for, 118
 - creating, 109-112
 - default, 94
 - jukebox, 89
 - overriding for volume, 154
 - preconfigured, 86-87
 - Decade, 87
 - Month, 86-87
 - Quarter, 87
 - Week, 86
 - Year, 87
 - relative size, 86
- browsing index, 81
- buttons, generic, 42-43
- ## C
- calendar in Schedules window, 103
- calibrating device, 304
- Cancel* button, 42
- change* button, 42
- Change Server* speedbar button, 134
- changing
- administrator, 133-134
 - NetWorker server, 134-135
 - time for backup group, 100
 - volume mode, 151-153
- cleaning jukebox devices, 304
- client, 3
- access, restricting, 371
 - adding too many, 122
 - backing up files, 121
 - binding to server, 17
 - browsing files, 121
 - cloning several automatically, 250-256
 - configuring
 - archive services, 263-265
 - on network with NIS, 17-18
 - viewing window, 50-52
 - configuring on server, 115-126
 - alias, 121
 - archive, 118
 - archive users, 121
 - backup group, 119
 - browse and retention policies, 118
 - directives, 119
 - password, 120

- recover access, 123-124
- remote access, 119
- remote user, 120
- removing, 126
- save set, 119
- schedule, 118
- selecting files for backup, 122-123
- CPU speed, 314
- creating an archive, 272
- data generation rate, 313-314
- in Pools window, 183
- license required, 12
- manual backup, 129
- maximum parallelism, 129
- NetWorker server as client of itself, 93, 117
- parallelism, 130-131
- performance, 313-314
- preconfigured settings, 94-95
- recognizing the server, 17
- recover access, 123-124, 318
- reference pages, 377
- removing, 126
- renaming, 317-318
- retries, 98-99
 - setting, 98
- save session, 132
- selecting files for backup, 122-123
- server as client of itself, 93, 117
- software for IRIX 4.0.x, 8
- software installation, 19-21
- subsystems for, 20
- too many, 122
- transfer rate, 306-307
- Client Connection, 12
- ClientPak Software Module, 13
- Clients menu, 29-30
- Clients window, 50-52, 117
 - lower part, 120
- client/server model, 3, 319
- clone
 - backup volume, 156-164
 - data for group, 98
 - pool
 - for group, 98
 - for save set, 238
 - preconfigured, 238
 - selecting, 243
 - volume
 - changing mode, 151-153
- Clone Volume window, 160
- cloned data
 - recovering, 158
 - removing entries from index, 155
 - retention policies, 236
- cloning
 - archive, 267
 - save set, 259
 - volume, 271
 - data
 - automatically, 247
 - minimizing network traffic, 160, 237
 - see also* save set, cloning
 - selecting save sets, 241
- commands, 373-379
- compressasm* directive, 69-71
 - and backup device, 113
 - reference page, 376
- compressing
 - backup data, 113
 - files, 69, 71
- concurrent device, 12, 57, 131-133, 193
 - distributing backups, 132
 - maximum, 129, 132
 - reducing backup time with, 70
- configuring
 - backup group, 95-101
 - permission, 99
 - client on server, 115-126
 - jukebox, 278-285
 - NetWorker server, 93-135

- conventions, xxv
- CPU speed
 - client, 314
 - server, 312
- crash recovery, *see* disk crash, recovering from
- Create button, 42
- Create Multiple button, 42
- customer support, xxvii
- Customize menu, 31-32

