



Data General Corporation, Westboro, Massachusetts 01580

Customer Documentation

Legato NetWorker Administrator's Guide

069-100495-01

A V I I O N[®]
P R O D U C T L I N E

Legato NetWorker

Administrator's Guide

069-100495-01

For the latest enhancements, cautions, documentation changes, and other information on this product, please see the Release Notice (085-series) supplied with the software.

All Rights Reserved
Printed in the United States of America
Rev. 01, January 1994
Ordering No. 069-100495-01

NOTICE

DATA GENERAL CORPORATION (DGC) HAS PREPARED AND/OR HAS DISTRIBUTED THIS DOCUMENT FOR USE BY DGC PERSONNEL, LICENSEES, CUSTOMERS, AND PROSPECTIVE CUSTOMERS. THE INFORMATION CONTAINED HEREIN IS THE PROPERTY OF THE COPYRIGHT HOLDER(S); AND THE CONTENTS OF THIS MANUAL SHALL NOT BE REPRODUCED IN WHOLE OR IN PART NOR USED OTHER THAN AS ALLOWED IN THE APPLICABLE LICENSE AGREEMENT.

The copyright holders reserve the right to make changes in specifications and other information contained in this document without prior notice, and the reader should in all cases determine whether any such changes have been made.

THE TERMS AND CONDITIONS GOVERNING THE SALE OF DGC HARDWARE PRODUCTS AND THE LICENSING OF DGC SOFTWARE CONSIST SOLELY OF THOSE SET FORTH IN THE WRITTEN CONTRACTS BETWEEN DGC AND ITS CUSTOMERS, AND THE TERMS AND CONDITIONS GOVERNING THE LICENSING OF THIRD PARTY SOFTWARE CONSIST SOLELY OF THOSE SET FORTH IN THE APPLICABLE LICENSE AGREEMENT. NO REPRESENTATION OR OTHER AFFIRMATION OF FACT CONTAINED IN THIS DOCUMENT INCLUDING BUT NOT LIMITED TO STATEMENTS REGARDING CAPACITY, RESPONSE-TIME PERFORMANCE, SUITABILITY FOR USE OR PERFORMANCE OF PRODUCTS DESCRIBED HEREIN SHALL BE DEEMED TO BE A WARRANTY BY DGC FOR ANY PURPOSE, OR GIVE RISE TO ANY LIABILITY OF DGC WHATSOEVER.

IN NO EVENT SHALL DGC BE LIABLE FOR ANY INCIDENTAL, INDIRECT, SPECIAL OR CONSEQUENTIAL DAMAGES WHATSOEVER (INCLUDING BUT NOT LIMITED TO LOST PROFITS) ARISING OUT OF OR RELATED TO THIS DOCUMENT OR THE INFORMATION CONTAINED IN IT, EVEN IF DGC HAS BEEN ADVISED, KNEW OR SHOULD HAVE KNOWN OF THE POSSIBILITY OF SUCH DAMAGES.

All software is made available solely pursuant to the terms and conditions of the applicable license agreement which governs its use.

AViiON is a U.S. registered trademark of Data General Corporation.

DG/UX is a trademark of Data General Corporation.

CLARiiON is a trademark of Data General Corporation.

Restricted Rights Legend: Use, duplication, or disclosure by the U.S. Government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at [DFARS] 252.227-7013 (October 1988).

DATA GENERAL CORPORATION
4400 Computer Drive
Westboro, MA 01580

Legato NetWorker Administrator's Guide

069-100495-01

Effective with Legato NetWorker, Revision 4.0

Copyright 1993, Legato Systems, Inc. All Rights Reserved. Legato NetWorker is a trademark of Legato Systems, Inc. All other trademarks are the property of their respective owners.

LICENSE AGREEMENT

THIS PRODUCT CONTAINS CERTAIN COMPUTER PROGRAMS AND OTHER PROPRIETARY MATERIAL, THE USE OF WHICH IS SUBJECT TO THIS LICENSE AGREEMENT. IF YOU DO NOT AGREE WITH ALL THE TERMS, YOU MUST RETURN THIS PRODUCT, ALL MANUALS AND DOCUMENTATION, AND PROOF OF PAYMENT, TO THE PLACE YOU OBTAINED THEM FOR A FULL REFUND WITHIN 30 DAYS OF FIRST ACQUIRING THIS PRODUCT. YOUR WRITTEN APPROVAL IS NOT A PREREQUISITE TO THE VALIDITY OR ENFORCEABILITY OF THIS AGREEMENT AND NO SOLICITATION OF ANY SUCH WRITTEN APPROVAL BY OR ON BEHALF OF LEGATO SHALL BE CONSTRUED AS AN INFERENCE TO THE CONTRARY.

LICENSE AND TERM

Legato and any applicable sublicensors grant to you a non-exclusive, non-transferable license to use the server software programs and related documentation in this package (collectively referred to as the "Software") on one licensed server and to use the client software programs and related documentation in their package (also referred to as the "Software") on licensed client processing unit(s). If the Legato product that you acquired is an upgrade, then the terms and conditions of this agreement are transferred to the upgraded product and the licensing agreement related to the preceding product is null and void. Any attempted sublicense, assignment, rental, sale or other transfer of the Software or the rights or obligations of this Agreement without the prior written consent of Legato shall be void. This License will automatically terminate without notice to you if you fail to comply with its terms. This Agreement will be governed by the laws of the State of California.

The Software and documentation are copyrighted. You may make copies of the Software only for backup and archival purposes. Unauthorized copying, reverse engineering, decompiling, disassembling, and creating derivative works based on the Software are prohibited. Title to the Software is not transferred to you by this license. Ownership and title to the Software and to the actual contents of this package, including the copy of the Software and the media on which it is stored and the associated documentation, are retained by Legato.

SOFTWARE LIMITATIONS

Legato does not warrant that the Software will be free from error or will meet your specific requirements. You assume complete responsibility for decisions made or actions taken based on information obtained using the Software. Any statements made concerning the utility of the Software are not to be construed as unexpressed or implied warranties. LEGATO MAKES NO WARRANTY, INCLUDING BUT NOT LIMITED TO ANY IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE REGARDING THE SOFTWARE AND MAKES THE SOFTWARE AVAILABLE SOLELY ON AN "AS IS" BASIS. NEITHER LEGATO NOR ANY SUBLICENSOR SHALL BE RESPONSIBLE FOR INCIDENTAL, SPECIAL, OR CONSEQUENTIAL DAMAGES, EVEN IF LEGATO HAS BEEN APPRISED OF THE LIKELIHOOD OF SUCH DAMAGES.

LIMITED WARRANTY ON MEDIA AND DAMAGES DISCLAIMER

Subject to the conditions and limitations on liability stated herein, Legato warrants for a period of thirty (30) days from the delivery of the first copy of each type of SOFTWARE, as so delivered, will materially conform to Legato's then current documentation for such SOFTWARE. This warranty covers only problems reported to Legato during the warranty period. ANY LIABILITY OF LEGATO WITH RESPECT TO THE SOFTWARE OR THE PERFORMANCE THEREOF UNDER ANY WARRANTY, NEGLIGENCE, STRICT LIABILITY OR OTHER THEORY WILL BE LIMITED EXCLUSIVELY TO PRODUCT REPLACEMENT OR, IF REPLACEMENT IS INADEQUATE AS A REMEDY OR, IN LEGATO'S OPINION, IMPRACTICAL, TO REFUND OF THE LICENSE FEE. EXCEPT FOR THE FOREGOING, THE SOFTWARE IS PROVIDED "AS IS" WITHOUT WARRANTY OF ANY KIND INCLUDING WITHOUT LIMITATION, ANY WARRANTY OF FITNESS OR MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. FURTHER, Legato DOES NOT WARRANT, GUARANTEE, OR MAKE ANY REPRESENTATIONS REGARDING THE USE, OR THE RESULTS OF THE USE, OF THE SOFTWARE OR WRITTEN MATERIALS IN TERMS OF CORRECTNESS, ACCURACY, RELIABILITY, OR OTHERWISE. The Licensee understands that Legato is not responsible for and will have no liability for hardware, software, or other items or any services provided by any persons other than Legato. Legato shall have no liability for delays or failures beyond its reasonable control.

LEGATO SHALL NOT BE RESPONSIBLE OR LIABLE WITH RESPECT TO ANY SUBJECT MATTER OF THIS SOFTWARE LICENSE AGREEMENT OR THE AGREEMENT OF WHICH IT IS A PART OR ANY ATTACHMENT, PRODUCT ORDER, SCHEDULE OR TERMS OR CONDITIONS RELATED THERETO UNDER ANY CONTRACT, NEGLIGENCE, STRICT LIABILITY OR OTHER THEORY: A) FOR LOSS OR INACCURACY OF DATA OR (EXCEPT FOR RETURN OF AMOUNTS PAID TO LEGATO THEREFOR), COST OF PROCUREMENT OF SUBSTITUTE GOODS, SERVICES OR TECHNOLOGY, B) FOR ANY INDIRECT, INCIDENTAL OR CONSEQUENTIAL DAMAGES INCLUDING, BUT NOT LIMITED TO LOSS OF REVENUES AND LOSS OF PROFITS; OR C) FOR ANY MATTER BEYOND ITS REASONABLE CONTROL.

**Legato NetWorker Administrator's Guide, UNIX Version
October 1993 - P/N 22-1627-01**

This document was printed on recycled paper.



LEGATO

Preface

The *Legato NetWorker Administrator's Guide* explains how to configure and manage NetWorker's backup and recover software. This manual provides enough information for the basic use and administration of NetWorker installed in a typical networked environment. For more detailed information about the NetWorker commands, refer to the on-line manual pages after you have installed NetWorker. Appendix C, "Theory of Operations," provides an overview of how NetWorker works.

Audience

This manual is intended for administrators who install software and maintain the servers and clients on a network. This manual may also be used by operators who monitor the daily backups of the machines on a network and maintain the backup media and devices.

Conventions

There are some conventions used in this manual to make information easy to access and understand. The following list describes the conventions:

- ❑ Command names are displayed in a **bold** typeface. For example:

To invoke NetWorker, enter the **networker** command at the system prompt.

- ❑ Examples, shell prompts, and information displayed on the screen are displayed as a fixed-width typeface. For example:

```
NetWorker successfully installed on 'mars'!
```

- ❑ Examples of what *you* type are displayed as a bold fixed-width typeface. For example:

```
mars% networker &
```

- ❑ Names of buttons, displays, menus, and windows are displayed as a Helvetica typeface. For example:

Press the Mount button.

Requests for mounting the correct backup media appear in the Pending display.

- ❑ Important pieces of information and cautionary notes that prevent you from making a mistake are marked as "Important." For example:

Important ⇨ Using a label that is already assigned to a backup volume destroys the records currently in that volume.

- ❑ 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 as a "Tip." For example:

Tip 📌 You may want to store a copy of the weekly or monthly full backup off-site.

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

```
mars% man nsr
```

- Directory *pathnames*, machine names, and new terms defined in the glossary are displayed as an italic typeface. For example:

The NetWorker programs are in the directory */usr/opt/networker/bin*.

Log in to the server, named *mars*.

The *media pool* is the collection of backup media recognized and managed by NetWorker.

Contacting Data General

Data General wants to assist you in any way it can to help you use its products. Please feel free to contact the company as outlined below.

Manuals

If you require additional manuals, please use the enclosed TIPS order form (United States only) or contact your local Data General sales representative.

Telephone Assistance

If you are unable to solve a problem using any manual you received with your system, free telephone assistance is available with your hardware warranty and with most Data General software service options. If you are within the United States or Canada, contact the Data General Customer Support Center (CSC) by calling 1-800-DG-HELPS. Lines are open from 8:00 a.m. to 5:00 p.m., your time, Monday through Friday. The center will

put you in touch with a member of Data General's telephone assistance staff who can answer your questions.

For telephone assistance outside the United States or Canada, ask your Data General sales representative for the appropriate telephone number.

Joining Our Users Group

Please consider joining the largest independent organization of Data General users, the North American Data General Users Group (NADGUG). In addition to making valuable contacts, members receive FOCUS monthly magazine, a conference discount, access to the Software Library and Electronic Bulletin Board, an annual Member Directory, Regional and Special Interest Groups, and much more. For more information about membership in the North American Data General Users Group, call 1-800-253-3902 or 1-508-443-3330.