- D**

- daemons
 - error messages, 319
 - NetWorker, starting, 316
 - nsrd*, 374
 - nsrexecd*, 325, 374
 - nsrmmdbd*, 375
 - on server, 351-353
 - restarting, 319
 - savefs*, 358-359
 - saveindex*, and disaster recovery, 361
 - starting, 11, 325
- data
 - archiving, 259-260
 - cloning, *see* save set, cloning
 - generation rate of client, 313
 - integrity, and cloned save set, 237
 - recovering large amounts, 224
 - recovering unbrowsable, 224
 - retrieving, 260-261
 - scanning, 224
 - verification, 260, 267, 271
- Decade policy, 87
- Default
 - backup group, 94, 95
 - backup schedule, 74-75
 - Clone label template, 166
 - directive, 94
 - label template, 166
- Delete button, 42
- deleting
 - administrator, 134
 - backup device, 196
 - backup volume, 155-156
 - client, 126
 - volume pool, 179, 182
- Details button, 38, 55
- details, displaying, 38-39
- device
 - active backup, maximum, 129
 - backup, media type, 195
 - concurrent, 57, 131-133
 - maximum, 129
 - jukebox, *see* jukebox, autochanger
 - number of save sessions, 129
 - reserved for recovery, 129
 - tape, no-rewind, 5
- devices display in NetWorker Administrator window, 27
- Devices window, 194
- df* command and disk recovery, 331
- Digital Linear Tape
 - capacity, 308
 - data rate, 306
 - see also* DLT2700
- directive, 112-115, 376
 - and archiving, 267
 - compressasm*, 69-71
 - and backup device, 113
 - reference page, 376
 - configuring client on server for, 119
 - creating, 113-115
 - default, 94
 - DOS standard, 113
 - logasm* reference page, 376
 - mailasm* reference page, 376

NetWare standard, 113
nsrindexasm, 374
 reference page, 378
nsrmmdbasm, 375
preconfigured, 94, 113
skip, 70
swapasm reference page, 376
uasm, 324, 375
 reference page, 376
Unix standard, 113
Unix with compression, 113
xlateasm reference page, 376

Directives window, 114

directory, recovering save set, 230

disk

- crash
 - preparing for, 329-333
 - recovering from, 327-347
 - jukebox, 346-347
 - primary, 339-340
 - secondary, 338-339
 - shortcut, 328
 - theory of operations, 368-369
- information, obtaining and filing, 331-333
- multiple, 314
- primary, recovering, 334, 339-340
- recover partition, example, 234-236
- recovering, 224
- reducing amount used, 223
- secondary, recovering, 338-339
- space
 - determining amount needed, 308-309
 - event notification, 218
 - for Decade policy, 87
 - for indexes, 9
 - jukebox, 300
 - monitoring usage, 208-212
 - reclaiming, 81, 84, 213-215
 - reducing amount needed, 81, 85
 - required for backup level, 308

displaying backup volume information, 144
dlt media type in Devices window, 195

DLT2700

- data rate, 308
- stacker, non-Silicon Graphics, 5
- time requirement, 69

DNS

- and configuring, 17, 19
- and nsrhost alias, 19

E

enabling

- NetWorker products, 12-16
- scheduled backup, 52-54, 98

error message

- binding to server, 319
- cannot print bootstrap information, 322
- file changed during backup, 321
- saving remote filesystem, 320-321

/etc/lvtab file and disk recovery, 331

event notification, 215-219

- customizing, 217
- index size, 218
- log default, 218
- preconfigured, 218-219
- registration, 218
- Savegroup completion messages, 319-322
- Savegroup completion warning messages, 322-323
- theory of operations, 364-365

event processing, 364-365

EXABYTE

- 8200, data rate, 306
- 8500
 - data rate, 306
 - time requirement, 69
- EXB-10i stacker
 - data rate, 308
 - example policies, 89-90

excluding data from recovery, 230
Exit command, 29
exiting NetWorker, 29

F

failed backup, 202
file
 and retention policies, 85
 changed during backup, error message, 321
 compressing, 69
 index, 2, 80
 and browse and retention policies, 86
 contents, 80
 disk space, 80
 removing entries, 81
 name length, 315
 open, 205
 recover
 entire save sets, 230
 from interrupted backup, 205, 324
 individual, example, 232-234
 recovered
 overwriting, 232
 renaming, 232
 removed from index, recovering, 84
 unbrowsable, recovering, 224
 version
 and disk space, 82
File menu, 29
filename conflict, 232
filesystem
 backing up specific, 183, 192
 dividing for backup, 124-126
 in Pools window, 183
 saving remote, error message, 320-321
 traversing, 313
font conventions, xxv
full backup, 63

 and incremental and level, 68
 and incremental backup, 65
 setting, 104
Full Every Friday backup schedule, 75
Full label template, 166
full mode, 146
Full on 1st of Month backup schedule, 77-78
Full on First Friday of Month backup schedule, 76-77

G

generic buttons, 42-43
grooming, 260, 268
group
 default, 94
 preconfigured, 94
 see also backup group
Group Control Details window, 201
Group Control window, 55, 200
Groups window, 97

H

hardware configuration, 305-310
 client, 306-307
 jukebox, 307-309
 network, 307
 server, 307
 tape drives, 306
Help button, 34
Help menu, 34
Hierarchical Storage Management, 258
himt media type in Devices window, 195
HSM, 258

I

in-pro save set status, 148
incremental backup, 63
 and full and level, 68
 and full backup, 65
 setting, 104
index, 85
 browsing, 81
 cloned save set, 157, 237
 disk space, 80, 208-210
 requirements, 9
 entry
 recovering, 223
 removing, 237
 file, 2, 80
 and browse and retention policies, 86
 contents, 80
 percent used, 209
 recovering, 219-220
 inserting entries, 81
 management, theory of operations, 367-368
 media, 2, 80
 and retention policies, 85
 moving, 324
 online, 80
 managing, *see* browse and retention policies,
 creating
 store, 184
 theory of operations, 367-368
 reducing, 213-215, 223
 removing
 entry, 81, 157
 oldest cycle, 212-213
 size, 80
 notification, 218
Indexes window, 209
InSight manuals, xxvii, 20
installation
 on client, 19-21

 on server, 9-19
 requirements, 9
 subsystems
 on client, 20
 on server, 10
Instances button, 211
Instances window, 211
inventorying jukebox, 302-303
IRIX version required, 2, 20

J

jbexercise command, 278, 374
 reference page, 378
jbm_enabler command, 15
 reference page, 378
jukebox, 57, 277-304
 accidentally relabeling media, 289
 adding device, 282-284
 and scheduling, 70
 and volume pool, 304
 backup volume, mounting after errors, 300
 bar code, *see* bar code
 calibrating devices, 304
 capacity, 88-90
 changing available slots, 285
 checking space, 300
 choosing, 90-91
 cleaning devices, 304
 configuring, 278-285
 control port, 281
 determining capacity, 206-207, 300
 device pathname, 281
 order, 284
 disaster recovery, 346-347
 installing and activating license, 15-16
 inventory, 302-303
 labeling status message, 296-297
 license required, 13

- loading mechanism, 281
- model, 280
- notification, 299-300
- operating, 299-304
- policies, 88-91
- recycling volume manually, 289
- request, 219
- slot, 281
 - controlling, 285
 - labeling, 287-289
 - number, 16
 - specifying, 16
- space, 300
- testing connection, 278
- transfer rate, 307-309
- unavailable tape, allowing for, 90, 91, 309
- volume pool example, 192
- volume, mounting, 297-298

Jukebox Mounting window, 207

L

- Label* speedbar button, 139
- label template, 164-177, 286-290
 - and volume pool, 138
 - Archive, 165
 - Archive Clone, 165
 - choosing, 167-168
 - components, 177
 - consistent with volume label, 287
 - creating, 169-171, 172-174, 175-177
 - Default, 166
 - Default Clone, 166
 - Full, 166
 - in Pools window, 183
 - matching pool name, 171
 - naming, 164
 - NonFull, 166
 - Offsite, 166
 - preconfigured, 165-167

- Two Sided, 166
 - example, 176-177
- Label Templates window, 169
- label, adhesive, 138, 168, 297
- labeling
 - accidentally relabeling, 286
 - and mounting volume, 141-142
 - and not mounting volume, 141
 - consistently with templates, 287
 - matching bar code and volume names, 293
 - relabeling, 141
 - with bar codes, 290-297
 - identical, 290, 295
- license
 - displaying, 58
 - installing and activating, 13-16
 - NetLS, 12
 - permanent, 13
 - problems, 317
 - required for NetWorker products, 12
 - temporary, 7, 13
- list* command, 340
- log default notification, 218
- logasm* directive, reference page, 376
- ls_nodelock* command, 16

M

- mailasm* directive, reference page, 376
- man recyc mode, 146
- manual save, 129
- matching label names, 287
- media
 - accidentally relabeling, 289
 - device, maximum number for concurrent use, 129
 - format, 351
 - index, 2, 80
 - and cloned save set, 236