End of Preface

Contents

Preface	v
Audience	v
Conventions	v
Contacting Data General	vii
Manuals	vii
Telephone Assistance	vii
Joining Our Users Group	viii
Chapter 1 – Introduction	17
About this Guide	17
Starting NetWorker	19
The NetWorker Main Window	20
Using NetWorker Menus	21
The File Pull-down Menu	21
Changing the NetWorker Server	21
Changing the Polling Interval	22
Exiting NetWorker	23
The Operation Pull-down Menu	23
The Administration Pull-down Menu	24
The Help Pull-down Menu	25
Selecting a Tabular View	26
Chapter 2 – NetWorker Daily Operations	29
Using the Pre-configured Settings	29
Performing the NetWorker Daily Tasks	33

Checking the NetWorker Server Status	33
Labeling and Mounting Backup Volumes	35
Labeling a Backup Volume	36
Labeling and Mounting a Backup Volume	38
Labeling a Backup Volume without Mounting	39
Mounting and Unmounting Backup Volumes	40
Finding a Backup Volume Name	42
Starting a Network-wide Backup	42
Enabling a Scheduled Backup	43
Starting a Scheduled Backup Immediately	44
Viewing the Group Status	45
Monitoring a Scheduled Backup	46
Starting a Scheduled Backup Immediately	47
Stopping a Backup Group	48
Re-starting a Backup Group	48
Displaying the Backup Details	48
Handling Open Files	50
Displaying Backup Volume Information	50
Managing the Indexes	52
Checking NetWorker Notices	54
Chapter 3 – Configuring the NetWorker Server	55
Configuring Backup Groups	56
Creating a Backup Group	57
Re-trying a Failed Client Backup	59
Changing a Backup Time	60

Deleting a Backup Group	61
Setting up the Backup Schedules	61
The Logic Behind the Scheduling Schemes	64
Understanding the Backup Levels	65
Using the Pre-configured Schedule	67
Suggestions for Easier Operation	67
Creating a New Backup Schedule	68
Example 1 – A Weekly Schedule	69
Example 2 – A Monthly Schedule	71
Summary	74
Using Label Templates	75
Using the Pre-configured Templates	75
Choosing a Label Template	77
Creating a Custom Label Template	77
Example 1 – A Storage-oriented Template	80
Example 2 – A Sequential Label Template	83
Using Volume Pools	85
Using the Pre-configured Pools	91
Creating a New Pool	94
Pool Examples	98
Example 1 – Separate Departments	98
Example 2 – Secure Device	98
Example 3 – Targeted Backups	99
Example 4 – Specific Filesystem	99
Creating Index Policies	100
Adding and Deleting Backup Devices	106

Selecting Server Options	109
Changing Server Parallelism.....	110
Using Concurrent Devices	112
Adding or Changing Administrators.....	114
Using Directives	116
Setting Event Notification.....	119
Log Default.....	122
Index Size	122
Savegroup Completion	123
Backup Media Request Notices	123
Configuring the Server's Clients.....	124
Adding a New Client	128
Selecting Policies for a Client.....	130
Scheduling Large Client Filesystems.....	131
Removing a Client.....	135
Securing the NetWorker Environment	135
Restricting Client Access	136
Chapter 4 – Managing the On-line Indexes	137
On-line Indexes.....	137
Save Sets.....	138
Index Actions.....	140
Managing the Indexes Manually	140
Removing the Oldest Cycle	144
Reclaiming Index Space	145
Viewing Save Set Details	147

Managing the Volumes Manually	150
Viewing the Volumes Scrolling List	152
Viewing the Save Sets Scrolling List	153
Managing Backup Volumes	154
Setting the Location for Your Backup Volumes	155
Changing the Mode of a Backup Volume	156
Removing a Backup Volume	159
Chapter 5 – Recovering from a Disk Crash	161
Preparing for a Crash	163
File the Bootstrap Information	164
File the Disk Information	165
Recovering the NetWorker Index Files	168
Replacing the /nsr/res Directory	171
Finding the Bootstrap Save Set ID	172
Recovering a Filesystem	173
Summary	174
Appendix A – Theory of Operations	175
Introduction	175
NetWorker Media Format	176
The NetWorker System at Rest	177
Browsing and Modifying Resources	179
Backing up Files	181
Recovering Files	183
Backing up the Whole Machine	185
Automatic Network-wide Backups	187

Hints for Efficient Network Saves	189
Event Processing	191
Media Management.....	192
Index Management	194
Disaster Recovery	195
Appendix B – Managing the NetWorker Environment	197
NetWorker Backup Schedules.....	197
Backup Levels	197
Full Backups versus Incremental Backups.....	198
Using Level Backups	199
A Typical Monthly Backup Schedule	202
Backups Take Time	203
Staggering the Backup Schedules.....	204
Convenience versus Security	205
Backup Schedule Examples	206
Default	207
Full Every Friday	208
Full on 1st Friday of Month	209
Full on 1st of Month.....	210
Quarterly.....	211
NetWorker Browse and Retention Policies.....	213
Browse Policy.....	213
How Browse Policies Work	214
Reclaiming Disk Space	215
Recovering Files Removed from the Index.....	216

Media Retention Policy.....	216
Policy Examples.....	218
Week Policy.....	219
Month Policy.....	220
Quarter Policy.....	221
Year Policy.....	222
Decade Policy.....	223
Summary.....	224
Guidelines for Choosing a Configuration.....	224
Tape Drives.....	225
Clients.....	225
Network.....	226
Server.....	226
Measuring Performance.....	227
Server Performance.....	227
Backup Device Speed.....	227
CPU Speed.....	228
Memory.....	228
Client Performance.....	228
Filesystem Traversing.....	229
Data Generation Rate.....	229
Data on Multiple Disks.....	229
CPU Speed.....	230
Appendix C – Command Summary.....	231
Appendix D – Troubleshooting.....	237

Checking the NetWorker Daemons.....	237
Using the networker -x Option	237
Renaming a Client	238
Recover Access Issues	239
Previewing a Backup	239
Halting a Network Backup	240
Backup Media Capacity.....	241
Savegroup Completion Messages.....	242
Binding to Server Errors	242
Saving Remote Filesystems	243
File Changed During Backup	245
Cannot Print Bootstrap Information	245
Maximum Filename Length.....	245
Errors in Executing External ASMs.....	246
Savegroup Completion Warning Messages	246
Moving Indexes	247
Recovering Files from an Interrupted Backup	247
Determining the NetWorker Server.....	248
Using nsrexecd.....	248
Glossary	249
Index	257

Chapter 1

Introduction

Everyone knows about 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.

Legato 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. And NetWorker creates a database of your backups, making it easy to recover files. As the network and volume of files expands, NetWorker has the capacity and performance to handle the load.

About this Guide

The *Administrator's Guide* contains information on the day-to-day NetWorker tasks such as managing the backup volumes and monitoring the backups, and the less frequent system administration tasks involved in configuring the backup environment using NetWorker. The more frequent tasks, as well as how to use NetWorker with the pre-configured settings, are documented in Chapter 2, "Daily Operations." The custom configuration tasks are documented in Chapter 3, "Configuring the NetWorker Server."

1

NetWorker's X Window System interface will greatly facilitate your operation and system administration of NetWorker. Through NetWorker's window interface, many system configuration and operation tasks will be intuitive. In addition, every window contains on-line help so you can easily get information about the tasks you are trying to complete.

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 pre-selected 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. 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.

This *Administrator's Guide* assumes that NetWorker has already been installed on your server and clients. If it is not installed, refer to the NetWorker release notice for installation instructions.

The day-to-day tasks are described in Chapter 2, "Daily Operations." These tasks include the following:

- using the pre-configured settings
- checking the server status
- managing the backup volumes
- labeling and mounting backup volumes
- monitoring the indexes

- ❑ checking NetWorker notices

The configuration tasks are described in Chapter 3, "Configuring the NetWorker Server." Typical system configuration tasks include:

- ❑ adding backup devices to the NetWorker server
- ❑ setting up notification from the NetWorker server about NetWorker events
- ❑ creating backup schedules for the clients for automatic network-wide backups
- ❑ creating policies for the clients for automatic index management
- ❑ adding and configuring the clients
- ❑ creating label templates
- ❑ using and creating volume pools

Starting NetWorker

To start NetWorker, enter the **networker** command at the system prompt.

```
# networker &
```

Or, if the NetWorker icon is displayed on your screen, select the icon to open the Main window.



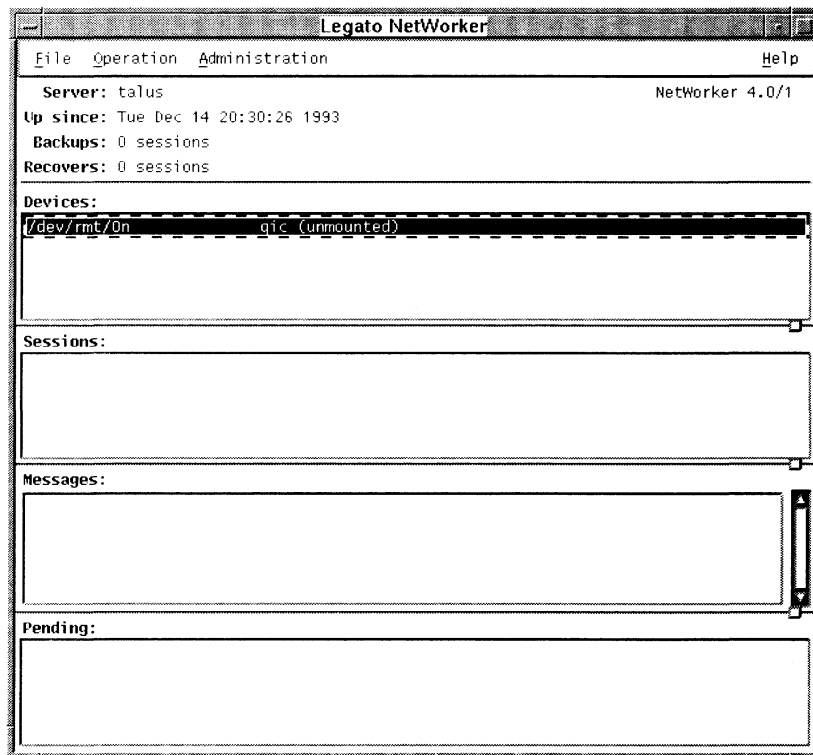
The NetWorker icon

The Main window appears when you start NetWorker.

The NetWorker Main Window

The NetWorker Main window is a dual purpose window: it serves as both a starting point for all NetWorker tasks and as a monitor for server activity.

The Main window, shown on the next page, contains the pull-down menus for all operation and administration tasks. You select all NetWorker tasks from the pull-down menus in the Main window.



Using NetWorker Menus

This section includes information about the pull-down menus in the Main NetWorker window and how to use them.

The File Pull-down Menu

The File menu has selections for changing the current NetWorker server, changing the polling interval, and for exiting NetWorker.

File	Operation
	Change Server...
	Change Polling...
	Exit

Changing the NetWorker Server

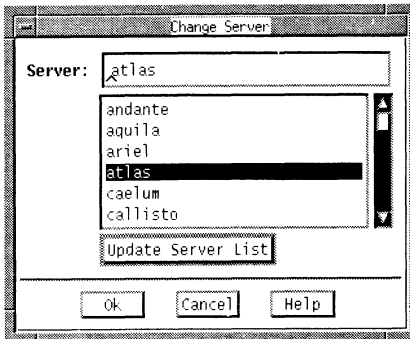
Select Change Server to choose another NetWorker server to monitor or configure. Most networks need only one NetWorker server to back up all the systems on the network. If you have more than one NetWorker server, however, choose the one you want to monitor or configure by selecting its *hostname* in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

To change NetWorker servers:

1. Select Change Server from the File menu in the Main window and the Change Server dialog box will appear.
2. Use the Update Server List button to display all of the NetWorker servers on your network.
3. Highlight the *hostname* of the desired server in the Server scrolling list, or enter a valid *hostname* in the Server field.

1

4. Press Ok to monitor and configure the new server in all the NetWorker windows.




NOTE: All operation and administration tasks apply to the server selected in the Main window.

Changing the Polling Interval

Change Polling allows you to change the polling interval. Polling controls how often, in seconds, the data in the Main window is updated.

To change the polling interval:

1. Select Change Polling from the File menu in the Main window.



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

Exiting NetWorker


To exit NetWorker and return to the system prompt, select Exit from the File menu or press both the [Alt] key and the [F] key at the same time, release them, and then press [X].

The Operation Pull-down Menu

The Operation pull-down menu lists the commands for operating NetWorker on a daily basis. These commands are explained in this manual and in the *Legato NetWorker User's Guide*. All of the commands on the Operation menu have a speed key sequence that is represented by a line underneath a letter in each command. For the commands in the Operation menu, hold down the [Alt] key and the [O] key at the same time, release them, then press the key which corresponds to the underlined letter in the command name.


The Operation pull-down menu contains the volume mounting and labeling commands (described in Chapter 2), the manual backup and recover commands (described in the *Legato NetWorker User's Guide*), and the commands for managing the indexes and volumes (described in Chapter 4).

The Operation pull-down menu is shown below:



The Administration Pull-down Menu

The Administration pull-down menu contains the commands for administering and configuring the NetWorker server and clients (explained in Chapter 3).




Each of the commands on the Administration menu has a corresponding speed key sequence for selecting it. Hold down the [Alt] key and the [A] key at the same time, release them, and then press the key which corresponds to the underlined letter in each command.

NOTE: The first and last two selections on the Administrative pull-down (Registration, Label templates, and Jukeboxes) refer to options that are not available on AViiON systems.

The Help Pull-down Menu

All NetWorker windows contain a Help pull-down menu or button. The Help menu offers you on-line help on the window as well as a menu for selecting other types of help.

The Help pull-down menu is shown below:



The Help pull-down menu has four commands:

- On Window – displays a help screen with information about the currently selected window.
- Topics – to select help from a list of topics.
- On Help – gives you information on how to use the Help facility.
- On Version – displays the NetWorker software release version for the server.

After this introduction to the Main window and the pull-down menus, you are ready to operate and configure NetWorker using the menus.

1

Selecting a Tabular View

NetWorker lets you display information in a compact, tabular format. The tabular view is especially helpful where there is lots of information to be displayed.