- and volume label, 137
- changing volume mode, 151-153
- disk space, 80
- recycling volume, 154
- removing volume entries, 155
- setting volume location, 150
- viewing, 144
- management, theory of operations, 365-366
- manager, 2
- relabeling, 141
- request notice, 219
- type of backup device, 195
- verification, 260
- Media menu, 30-31
- memory and performance, 313
- Messages display in NetWorker Administrator window, 27
- mminfo* command, 219, 373, 374
 - reference page, 376
- mmlocate* command, 374
 - reference page, 376
- mmpool* command, reference page, 376
- monitoring
 - backup, 199-205
 - device, 206-207
- Month policy, 86-87
- Mount* speedbar button, 142
- mounting volume, 141-142
- multiple configurations, 42

N

- NetLS, 12
 - installing and activating license, 12-16
 - license, 12
 - information sheet, 13
 - problems, 317
 - subsystem required, 14

- network
 - data transfer rate, 312
 - speed, 312
 - traffic, and cloning volumes, 237
- NetWorker
 - Administrator program, *see* Administrator program
 - AutoChanger Software Module, 57
 - Autochanger Software Module, 88
 - commands, 373-379
 - components, diagrammed, 350
 - enabling, 12-16
 - installation, 7-21
 - media format, 351
 - operating, 197-220
 - other documentation, xxvi
 - performance, 305-314
 - products, 12
 - programs, 23
 - starting, 23-28
 - reference pages, 352-354, 356, 358, 359, 364, 365, 366, 367, 369, 376-379
 - removing, 21
 - resources, 353-354
 - server, client of itself, 93, 117
 - software distribution CD, 7
 - part number, 9
 - version, xxiii, 2
 - displaying, 34
 - networker* command reference page, 376
 - networker start* command, 316, 318
 - networker stop* command, 318
- NIS
 - and configuring, 17-18
 - and nsrhost alias, 19
- no-rewind tape device, 5
- NonFull label template, 166
- normal save set status, 149
- notice, *see* notification

- notification, 215-219
 - jukebox, 299-300
 - preconfigured, 218-219
 - see also* event notification
 - Notifications window, 217
 - nsr* command, 374
 - reference page, 378
 - /nsr* directory for indexes, 316
 - recovering, 340-345
 - nsr_crash* command, 374
 - reference page, 378
 - nsr_getdate(5)* reference page, 377
 - nsr_shutdown* command, 21
 - nsradmin* command, 373, 374
 - reference page, 378
 - nsrarchive* command, 260
 - nsrcap* command, 16
 - reference page, 378
 - nsrck* command, 373, 374
 - reference page, 378
 - nsrd* command, 319
 - reference page, 376, 378
 - nsrd* daemon, 374
 - nsrexecd* command, 203, 320, 325, 373
 - and security, 371
 - reference page, 376
 - nsrexecd* daemon, 374
 - nsrhost* alias, 17-19
 - nsrim* command, 374
 - reference page, 378
 - nsrindexasm* directive, 374
 - reference page, 378
 - nsrindexd* command, 375
 - reference page, 378
 - nsrjb* command, 207, 301, 346, 373, 375
 - reference page, 378
 - nsrls* command, 373, 375
 - reference page, 378
 - nsrmm* command, 373, 374
 - reference page, 376
 - nsrmmmd* command, 357, 365-366, 375
 - reference page, 378
 - nsrmmdbasm* directive, 375
 - reference page, 378
 - nsrmmdbd* command, 318, 365
 - reference page, 378
 - nsrmmdbd* daemon, 375
 - /nsr/res* directory, recovering from crash, 343-344
 - nsrretrieve* command, 260
 - and *nwrecover*, 264
 - nsrwatch* command, 339, 340, 374, 375
 - reference page, 376
 - NW Backup
 - icon, 27
 - starting, 27
 - NW Recover
 - icon, 28
 - starting, 28
 - nwadmin* command, 24, 197, 374
 - nwbackup* command, 27
 - nwrecover* command, 28
- O**
- Offsite label template, 166
 - Ok* button, 42
 - operator group
 - and permissions, 318
 - and security, 370
 - optical media type in Devices window, 195
 - Options menu, 32-34
 - override
 - backup level, 103
 - setting, 104
 - Overwrite* button, 42