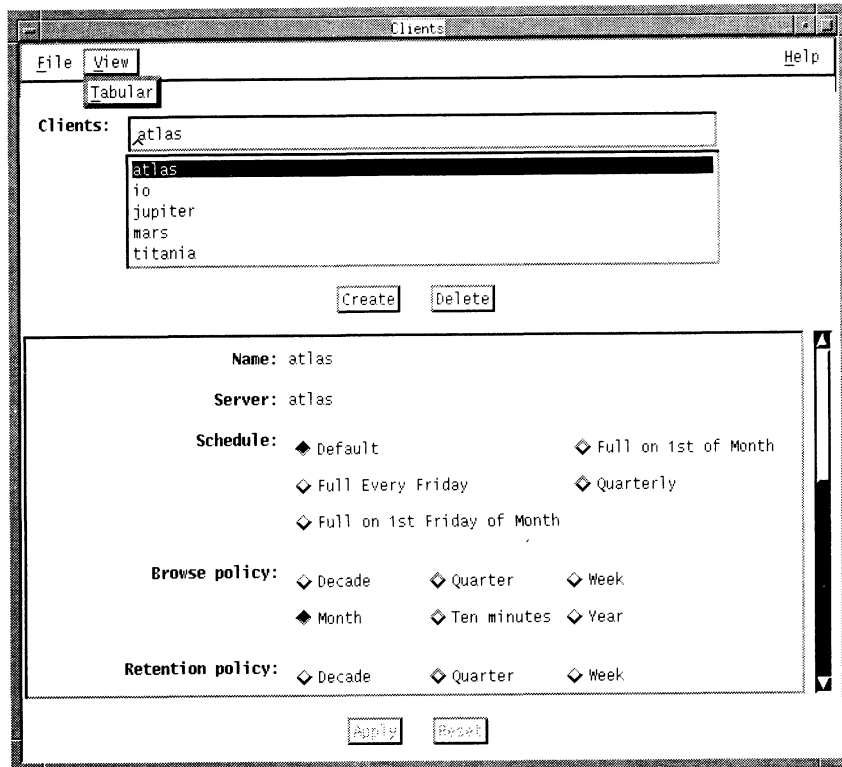
For example, in the Clients window provides two ways to view client information:

- you can see configuration information for one client at a time
- or, you can choose to view configuration information for all clients in one tabular view.

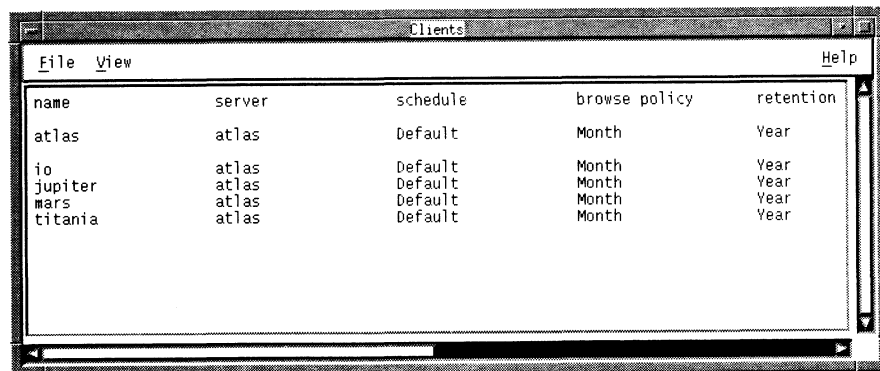
To choose the tabular view for the Client window, follow these steps:

1. Select the Clients command from the Administration pull-down menu. The Clients window appears.

The Tabular command is in the View pull-down menu



2. Select the Tabular command from the View pull-down menu. The Clients window changes to a tabular view, as shown below.



1

Use the scroll bars in the Clients tabular view window to see all of the client configuration information.

To return to the normal Clients dialog box, select the Normal command from the View pull-down menu.

The tabular view is available for all of the administrative commands on the Administration pull-down menu except for the Schedules command.

Chapter 2

NetWorker Daily Operations

This chapter describes how to operate NetWorker using the pre-configured settings. It also describes the everyday tasks for operating NetWorker.


Using the Pre-configured Settings

NetWorker is shipped with easy-to-use pre-configured settings. Pre-configured settings are existing choices NetWorker has established for backing up the clients on your network. You are free to use these pre-configured selections, change them, or create your own configurations. By using these pre-configured settings you can start backing up your data as soon as you install NetWorker.


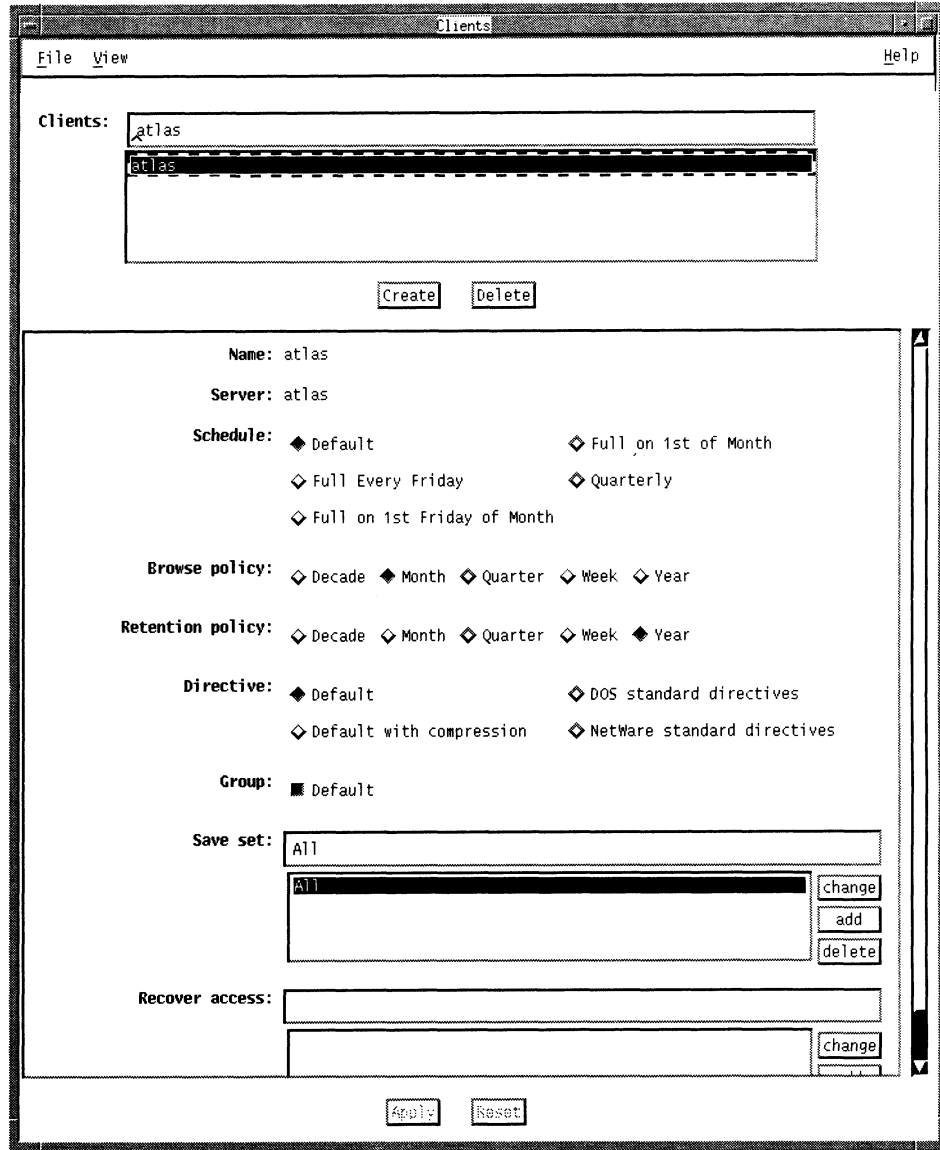
After installation on a server, NetWorker backs up the NetWorker server as a client of itself using these pre-configured settings. If you are backing up other client machines to your server, all you have to do (after installing NetWorker on the clients) is add the clients by using the Clients window. NetWorker may then back up your clients using the pre-configured settings.

Open the Clients window by selecting the Clients command from the Administration pull-down menu.

2



The Clients window with the pre-configured settings is shown below:



NOTE: The NetWorker server will be listed in the Clients scrolling list prior to creating any clients. To back up the NetWorker server as a client of itself using the pre-configured settings, all you have to do is Enable the Autostart feature of the Default group in the Groups window. The server is automatically created as a client of itself when NetWorker is installed.

Your NetWorker server and newly created clients have the following pre-configured settings, unless you select otherwise:

- ❑ **Schedule** – The client uses the Default schedule of a full backup every Sunday and an incremental backup every other day of the week. The Default schedule is described in detail in the section “Backup Schedule Examples” of Appendix B.
- ❑ **Browse policy** – The Browse policy for entries in the client’s on-line 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 being automatically removed. For information about the Browse policy see Chapter 3 “Creating Index Policies.”
- ❑ **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 core and /tmp files and backs up mail files without 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 morning, change the Autostart feature to Enabled in the Groups window.

- ❑ Save set – The pre-configured selection is “All.” All of the local client filesystems will be backed up.
- ❑ Recover access – This field is blank, meaning that only the client itself can recover its backed-up files.

Performing the NetWorker Daily Tasks

2

Generally, there are five daily tasks required to operate NetWorker:

- ❑ checking the server status
- ❑ managing the backup volumes
- ❑ labeling and mounting backup volumes
- ❑ monitoring the indexes
- ❑ checking NetWorker notices

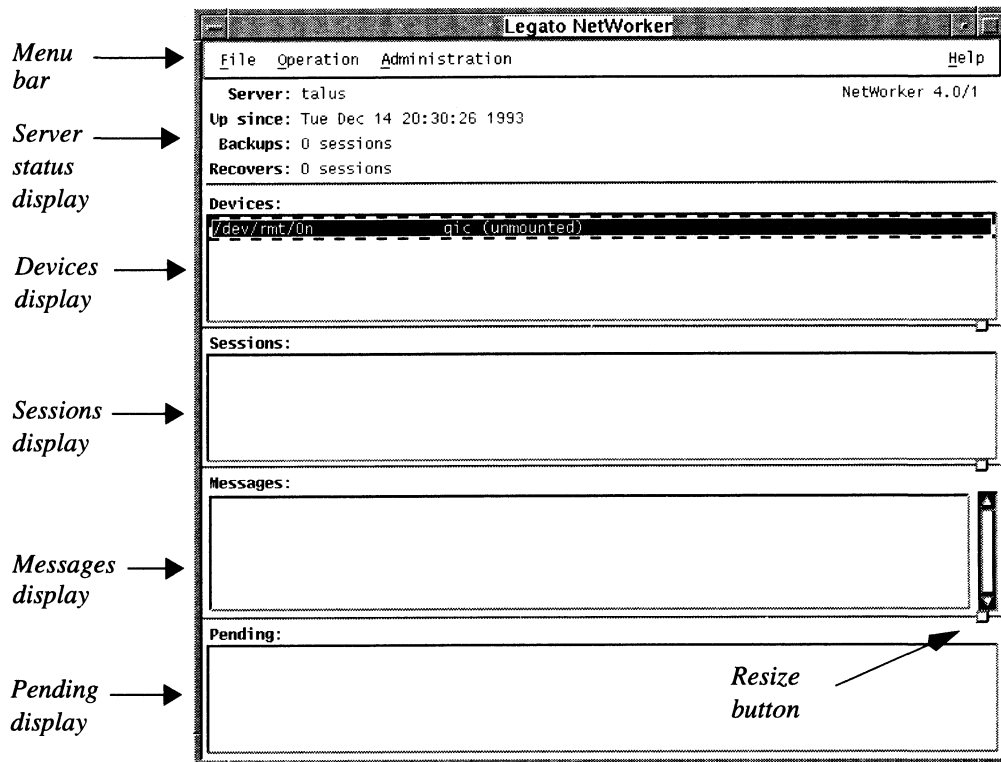
This section describes each of these tasks.

Checking the NetWorker Server Status

The Main window displays the status of the NetWorker server in five displays. If you do not already have the Main window on your screen, enter the **networker** command at the system prompt to start NetWorker and to display the Main window.

```
# networker &
```

The NetWorker Main window appears, as shown below.



The Main window has a menu bar and five displays:

A scroll bar appears to the right of the display area when there is more information to see. Use the resize buttons for resizing the display areas.

- ❑ Menu bar – contains the File, Operation, Administration, and Help pull-down menus.
- ❑ Server status display – tells you how long the current server has been on-line, the number of backups and recovers the server has completed since it has been on-line, and the amount of

data contained in the backups and recovers. It also contains the name of the NetWorker product you are using.


- ❑ Devices display – shows you the devices known to the server and which backup volume, if any, is mounted on the device.
- ❑ Sessions display – provides current information on which clients are 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.

2

Labeling and Mounting Backup Volumes

To track backup volumes, NetWorker requires that you label and mount a backup volume before it can be used to back up or recover files. Use the Operation menu in the Main window to mount, unmount, and label backup volumes.

The Operation pull-down menu is shown below:



The labels of the backup volumes are tracked by the NetWorker 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 which backup volume to mount when you want to recover lost data.

Labeling a Backup Volume

2

Before using a backup volume for backing up files, you need to label it. Labeling a backup volume provides NetWorker with a unique name so that it can be recognized later.

Every backup volume belongs to a volume pool. You may use the pre-configured pools or create your own. If you choose not to separate your backup volumes into pools, NetWorker will automatically use 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 pre-configured label templates, or create a custom label template using the Label Templates window.

NetWorker will automatically label the backup volumes with the next sequential label from a template unless you choose to override the label with a label not associated with a pool. Volume pools and label templates are described in detail in Chapter 3.

Each backup volume should have an adhesive label attached to it which matches the internal NetWorker label. There are no restrictions for labeling the backup volumes except that each volume requires a unique label. If you choose not to use a pre-configured template, you can label the volumes with names appropriate to your network. For example, you may choose to label your volumes 1, 2, 3, or A,B,C.

Tip 

In the NetWorker environment, each NetWorker server has a 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 this backup volume is from the NetWorker server named *atlas*

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

Naming schemes vary from site to site. At Legato we name the backup volumes simply with the name of the NetWorker server followed by a three-digit number. For example:

atlas.001 atlas.002 ... atlas.999

No matter what naming scheme you choose, keep in mind that a convention should be easily understood by operators and administrators.

Important ⇨ Backup volumes are not overwritten, but 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 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.

Use the `Label` and `Mount` command if you want to label a backup volume and immediately mount it in the backup device. This volume will be available for backups.

If you want to label a volume without mounting it, use the `Label` command. Use this choice if you would like to label several backup volumes at once without mounting any of them.


Labeling and Mounting a Backup Volume

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

1. Place a blank backup volume in the NetWorker server backup device.

Your NetWorker server may have one or more backup devices connected to it. The device currently highlighted in the Devices display of the Main window will appear in the Label and Mount dialog box. To select another device, press Cancel to close the Label and Mount dialog box, highlight the new device in the Main window, and select Label and Mount again.

2. In the Main window, select Label and Mount from the Operation pull-down menu to display the Label and Mount dialog box.



3. Choose a volume pool from the Pool choices. Only the pools that have been enabled in the Pools window appear in the Label and Mount dialog box. Both the Default and the Archive pools are pre-configured to be enabled. NetWorker will automatically use the label template associated with the pool to label the backup volume. For more information about volume pools, see "Using Volume Pools" in Chapter 3.
4. If you do not want NetWorker to label the backup volume by following the label template associated with the volume pool,

delete the contents of the Volume name field and enter a name for the backup volume into the field.

5. Press the Ok button when you are finished naming a backup volume, or if you make a mistake press the Cancel button to cancel the label operation.

Pressing Ok labels and mounts the backup volume. The name of the backup volume appears in the Devices display of the Main window beside the path of the device.


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

Important ⇨ Re-labeling a backup volume destroys all record of its contents under the old label. That is why NetWorker asks for confirmation with the question "About to re-label." For more information about re-labeling backup volumes, see Chapter 4, "Managing the On-line Indexes."

Labeling a Backup Volume without Mounting

You may want to label several backup volumes in one session and not mount any of them. To label a backup volume without mounting it, follow these steps:

1. In the Main window, select Label from the Operation pull-down menu to display the Label dialog box.



The device that is currently highlighted in the Devices display of the Main window is selected for labeling volumes.

2. Choose a volume pool from the Pool choices.

NetWorker will use a label template associated with the pool and automatically label the backup volume with the appropriate name.

3. If you do not want NetWorker to use a label template, delete the contents of the Volume name field, and enter the name for the backup volume into the Volume name field.
4. Press the Ok button.
5. After each labeling operation is complete remove the newly labeled volume from the drive and repeat the label procedure for each volume that you wish to label without mounting.

Pressing the Ok button labels the backup volume in the currently selected device in the Devices display. You will see the name of the backup volume in the Devices display of the Main window with an (unmounted) message following the media type. You must mount the newly labeled volume before sending backups to it. (See "Mounting and Unmounting Backup Volumes," in this chapter for instructions on mounting volumes.)

Important ⇨ You cannot use the same name for a volume label more than once. You can, however, re-label a volume. Relabeling a volume destroys the records for the volume under the previous label.

Mounting and Unmounting Backup Volumes

Before you can back up files to a backup volume or recover files from one, you must mount the volume in the server backup device. You may 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.

To mount a backup volume in a specific device:

1. Select the device name in the Devices display of the Main window.
2. Select Mount from the Operation pull-down menu.

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

Tip 

NetWorker allows you to pre-mount 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:

1. Select the name of the device in the Devices display.
2. Select Unmount from the Operation pull-down menu.

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 the following message in the Pending display:

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

A writable volume is one that is labeled and has room for data.

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


```
media waiting: recover waiting for 8mm 5GB tape volume
name
```

If you need more than one backup volume to recover the file(s), NetWorker gives you a list of all the backup volumes in the order

they will be needed in the Pending display, so you can prepare to mount all of them at once. As the files are recovered, NetWorker requests each backup volume it needs, one at a time.

Finding a Backup Volume Name

If the physical label on the backup volume is missing or illegible, you can find out its name by placing it in a device and pressing the Label button. When the Label dialog box appears you can see the name of the label in the Volume name field. Once you have determined the name of the volume press Cancel, because you do not want to rename the volume, you are just determining the existing label name.



Starting a Network-wide Backup

In this section you will learn how to start a network-wide backup. NetWorker provides you with the flexibility of backing up your files in several different ways. Using the Groups window, you can enable or disable a scheduled backup, or initiate a scheduled backup immediately. Using the Group Control window, you can start, stop, or re-start a scheduled backup.

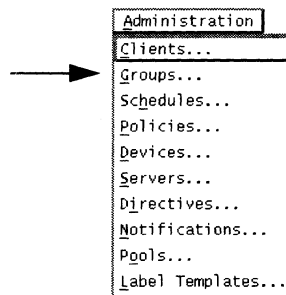
Enabling a Scheduled Backup

NetWorker is shipped with a Default group with an automatic backup start time of 3:33 a.m. After you install NetWorker and enable the Default group, it will automatically initiate a network-wide backup at this time every day, and back up the clients according to the Default backup schedule.

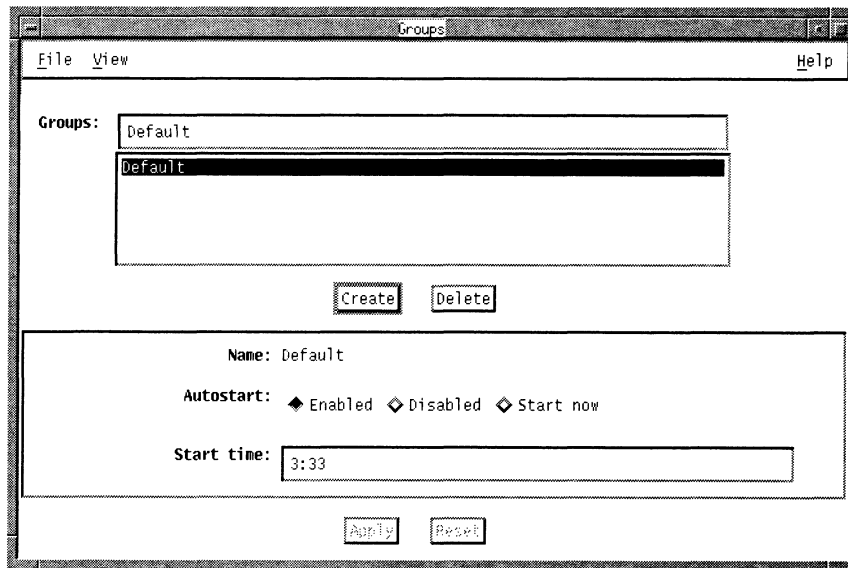
First you need to label the backup volumes and mount them in the NetWorker server's backup devices before you go home.

To enable a scheduled backup, follow these steps:

1. Select Groups from the Administration pull-down menu.



The Groups window appears, as shown below.



2

2. Select the Enabled from the Autostart choices.

NetWorker will begin backing up the clients in the Default group at the time shown in the Start time field.

Starting a Scheduled Backup Immediately

If you want to start a scheduled backup immediately, you may do so with the Groups window. This is useful, for example, if you must disable the network for repair and want to back up all the machines immediately.

To start a scheduled backup immediately, follow these steps:

1. Select Groups from the Administration pull-down menu.
2. Select the name of the group of clients you want to back up in the Groups scrolling list. You can select one group at a time.

3. Select Start now from the Autostart choices.

Name: Accounting
Autostart: Enabled Disabled Start now

4. Press the Apply button.

NetWorker will start the scheduled backup of the clients in the selected backup group. For more information about backup groups, see "Creating a Backup Group" in Chapter 3. The backup level is determined by the schedule in effect for the clients in the selected group. For more information about backup schedules, see "Setting up the Backup Schedules" in Chapter 3.

NOTE: You can also start a backup immediately by selecting the Start button in the Group Control window. For more information on the Group Control window see the next section in this chapter "Viewing the Group Status."

Tip 

Make sure the NetWorker server has a labeled backup volume mounted in its device before you initiate a backup using Start now. If you do not, NetWorker requests a volume in the Pending display of the Main window so it can start the backup for the clients in the group. See "Labeling and Mounting a Backup Volume" in this chapter for more information.

Viewing the Group Status

You can use the Group Control window to:

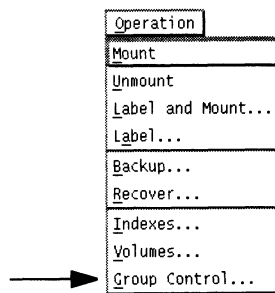
- monitor a scheduled backup
- view the details of a scheduled backup

- ❑ start a scheduled backup immediately
- ❑ stop a scheduled backup
- ❑ re-start a scheduled backup which you have stopped

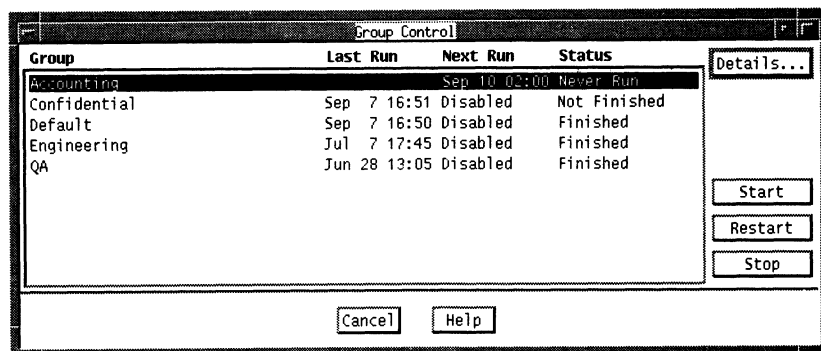
The next sections describe how to use the Group Control window for these tasks.

Monitoring a Scheduled Backup

Open the Group Control window by selecting the Group Control command from the Operation pull-down menu.



The Group Control window appears, as shown below.



The Group Control window displays information about all of the backup groups known to the NetWorker server.

- ❑ Group – displays the name of the backup group
- ❑ Last Run – displays the date and time the group was last backed up
- ❑ Next Run – displays the date and time scheduled for the next backup of the group, or whether the scheduled backup is disabled
- ❑ Status – displays the progress of the backup. There are four progress messages:

Never Run – the group has never been backed up

Finished – the backup has finished

Not Finished – the backup is still running or the backup has exited prematurely without finishing

Starting a Scheduled Backup Immediately

You can start a scheduled backup immediately by using the Start now choice in the Groups window, or use the Start button in the Group Control window. Both will initiate a scheduled backup immediately of the group selected in the window.

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

1. Select the group you wish to back up immediately in the Group scrolling list.
2. Press the Start button.

NetWorker will back up the clients in the group immediately according to the backup schedule, and display “running” in the Status field of the Group Control window.

Stopping a Backup Group

If you need to stop a scheduled backup which is in progress, you may do so by using the Stop button in the Group Control window.

Follow these steps:

1. Select the group you wish to stop in the Group scrolling list of the Group Control window.
2. Press the Stop button.

NetWorker will halt the scheduled backup, and display “not finished” in the Status field of the group in the Group Control window.

Re-starting a Backup Group

If you wish to re-start a scheduled backup that you have stopped, use the Restart button in the Group Control window.

Follow these steps:

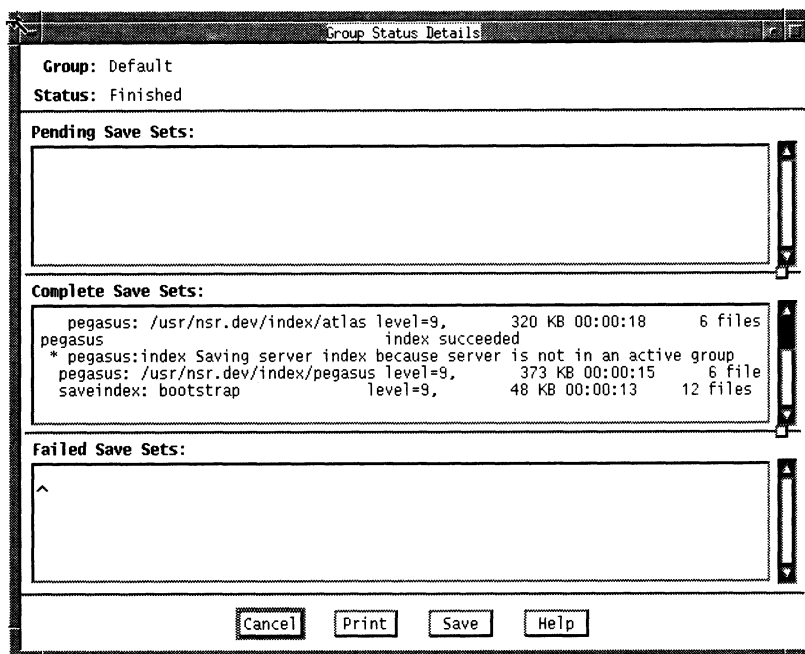
1. Select the name of the halted backup group in the Groups list of the Group Control window.
2. Press the Restart button.

NetWorker will resume the scheduled backup for the group, and display “running” in the Status field for the group.

Displaying the Backup Details

To view more detailed information about the backup of the group currently selected in the Group Control window, press the Details button.

NetWorker displays the Group Status Details window, as shown below.