overwriting recovered file, 232

P

parallel backup, 4, 69, 314

parallelism, 129, 130-131

- maximum, 129

- turning off, 131

password, configuring for client on server, 120

pathname

- of jukebox device, 281

- of tape device, 5, 195

- order for jukeboxes, 284

Pending display in NetWorker Administrator window, 27

performance, 305-314

- client, 313-314

- server, 311-313

Period buttons in Schedules window, 103

permissions, 133-134, 318

- user and user group, 134

policies

- autochanger, 88-91

- browse and retention, 80-87

 - preconfigured, 86-87

- changing, 112

- creating, 111-112

- deleting, 111

- jukebox, 88-91

- naming, 111

- preconfigured, 110

Policies window, 110

polling interval, changing, 33

pool, *see* volume pool

Pools window, 168, 181

portmapper registration service, 319

preconfigured settings

- clone pools, 238

- deleting, 112

- directives, 113

- event notifications, 218-219

- label template, 165-167

- schedules, 102

- volume pool, 185-186

Preview Run Status, 200

primary disk, *see* disk, primary

Print button, 43

Print command, 29, 40-41

printing window contents, 40-41

products, NetWorker, 12

prtvtoc command, 332

- and disk recovery, 331

Q

qic media type in Devices window, 195

Quarter policy, 87

Quarterly backup schedule, 78-80

R

.R extension, 232

reclaiming

- disk space, 84, 213-215

- index space, 209

recov save set status, 147

recover, 5

- access

 - default, 94

 - issues, 318

 - preconfigured, 94

- disk, 224

- disk partition example, 234-236

- file removed from index, 84

- files, from index, 219-220

- from disk crash, 327-347
 - shortcut, 328
 - from interrupted backup, 205, 324
 - from reserved device, 129
 - index, 340-344
 - jukebox, 346-347
 - large amounts of data, 224
 - policy, 62
 - preparation for, 329-333
 - primary disk, 339-340
 - program icon, 28
 - secondary disk, 338-339
 - theory of operations, 357-358
 - to new server after disk crash, 345-346
 - unbrowsable data, 224
 - using after disk crash, 333-337
- recover* command, 340, 375, 376
 - reference page, 376
- recoverindex* command, 340, 341, 343, 375
 - reference page, 378
- recyc
 - save set mode, 146
 - save set status, 148
- recycling volume manually, 140, 154
- reducing used disk space, 223
- reference page subsystems
 - client, 20
 - server, 10
- reference pages, xxvii, 352-354, 356, 358, 359, 364, 365, 366, 367, 369, 376-379
- registration notification, 218
- Registration window, 58
- relabeling
 - media, accidentally, 289
 - volumes, 141
- relocating recovered data, 232
- remote access, configuring for client on server, 119
- remote user, configuring for client on server, 120
- Remove oldest cycle command, 81
- removing
 - administrator, 134
 - backup device, 196
 - client, 126
 - client files and directories after archiving, 268
 - index entry, 157, 237
 - NetWorker, 21
 - oldest cycle, 210, 212-213
- renaming
 - client, 317-318
 - recovered file, 232
- rescheduling an archive request, 275
- Reset* button, 42
- resource, 353-354
 - type, 377
- restarting a backup group, 204
- retention policies, 85
 - and backup volume, 85
 - and indexes, 85
 - and save set, 85
 - for cloned save set, 236
 - for cloned volumes, 236
 - see also* browse and retention policies
- root
 - and permissions, 318
 - and security, 370
- rsh* command, 320

S

- Save* button, 43
- Save* command, 29, 41
- save* command, 355, 375
 - reference page, 376
- save session, 132
 - number per device, 129
- save set, 2, 73
 - and resource usage, 210-212

and retention policies, 85
archive, 264
 versus backup, 264
backup level, 246
browsing, 240-247
changing status, 149
cloned, index entries, 157
cloning
 automatic, 247-256
 several clients, 250-256
 index entry, 237
 integrity, 237
 manually, 248-249
 requirements, 238
 retention policies, 236
 selecting clone pool, 243
 status, 243-244
 timing, 237
 tracking, 236
 viewing details, 242-243
 volume, 236
 size, 236
configuring for client on server, 119
date created, 241
default, 94
details, 228
ID, 243, 329
 for bootstrap
 and recovering indexes, 340-342
 finding, 344-345
in Pools window, 183
level, and backup recovery, 334-336
listing in file index, 210
preconfigured, 94
recover
 entire, 230
 options, 231-232
 status, 229-231
 using after disk crash, 333-337
removing oldest cycle, 212-213
 status, 147-148
 normal, 149
 suspect, 149
 version, 226
Save Set Clone window, 240, 245
Save Set menu, 31
savefs command, 375
 reference page, 376
savefs daemon, 358-359
savegroup -p command, 122, 203
savegroup command, 360, 374, 375
 reference page, 378
Savegroup completion
 messages, 319-322
 notification, 218
 warning messages, 322-323
savegroup daemon, 368
saveindex command, 330, 360, 375
 reference page, 378
saveindex daemon, 368
 and disaster recovery, 361
saving
 manually, 129
 window contents, 41-42
scann save set status, 148
scanner command, 147, 156, 158, 224, 341, 344, 375
 reference page, 379
scanning data, 224
schedule
 configuring client on server, 118
 see also backup schedule, 118
Schedules window, 102
Search command, 40
secondary disk, *see* disk, secondary
security
 client recover access, 123
 secure devices, 191
 theory of operations, 370-371