2

There are three fields in the Group Status Details window which display information about each client in the backup group:

- Pending Save Sets – displays the filesystems which have not yet been backed up
- Complete Save Sets – displays the filesystems which NetWorker has successfully backed up
- Failed Save Sets – displays the filesystems which NetWorker was unable to back up

If NetWorker was unable to back up the filesystems for a client in the group, you need to determine the cause for the failed backup. Typical reasons for failed save sets include:

- the server crashed during the backup

- ❑ the client crashed during the backup
- ❑ the network connection failed during the backup

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

2

Handling Open Files

If a client's open file changes during a scheduled backup, NetWorker will back up the file and "notice" that it is changing. The file (save set) that changed during the backup will display a warning message in the Group Status Details window similar to this:


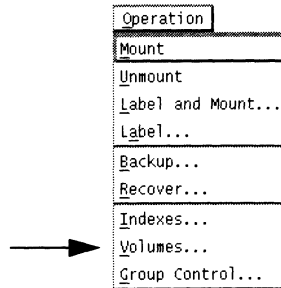
```
warning: file name changed during save
```

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

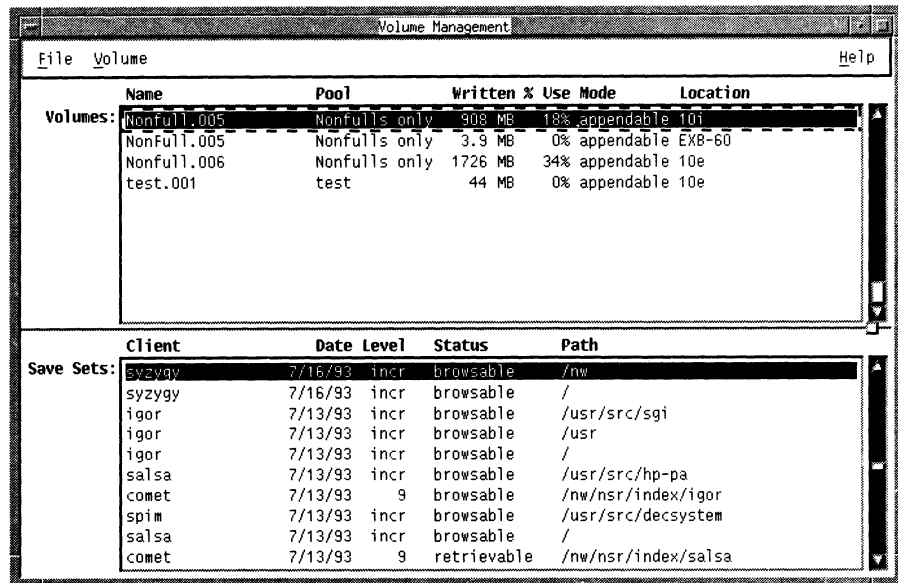
Displaying Backup Volume Information

The Volume Management window displays information about the backup volumes. It gives you a global look at the NetWorker server's collection of backup volumes.

To see the Volume Management window, select Volumes from the Operation pull-down menu.



The Volume Management window appears, as shown below.



The Volumes scrolling list displays the names and modes of the backup volumes with entries in the server media index. The mode indicates whether the volume is available to receive data. NetWorker uses volumes in appendable mode for backups. The Pool

2

column display the name of the pool to which the backup volume belongs. The Written column displays how much data is written to the media, and % Used displays how much of the media is available for more data. Location displays the location, if applicable, for the backup media.

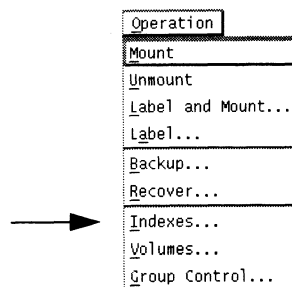
When you select a volume name in the Volumes scrolling list, its contents appears in the Save sets scrolling list. A *save set* is typically a filesystem backed up by NetWorker. This list displays the name of the client that generated the save set, the date of the backup, the backup level, the status of the save set, and path of the filesystem in the save set.

The Volume Management window is described in detail in Chapter 4, "Managing the On-line Indexes."

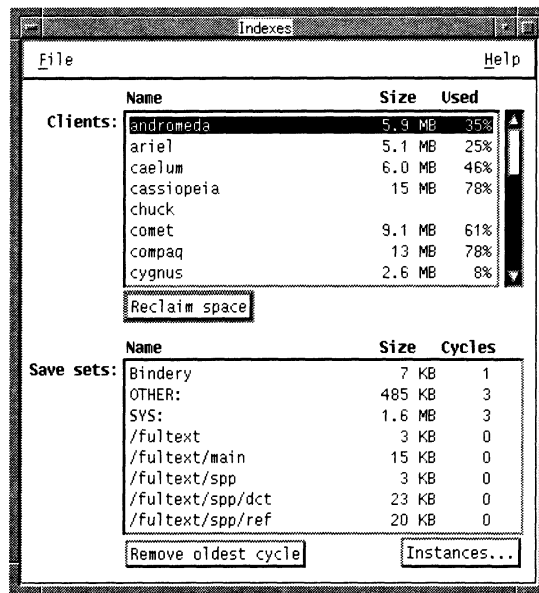
Managing the Indexes

Every time a backup completes, NetWorker creates entries for the files saved for each client in on-line indexes. These indexes require disk space and must be monitored to be sure they are not getting too large. Entries cannot 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.

To manually manage the indexes select Indexes from the Operation pull-down menu.



The Indexes window appears, as shown below.



2

Used indicates the percentage of the index file being used. If the percent listed is 100% that means 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 will grow automatically as the index size increases.

If the percentage is less than 100% you can reduce the size of the index by first selecting the client and then pressing the Reclaim space button. This will remove any holes in the index database created by the removal of index entries. Index entries are removed based on the browse and retention policies you have established for managing your on-line indexes. (For information on browse and retention policies, see the section entitled "Creating Index Policies" in Chapter 3.)

If the index requires further reduction, you may remove the oldest cycle of a save set for the client. Select the save set and then press the Remove oldest cycle button. After removing a cycle, press Reclaim space to remove the holes left by Remove oldest cycle.

For more information on managing the indexes, see Chapter 4, "Managing the On-line Indexes."

2

Checking NetWorker Notices

Checking NetWorker notices is an important daily task for 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 on-line indexes
- registration status

NetWorker is shipped with pre-configured notification instructions to send electronic mail to *root* about most events.

There are three ways for you to check NetWorker notices:

1. Log on as *root*, and read your electronic mail.
2. Look at the system console messages.
3. Use the NetWorker Main window to watch for notices in the Pending and Messages displays.

See "Setting up Event Notification" in Chapter 3 for more information.

Chapter 3

Configuring the NetWorker Server

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


- ❑ configuring backup groups
- ❑ setting up backup schedules
- ❑ creating index policies
- ❑ adding and deleting backup devices
- ❑ selecting server options
- ❑ using directives
- ❑ setting up event notification
- ❑ securing the NetWorker environment
- ❑ configuring the server's clients

Tip 

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

You select the system administration windows for the NetWorker server from the Administration pull-down menu in the Main window. The Administration menu is available for users who

have administration privileges for configuring the NetWorker server.

**3**

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 pre-configured backup group named "Default." The automatic backup of the Default group starts at 3:33 a.m. When a client is added, it automatically belongs to the Default backup group. If this backup start time suits your needs, you do not need to change it.

If you have a large number of clients, or if you prefer to start backing up certain machines at different times, you can easily do so. For example, you may want to start the backup of the engineering department machines at four o'clock in the morning, and all the other clients on the network at midnight.

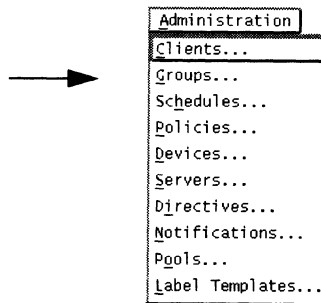
After creating a new backup group in the Groups window, you assign clients to the group in the Clients window.

You can have any number of backup groups on your NetWorker server. The clients in each backup group will start their automatic backups according to the start time of the group.

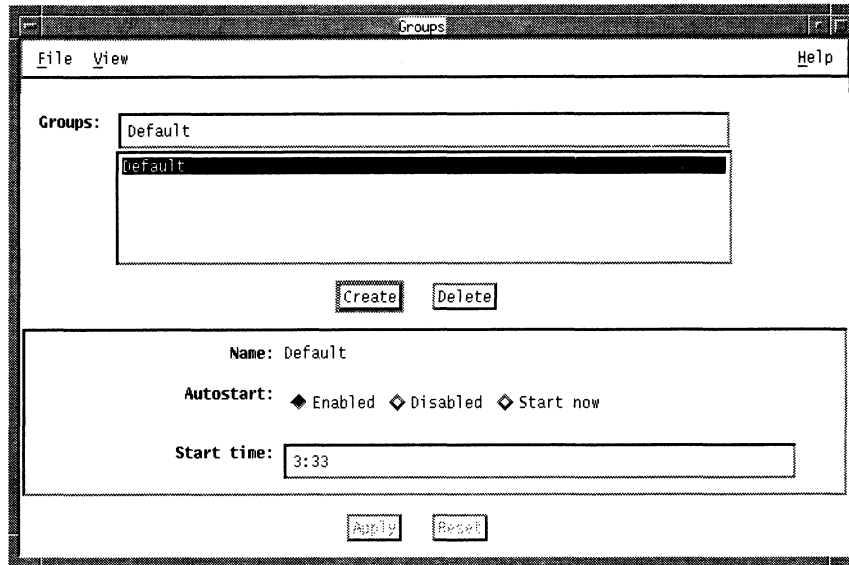
Creating a Backup Group

You create a backup group and set the group backup start time in the Groups window.

Select Groups from the Administration pull-down menu to open the Groups window.



The Groups window appears, as shown below.



3

The Groups scrolling list displays all the 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.

The Autostart choices allow you to disable the scheduled backup, or initiate a network-wide backup immediately. For example:

- Enabled – means allow the network-wide backup of the highlighted group to start at the scheduled time.
- Disabled – do not start the network-wide backup of the group at all.
- Start now – initiate the network-wide backup of the group immediately, instead of using the scheduled start time.

The Start time field displays the time NetWorker starts the backup of the clients in the group currently selected in the Groups scrolling list.

Re-trying a Failed Client Backup

NetWorker will normally try to back up a client in a group *once* during a scheduled backup. If NetWorker cannot back up a client in a group (because it is down, for example) it will not try again unless you configure it to do so.

Use the Client retries sliding bar in the Clients window to re-try failed clients during a scheduled backup:



1. Using the mouse, select the sliding bar and move it to the right or left.
2. The number of times NetWorker will re-try to back up the failed client appears in the Client retries field.

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

1. Open the Groups window.
2. Press the Create button.

Notice 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. 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.")
5. Use the sliding bar to set a value for client retries.
6. Press the Apply button to apply the changes.

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

NOTE: If you get the error message “user *user_name* needs to be on administrator’s list,” it means you do not have permission to make configuration changes. See “Adding or Changing Administrators” in this chapter for more information.

Your new group name will appear in the Groups scrolling list.

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 “Configuring the Server’s Clients” in this chapter for more information.



3

Changing a Backup Time

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

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

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

Deleting a Backup Group

NetWorker will 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. Select the name of the group you want to delete from the Groups scrolling list. Press the Delete button.

If any clients belong to the group you are trying to delete, you will see an "unable to delete" message. If no clients belong to the group, you will see an "ok to delete?" message.

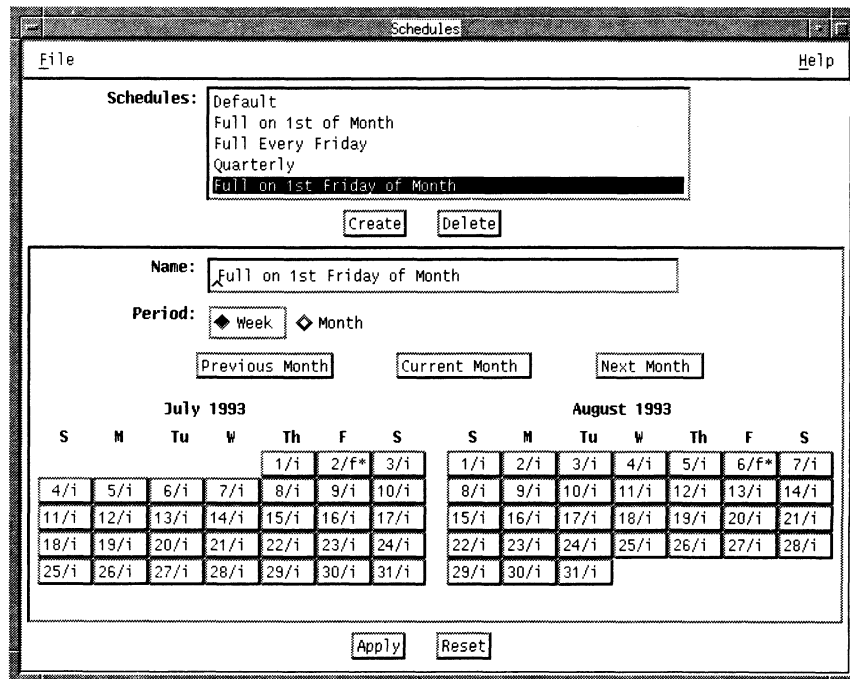

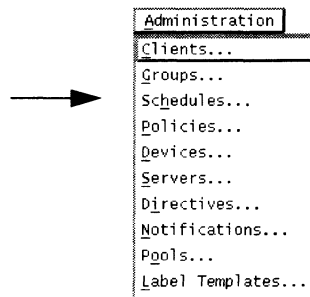
2. Select Ok to delete the group.

The group name disappears from the Groups scrolling list.

Setting up the Backup Schedules

NetWorker will automatically back 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.

You create backup schedules for the clients in the Schedules window. To open the Schedules window, select Schedules from the Administration pull-down menu.



Schedules scrolling list – displays the names of the schedules for automatic, network-wide backups known to the NetWorker server. Five pre-configured backup schedules are provided for your convenience. You may not delete the schedule named “Default.” The five pre-configured schedules are described in section “Backup Schedule Examples” of Appendix B.

Name field – displays the name of the currently selected schedule.

Period choices – let 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 will appear in all the calendar weeks. If you select *Month* and set the schedule for one month, the same schedule will apply to all the calendar months.

Previous Month and Next Month buttons – scroll to the previous and next months, one month at a time. Selecting the *Current Month* button displays the current month and next month.

The 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.

To create a new schedule, you name the schedule and describe it with the choices in the Schedules window. Then you apply the choices to the schedule. You can then apply the schedule to a client in the Clients window.

Before learning how to use this window to set up a schedule, you might want to decide what schedule best fits the needs of your network and also learn about the backup levels that NetWorker supports.

The Logic Behind the Scheduling Schemes

It is very easy to set up a schedule using NetWorker. The hard part is deciding on a schedule that best fits your requirements.

When creating backup schedules, you need to consider the following:

- 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 8mm tape cartridge is about 5 gigabytes, and the maximum transfer rate is around 400 kilobytes 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 gigabytes of data. To back up a network with a lot of data, you may use more than one schedule to stagger the full backups over several days.

Think about how many backup volumes you want to keep – this number will depend on how often the data changes, and how long you want to keep the on-line backups.

If you run only incremental backups every night, you will need more backup volumes to fully recover from a disk crash. If a site has 10 gigabytes of data, and 5% of all the data is modified each day, that means 500 megabytes of data need to be backed up every day. At 400 kilobytes per second, 500 megabytes will take

about 25 minutes to back up, and fill up about one tenth of an 8mm cartridge tape. If you are maintaining the backups for three months, you will have about 12 tapes of backups that you need to keep on your shelves.

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 will have to maintain *all* the backup volumes for the three-month period. On the other hand, if the users only expect to be able to recover 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.

Understanding the Backup Levels


NetWorker supports four kinds of 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. 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 on-line, you may 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.

The diagram below illustrates how backup levels work.



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 4 backup volumes. To recover from a disk crash, you will need all four of them – the one with the full backup (day1), and all the volumes with incremental backups.

On day 5, the level 8 backs up everything that has changed *since the full backup*. You no longer need the data on the backup volumes from day 2, 3, or 4. To do a full recovery, 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 only need 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.

Using the Pre-configured Schedule

NetWorker provides a pre-configured backup schedule named "Default" which you are not allowed to delete.

The Default schedule repeats on a weekly basis. Every week, a full backup takes place on Sunday, followed by an incremental backup every day of the week. This schedule is useful for a small-to-medium sized network where the scheduled backups fit onto one backup volume.

In addition to the Default schedule, NetWorker is shipped with four other pre-configured backup schedules which you can use, customize, or delete. These schedules are documented in the section "Backup Schedule Examples" of Appendix B.

If you have a large network with multiple file servers and want to stagger the full backups, or do unattended, scheduled backups during the night, you may do so by using multiple backup schedules.

Suggestions for Easier Operation

There are two ways that the total amount of data you are backing up can be reduced, possibly eliminating the need to change backup volumes even on the days when a full backup is performed.

First, NetWorker offers you the ability to compress files before they are backed up to a backup volume. 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 gigabytes may fit onto a single backup volume.

Second, NetWorker allows you to omit certain files from the backup entirely. For example, *core* files and “.o” files can often be skipped. The NetWorker *skip* directive provides an easy way to specify that files such as these be skipped. (See “Using Directives” in this chapter for more information.)

Finally, the NetWorker server can be configured with concurrent devices or with a jukebox for use during backups. Using these optional features allows you to perform unattended backups of large networks.

3

Creating a New Backup Schedule

This section illustrates setting up schedules with two examples. Read through the examples before you try to set up any new schedules on your own. There are four basic steps to set up a schedule:

1. Choose a name for the schedule.
2. Choose a period.
3. Choose backup levels for each day in the period.
4. Set the overrides, if any.

Note that when you set up one month in the Schedules window, that setup applies to all calendar months, and carries over from year to year. In the same way, if you set up a weekly schedule, it applies for all calendar weeks. This allows you to plan all the

weeks or months in one session, without having to repeat the process every week or month.

Important ⇨ You can make changes to an existing schedule, but you may not change its *name*. If you want to change the name of a schedule, you must delete it and re-create it, giving it a new name.

Example 1 – A Weekly Schedule

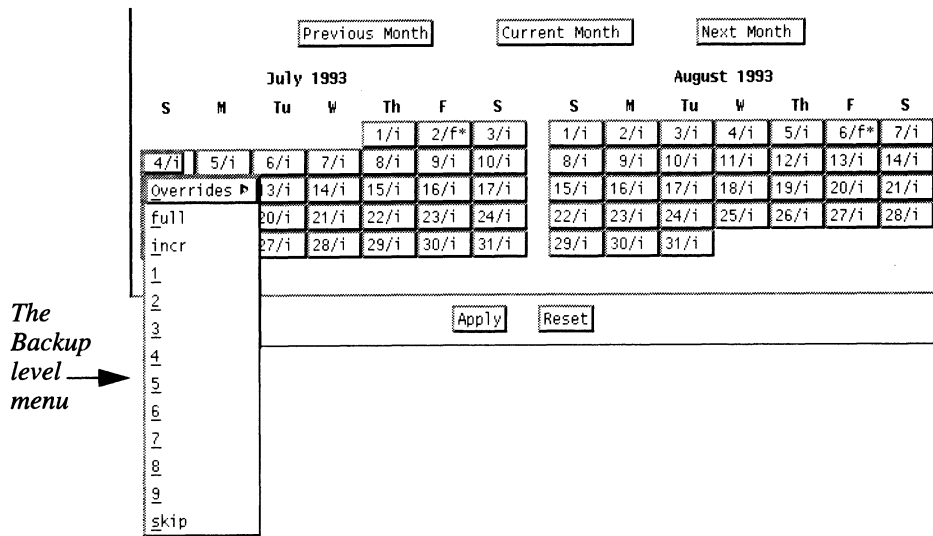
This schedule is suitable for NetWorker clients with files that change often. You want to do a full backup every Monday because a weekly full backup will make recovery from a disaster easier. Also, someone will be available to change backup volumes and monitor the server on Mondays. The other days of 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. Press the Create button.
3. Enter “Monday Full” into the Name field.
4. Select the Week button, if it is not already selected.

The calendars show a full backup on Sundays, followed by incremental backups all the other days of the month.

- In the current month calendar, select any Sunday and pull down the Backup Level menu, like this:



- 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:

- In the current month calendar, select any Monday, and pull down the Backup level menu.
- Select full from the menu.
- Press Apply to save the schedule.

NOTE: If you see the error message "user *user_name* needs to be on administrator's list," it means you do not have permission to make configuration changes. See "Adding or

Changing Administrators" in this chapter for more information.

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

Scroll through all the months, and notice how 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 will back up all the files that have changed since the full backup. All the other days of the month, an incremental backup will protect any daily changes. You will 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 May 15, because it is a company holiday, and force the backup on May 14 instead.

Follow these steps:

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

Notice the months are already set up to do a full backup on the first day.

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

1. In the current month, select the fifteenth day (15/i).
2. Select "5" from the Backup level menu.

Scroll through the months, and notice how the monthly schedule is carried through every month – a full backup on the first day, a level 5 on the fifteenth day, and incremental backups on the other days of the month.



Now you are going to override the scheduled level 5 backup on May 15 with a level "skip," to skip the backup for that day, and schedule the level 5 backup on May 14. Follow these steps:

1. Use the Previous Month or Next Month buttons to display the May calendar.
2. In the May calendar, select the fifteenth day.
3. Pull down the Backup level menu, and select Overrides.



3

- Slide the cursor to the right, to display the Overrides menu.



- Select skip* from the Overrides menu.
- Select the fourteenth day in the May calendar.
- Select 5* from the Overrides menu.
- Press Apply to save the schedule.

Notice the date on May 15 displays an "s*," indicating you have forced a "skip backup" instruction for that day. The date on May 14 displays a "5*" to indicate a level 5 backup will take place instead of the incremental. Any override must be scheduled explicitly each time. In other words, an override does not automatically repeat in a monthly or weekly period.

You may set up as many NetWorker backup schedules as you like, as long as you give each one a unique name. To assign a

client to a particular backup schedule, you must go to the Clients window. See "Configuring the Server's Clients" in this chapter for more information.

Summary

To create a backup schedule, follow these steps:

1. Choose Schedules from the Administration pull-down menu in the Main window.
2. In the Schedules window, press 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, press the Apply button to create the new schedule.

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 pre-configured label templates, or create a custom label template using the Label Templates window.

This section describes how to use label templates, gives a description of the pre-configured templates, describes how to create a custom label template, and provides examples of label templates for your use.

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.

Using the Pre-configured Templates

The pre-configured label templates shipped with NetWorker are: Archive, Default, Full, NonFull, Offsite, and Two Sided. These are provided so that you can easily start labeling your backup volumes. There are also pre-configured volume pools with corresponding names. The pre-configured volume pools will automatically use the pre-configured label template by the same name.

The number range for all of the pre-configured label templates starts at 001 and ends with 999 to allow for expansion of the volume pools.

The Archive label template is used only for clients that need to archive data. It has three fields each separated with a period. The first field contains the name of the NetWorker

server, the second field is "archive," and the third field contains a number.

For example:

```
server.archive.number  
  
space.archive.001  
space.archive.099  
atlas.archive.325
```

The Default template has two fields separated by a period. The first field contains the name of the NetWorker server and the second field contains a number.

For example:

```
server.number  
  
space.675  
space.800  
atlas.054
```

The three pre-configured label templates Full, NonFull, and Offsite use the same labeling conventions. The name of the label template appears in the first field, and the second field contains a number.

For example:

```
label name.number  
  
Full.076  
NonFull.003  
Offsite.120
```

The Two Sided template is for use with two-sided media such as optical media. When labeling two-sided media you need to be able to label both sides of the media. The first field contains the name of the server, the second a number, and the third field will either contain an "a" or "b" to differentiate between the two sides

of the media. See “Example 3 – Two-sided Media Label Template” in this chapter for an example of how to create your own two-sided media label.

```
server.number.side
```

```
phoenix.001.a
```

```
phoenix.001.b
```

Choosing a Label Template

Before you choose a label template for labeling your volumes, you must first choose the device that has tapes mounted for labeling from the Devices display in the Main window. Next, choose a label template in the Pools window, then start the labeling process by selecting Label from the Operation pull-down menu.

Select a label template from the Label template display in the Pools window.

Important ⇨ Whether you use a pre-configured 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.

Creating a Custom Label Template

To create and add a label template, open the Label Templates window by selecting Label Templates from the Administration pull-down menu.

Tip ⓘ If your label template matches your backup media storage system, it will help you organize and locate the media in

storage. For example, if you store backup media in bins or racks, you should include a place for the bin or rack number in the label template. (See "Example 1 – A Storage-oriented Template" in the next section.)

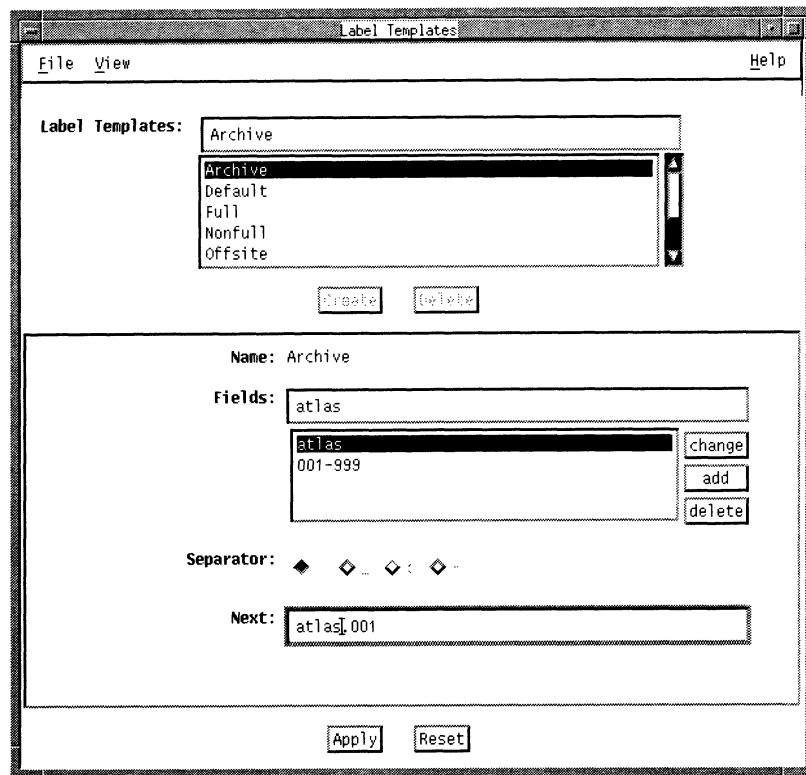
It is also helpful if the label template name matches the name of the corresponding pool.

For example:

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

3

The Label Templates window is shown below:



3

The Label Templates scrolling list displays the label templates known to the NetWorker server. The pre-configured label templates and any new templates you create appear in this list.

The Name field displays the name of the currently selected label template.

The Fields field displays the components of the label template. The components appear as one component per line in the display. There are four types of components:

- ❑ range of numbers – for example, *000-999*
- ❑ range of lowercase letters – *aa-zz*

- ❑ range of uppercase letters – AA-ZZ
- ❑ character string – for example, *Accounting, 001*

Each range includes a starting value, a dash (-), and an ending value. The starting value and the ending value must have the same number of characters – for example, use “aaa-zzz,” not “aa-zzz.” (This rule does not apply to a list of strings or words.) You may have as many components as you want, but it is best to keep the template simple, with few components. The total length of the label must be less than 64 characters.

Each component of the label template may be separated by a period, dash, colon, or underscore, represented by the Separator buttons. If none of the buttons are selected, the components will have no separators, like this: “AA00aa.”


The Next field displays the next label NetWorker will use to label the backup volume according to the template. If you want to force a label or start the labeling scheme at a certain point, you may enter the starting label into the Next field. NetWorker continues to generate labels from that point on, according to the rules of the template.

Example 1 – A Storage-oriented Template

In this example, assume your backup media are magnetic tapes 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.

NetWorker does not care how the volumes are labeled, as long as each one has a unique name. It tracks the backups and maps them to the backup volumes, so you do not have to remember which backups are on which volumes. When NetWorker needs a backup volume to fill a recovery request, it requests it by name in the Pending display of the Main window.

In this example, 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, left-to-right.