- versus convenience, 72
- server, 3
 - administrator permissions, 130, 133-134
 - backed up as client, 93, 117
 - changing, 46, 134-135
 - checking status, 197-199
 - client of itself, 93, 117
 - configuring, 93-135
 - daemons, 351-353
 - default backup as client, 50
 - determining, 324
 - displaying current, 26
 - errors, 319
 - hardware configuration, *see* hardware configuration
 - jukebox operations, 299-304
 - license, 12
 - installation, 12-16
 - list, 135
 - parallelism, 130-131
 - performance, 311-313
 - permissions, 130, 133-134
 - preconfigured settings, 94-95
 - settings, 127-135
 - software installation, 9-19
 - requirements, 9
 - software version on, 128
 - status display in NetWorker Administrator window, 26
 - subsystems for, 10
- Server menu, 30
- Server Setup command, 127
- Server window, 128
- sessions display in NetWorker Administrator window, 27
- skip
 - backup level, 63
 - backup, setting, 103
 - directive, 70
- slot, changing for jukebox, 285

- sort data, 178, 182
- speed
 - CPU
 - client, 314
 - server, 312
 - network, 312
- speedbar, 26
 - enabling, disabling, 33
- Speedbar* button, 33
- ssid, 329
- staggering backup schedules, 71-72
- starting
 - Administrator program, 24
 - backup program, 27
 - backups immediately, 203
 - programs, 23-28
 - recover program, 28
- status
 - backup group, 200
 - jukebox labeling, 296-297
- stopping a backup group, 204
- susp save set status, 148, 149
- swapsm* directive, reference page, 376

T

- Tabular* button, 37
- Tabular command
 - for printing window contents, 40
 - for saving window contents, 41
- tabular view, 35-37
- tape
 - device
 - no-rewind, 5, 195
 - performance, 311
 - drive
 - supported, 306
 - transfer rates, 306

Tape mount request, 299
tapeexercise command, 375
 reference page, 379
technical support, xxvii
temporary license, 7
theory of operations, 349-371
Tone button, 34
tone, enabling, disabling, 34
transfer rate, 62
 client, 306
 jukebox, 307-309
 network, 307
 server, 307
 tape drives, 306
troubleshooting, 315-325
TurboPak option, 12
Two Sided label template, 166

U

uasm directive, 324, 375
 reference page, 376
Unmount speedbar button, 142
Update Server List button, 135
user ID, special, 318, 355
/usr/etc/resolv.conf file, 19

V

/var/adm/SYSLOG file, 11, 199, 218, 316, 364
version
 IRIX and NetWorker, 2
 of NetWorker, xxiii
 software on NetWorker server, 128
Versions menu option and media information, 357
vhs media type in Devices window, 195

View menu, 35-40
volume
 capacity, 153
 clone pool for save set, 238
 cloned, changing mode, 151-153
 label, consistent with label template, 287
 location, setting, 150-151
 mode
 changing, 151-153
 Full, 152
 name, 140
 pool, 2
 and backup volume, 138
 and jukebox, 304
 and label template, 138
 choosing, 140
 creating for archiving, 265
 Default, 138
 deleting, 179, 182
 specific clients, 183
 Read Only mode, 152
 recycling, 154
 relabeling, 141
 see also backup volume
Volume Management window, 145
Volumes speedbar button, 144

W

Week policy, 86
window contents
 printing, 40-41
 saving, 41-42

X

X Window System, 4
XDR format, 351

xlateasm directive, reference page, 376
XView error messages, 323-324

Y

Year policy, 87