Using this labeling scheme, a tape labeled "2-3-54" will be stored in the second rack, on the third shelf, and will be 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 may 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.
2. Press 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" (do not enter the quotes).
5. Press the add button.

Now add the second and third components of the label template:


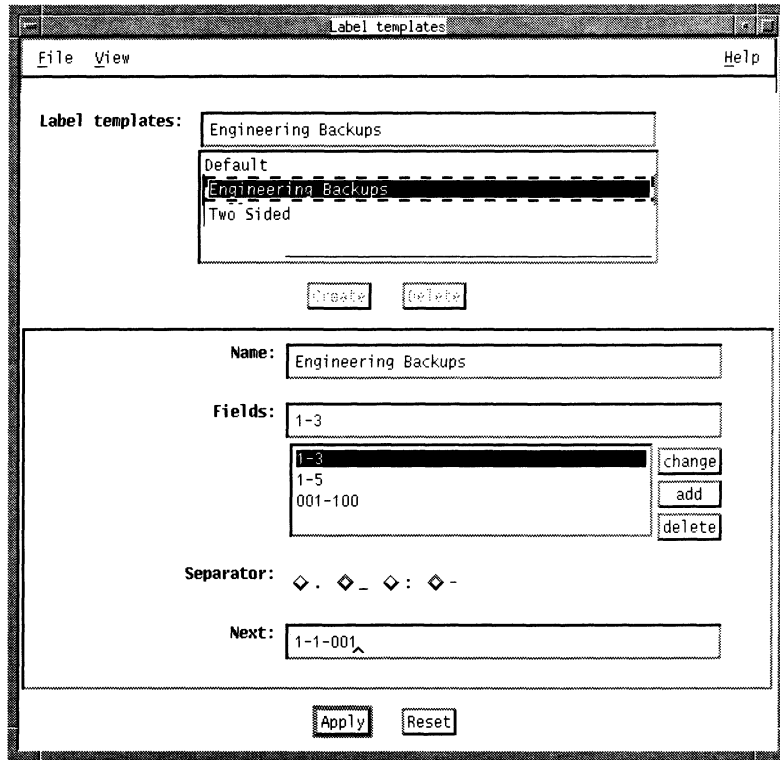
1. Enter the second component of the label template into Fields. This field holds the shelf number, so enter "1-5" into the field.
2. Press the add button.
3. Now enter the third component of the label template, the tape position numbers. Enter "001-100" into Fields.

NOTE: You must enter the same number of characters for any numbers in the label template. For example, use:

01-20	not	1-20
1-9	not	01-9

4. Press the add button.
5. Choose the separator for each component of the label template, by pressing the dash (-) Separator button.
6. Press the Apply button to apply the new template to the NetWorker server.

The Label Templates window now looks similar to this:



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 will generate *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 will generate 67,600 different labels (26^2 times 10^2). To give you an idea of how NetWorker will generate the labels, below is a sample of the labels:

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.
2. Press 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. Press the add button.

Now add the second component of the label template:

1. Enter the second component of the label template into Fields. This field contains double digits, so enter "00-99" into the field.
2. Press the add button.
3. Choose the separator for each component of the label template, by pressing the period (.) Separator button.
4. Press 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.

3

Using Volume Pools

This section describes volume pools, how to use pre-configured pools, how to create new pools, and provides several examples of why volume pools are useful.

Volume pools is 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. All NetWorker volumes belong to a pool – either the Default pool or a pool you have chosen.

If a backup is not sorted into either a pre-configured or a pool you have created, NetWorker will automatically send that backup data to the Default 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.

Leaving a choice for Groups, Clients, Save sets or Levels un-selected in the Pools window means that NetWorker will not sort the data by that choice. For example, if you do not choose a level from the Levels choices, NetWorker will not sort the data by levels for that pool. NetWorker will sort the data based on the other selections for that pool.

For each pool you must select either a group, client, or save set. You must have data to sort from one of these choices for a pool, or NetWorker will not allow you to create the pool.

For example, you can sort by:

- a backup group
- a backup client
- a 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:

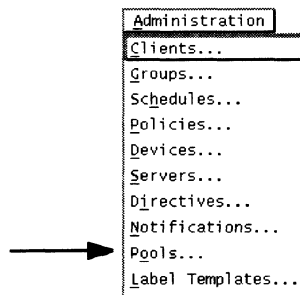
- choosing a special label template
- sending all the backup data to a particular backup device
- choosing to store entries for the files in the on-line index (or not, if you are sending the backup media to an off-site vault)

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.

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

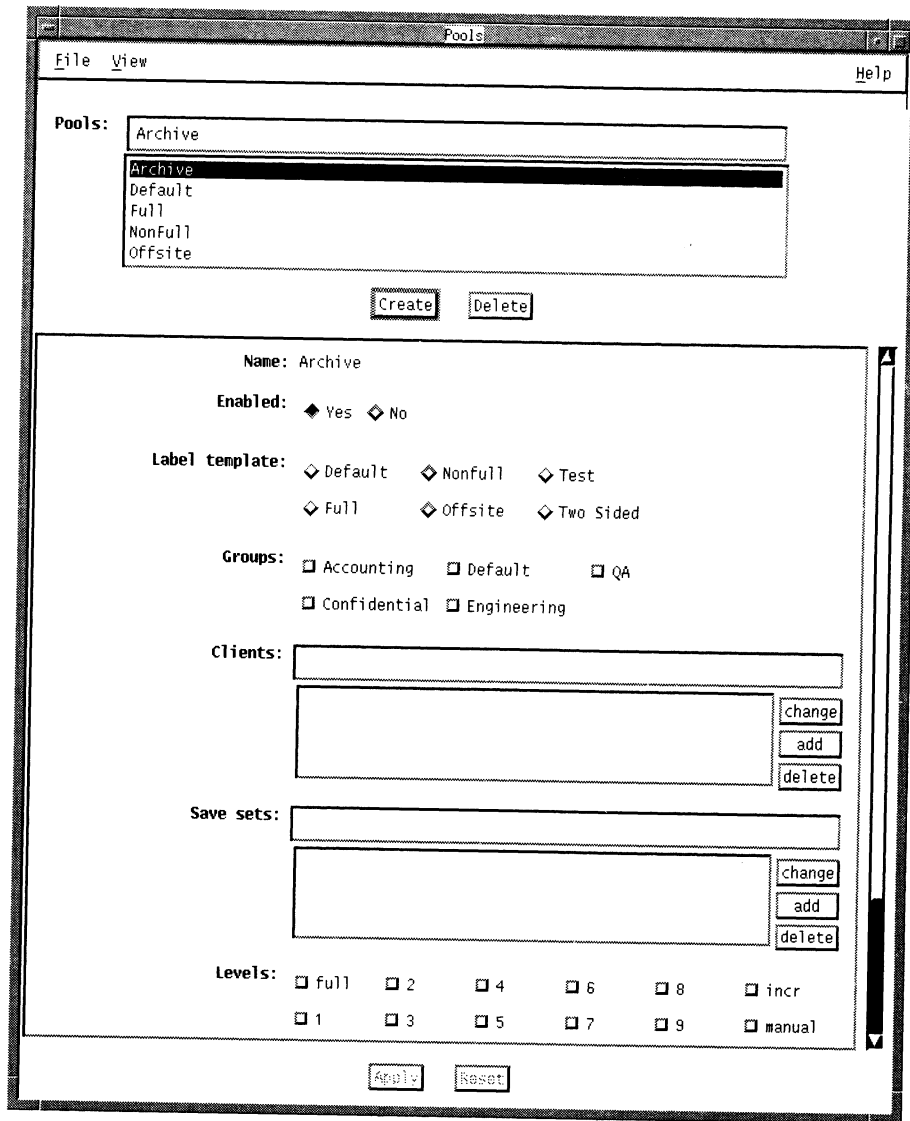
Using the Pools window, it is easy to create a new volume pool. However, you cannot create a new pool if NetWorker is busy doing a backup – this ensures that NetWorker will send the backup data to the correct pool. Nor can you delete a volume pool if any backup volumes still belong to that pool.

Each pool has an accompanying label template which you first create in the Label Templates window. You should use the same name for both your label template and volume pool to easily keep track of the backups and to which pools or tapes they belong. If you do not select a label template in the Pools window while creating a pool, NetWorker will automatically create a template for you with the same name as the pool. The advantage to creating your own label template in the Label Templates window is that you have more control over its design and content. To open the Pools window, select the Pools command from the Administration pull-down menu.



The Pools window appears, as shown below.

3



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 how to use them for modifying and creating custom pools.

The Pools field displays the pool currently selected in the pools scrolling list. You may enter the name of the pool in the Pools field or select it from the scrolling list. A scroll bar will appear to the right of the scrolling list if the list of pools extends beyond the display area.

Use the Create button to add a new pool.

Use the Delete button to eliminate a pool. (You cannot delete a pool as long as it still contains backup volumes.)

The Name field displays the name of the currently selected pool. It is also where you enter the name of a new pool after selecting the Create button.

If you want NetWorker to use a pool during the backup process, choose Yes from the Enabled choices. If you want to exclude a pool during the backup process, select No.

Select a label template from the Label template choices. Each new label template you create in the Label Templates window will appear as a choice. It is highly recommended that you use the same name for both the label template and the selected pool.

The Groups choices 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.

You can use the Clients field to:

- add an individual client(s) to a pool without selecting a group
- pick a specific client belonging to a group. Note that if you pick a specific NetWorker client from a group you will 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.

Use the Save sets field to:

- send a specific filesystem to a pool (see "Example 4 – Specific Filesystems" in this chapter for an example)
- send a specific filesystem, from a specific client, to a volume pool without selecting a group (see "Example 2 – Secure Device" in this chapter for an example)
- send a specific filesystem, from a specific client of a group, to a volume pool (see "Example 1 – Separate Departments" in this chapter for an example)

Use the scroll bar to view the remainder of the choices in the Pools window. The choices at the bottom of the Pools window are shown below:

Archive only: Yes No
Devices: /dev/rmt/0n
Store index entries: Yes No

The Archive only choices apply only if the NetWorker server is backing up data for archival. Typically clients that archive data are DOS, OS/2, and NetWare clients. If the volume pool is for archiving your clients' data, select Yes. Usually, the Archive only

choice is No. (Refer to the *Legato NetWorker Administrator's Guide, NetWare Version*, for information about archiving data.)

The **Devices** choices display the devices recognized by NetWorker. This is where you choose which 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 media devices on the same server. If you do not choose a devices, NetWorker will use any device.

Store file index entries allows you to choose, with the **Yes** and **No** choices, whether or not you want to include the index entries from the backup in the on-line index. If you choose **No** the index entries will not be included in the on-line index. For example, if you are using the pre-configured pool **Offsite**, you may not want the index entries on-line. However, an entry for the save set is still made in the media database.

Use the **Apply** button to apply all of the selections you have chosen for a particular pool.

Use the **Reset** button to restore your original selections, and override any new selections you may have chosen.

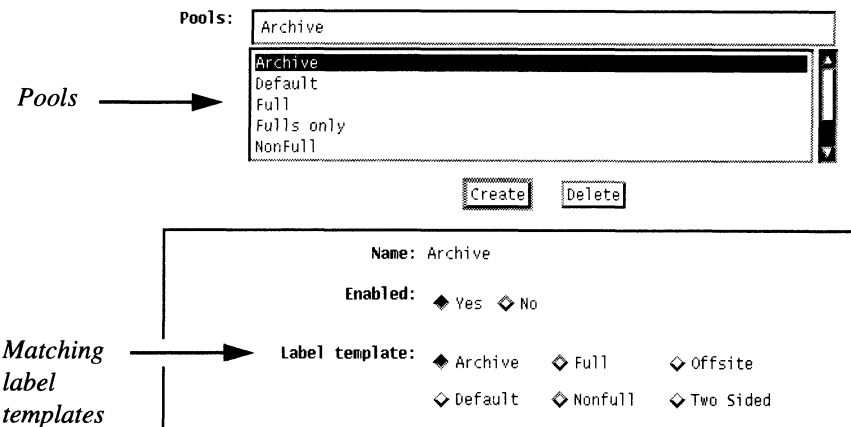
Use the **Pools** window to choose the pre-configured pools, create new pools, and change existing pools. Use this window for all administrative tasks relating to volume pools.

Using the Pre-configured Pools

NetWorker is shipped with five pre-configured pools and matching label templates. Each pre-configured volume pool has a set of unique pre-selected choices. If you do not choose a pool for your backups they will automatically be assigned to the pre-configured **Default** pool and will be labeled using the **Default** label template.

The pre-configured pools have been included for your convenience and provide a variety of ways for organizing your data.

The five pre-configured volume pools and their matching label templates are shown below. (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 and Archive pools without making any selections. To use the other pre-configured 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 five pre-configured pools are:

- ❑ **Archive** – for archiving client data only. This pool cannot be modified or deleted. The pre-configured settings are: Enabled – Yes, Label template – Archive, Archive only – Yes, Store Index entries – Yes. There are no selections for you to make for this pool.
- ❑ **Default** – automatically used if you do not choose a pool. If you decide not to use the pools feature, NetWorker will

automatically place all of your backup volumes in this pool. The Default pool cannot be deleted or modified. The pre-configured settings are: Enabled – Yes, Label template – Default, Archive only – No, Store Index entries – Yes. There are no selections for you to make for this pool.

The Full, NonFull, and Offsite pools are intended for sorting data by levels.

- ❑ Full – use this pool for full backups only. This pool separates all of your full backups from the incremental and level backups. Using the Full pool provides you with the ability to easily track and separate your full backups from the incremental and level backups. Typically you use this pool in conjunction with the NonFull pool. The pre-configured settings are: Enabled – No, Label template – Full, Levels – full, Archive only – No, Store Index entries – Yes.
- ❑ NonFull – use for any backups other than full backups. This pool includes all incremental and level backups. Use the NonFull pool to easily keep your incremental and level backups separate from the fulls. Typically you use this pool in conjunction with the Full pool. The pre-configured settings are: Enabled – No, Label template – NonFull, Levels – all level and incremental backups, Archive only – No, Store Index entries – Yes.
- ❑ Offsite – for volumes being stored off-site. The Offsite pool allows you to easily create a set of volumes to be stored off-site. If your on-site backup volumes are destroyed, you can still recover your valuable data with the volumes you have stored off-site. 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 Off-site pool. The pre-configured settings are: Enabled – No, Label template – Offsite, Levels – full, Archive only – No, Store Index entries – No.

Important ⇨ Remember to enable the pools you wish to have in effect during the scheduled backups by selecting Yes from the Enabled choices.

Creating a New Pool

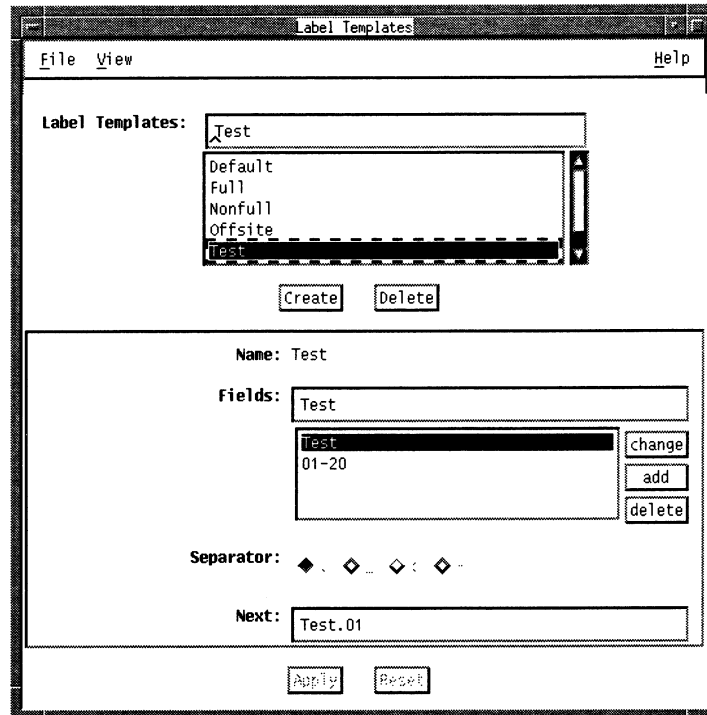
You can create as many pools as you need to effectively organize and sort your backup volumes. There is no limitation on the number of new pools you can create. This section uses an example to show you how to create a new pool.

3

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.

All NetWorker pools require a label template, described earlier in this chapter in the section "Using Label Templates." Before creating a new pool you should first create a label template with the same name, in this example "Test." It is strongly recommend that you keep the pool names and their corresponding label templates consistent. If you do not use the exact name, you should at least use names that have a logical connection.

The Label Templates window will look like this after you create the Test label template:



3

After creating the label template, open the Pools window and follow these steps to create a pool named "Test:"


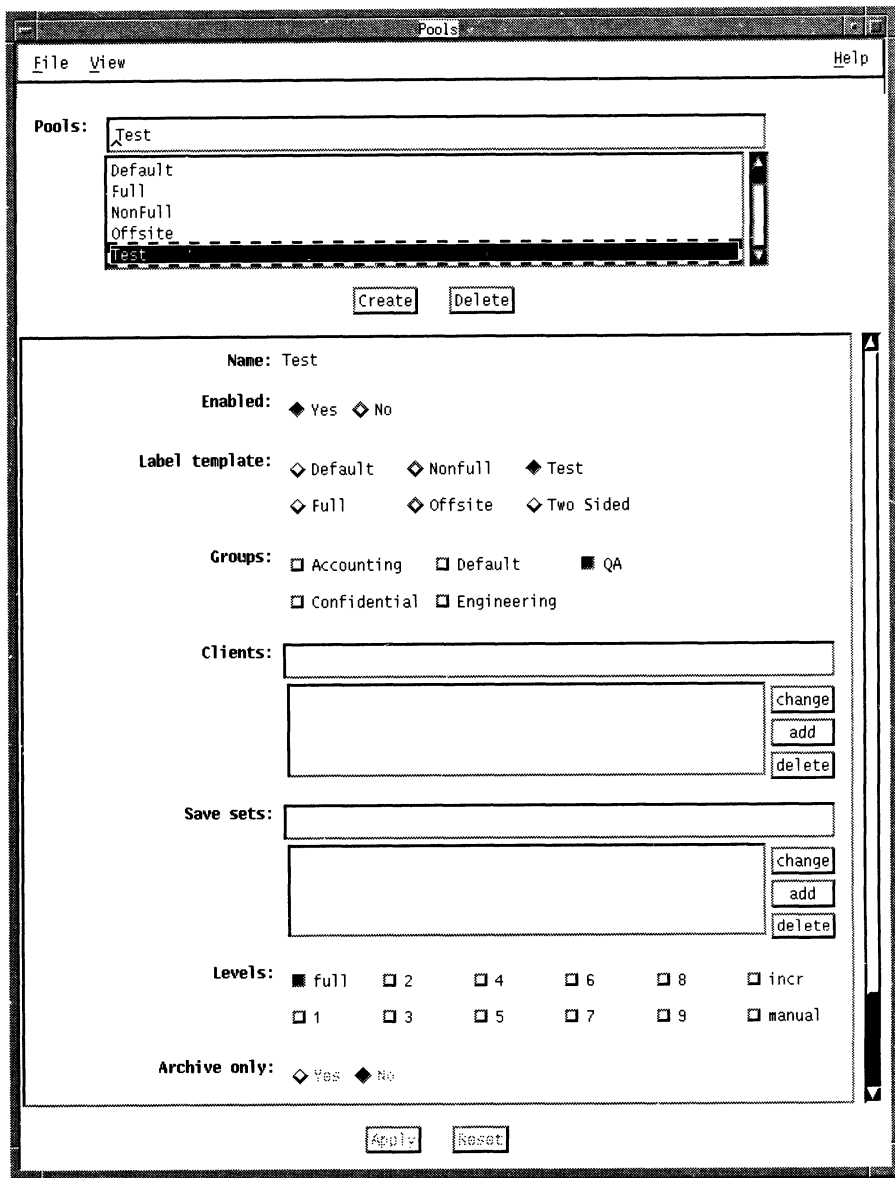
1. Press the Create button and enter "Test" into the Name field.
2. Select Yes from the Enabled choices. (Yes is the pre-configured choice.)
3. Choose the label template Test, which you previously created.
4. Choose the group QA, because it contains the clients you want to include in the Test pool.

5. You do not want to selectively include clients or save sets, so you do not need to use either the Clients or Save sets fields.
6. Select full from the levels choices.
7. Leave the Archive only choice as No.
8. If your NetWorker server has more than one device, you may select one of them from the Devices choices.
9. Select No from the Store index entries display. Selecting No means that you will not be including the index entries in the on-line index, saving space on your hard disk.
10. Press the Apply button to apply the selections.

**Important** ⇨

NetWorker will not allow you to create two different pools with duplicate choices for Groups, Clients, Save sets, Levels, or Archive Only in the Pools window. If you try to create two pools with the same choices for all of these items NetWorker will display an error message.

The Pools window will look like this after creating the pool "Test."



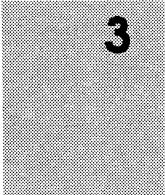
Once you create a new pool you can easily modify it by choosing the pool from the Pools scrolling list and changing the current

selections. When you have completed making the modifications press the Apply button. If you have changed your mind and do not want to implement the new selections press the Reset button to restore the Pools window to its original configuration.

Pool Examples

This section provides some 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, first create a label template with the name of the department in the Label Templates window, use the Pools window to create a new pool with the same name as the label template, and 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, but just a select set of files from a small number of clients in a pool named "Confidential."

To create the "Confidential" pool first create a label template named "Confidential," create a pool with the same name, then enter the name of each client in the Clients field, and enter the filesystems for each of the clients in the Save sets field. Finally select the device they will be backed up to from the Devices choices.

Example 3 – Targeted Backups

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

First create a label template named "Graphics Full," create a new pool by the same name, choose the group(s) that need to be included, select full from the Levels choices, then choose the jukebox device from the Devices choices where the graphic files will be backed up automatically.

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 will create a pool named "E-mail."

To create a pool named "E-mail" first create a label template named "E-mail," then create a pool with the same name. You do not need to select a specific group or client because you are backing up all of the client systems' electronic mail. In the Save sets field you would enter the following:

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

NetWorker will back up all the electronic mail from all the clients to the volumes in the "e-mail" pool.

Creating Index Policies

This section describes how to create policies that automatically manage the on-line indexes. You can also manage the indexes manually using the Indexes and Volumes windows. Manual index management and index management concepts are described in detail in Chapter 4, "Managing the On-line Indexes."

3


When NetWorker starts a backup, it creates entries for the saved files in the on-line indexes. NetWorker maintains two types of indexes: 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.

Every backup adds entries to the indexes, so the indexes require management in order to control their size and contents. NetWorker automatically manages the contents of the indexes with policies that you specify for each client. Automatic index management with policies preserves valuable disk space by removing old entries from the indexes. You may use pre-configured policies or customize your own using the Policies window.

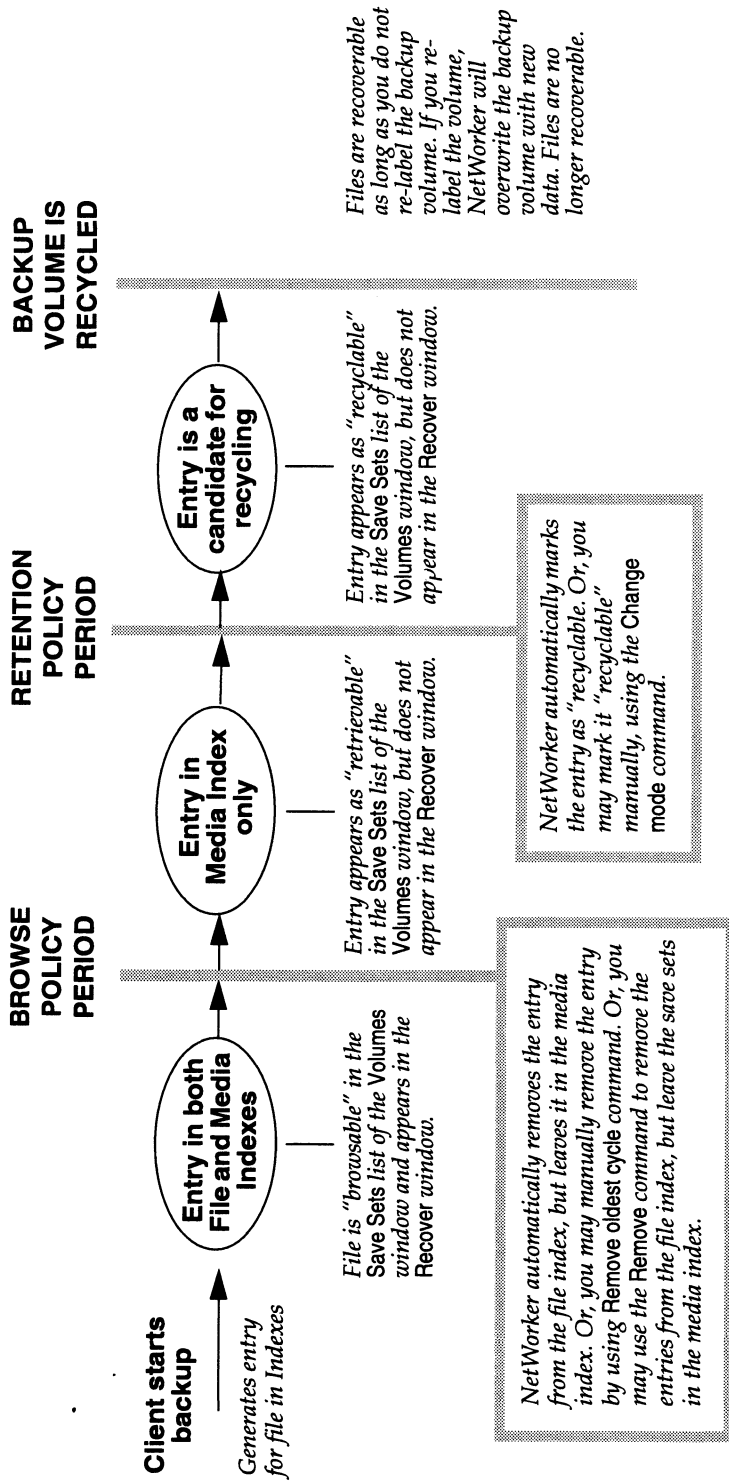
Each client requires two separate policies: a Browse policy and a Retention policy. The Browse policy determines how long entries for your files will remain in the on-line file index and thus be browsable in the Recover window. Entries older than the Browse policy are automatically removed from the on-line file index. The Retention policy determines how long entries will be retained in the

media index and thus be recoverable. Entries older than the Retention policy are marked as recyclable in the media index.

The diagram on the next page illustrates the interaction between an entry in the indexes and the index policies.



Index Policies and Index Management



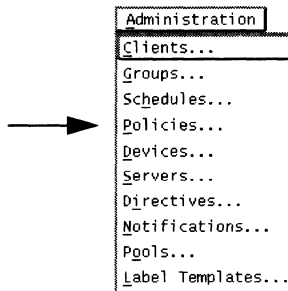
To recover a file which is "retrievable" or "recyclable," but not "browsable," you must:

1. Find the save set that includes your file in the Save Sets list. You may have to look through all your backup volumes for the save set with the date that you want.
2. You can use the Recover Index command or the Scanner command from the command line to recover the data from the backup volume and re-populate the file index. Your file will now be "browsable" in the Recover window.

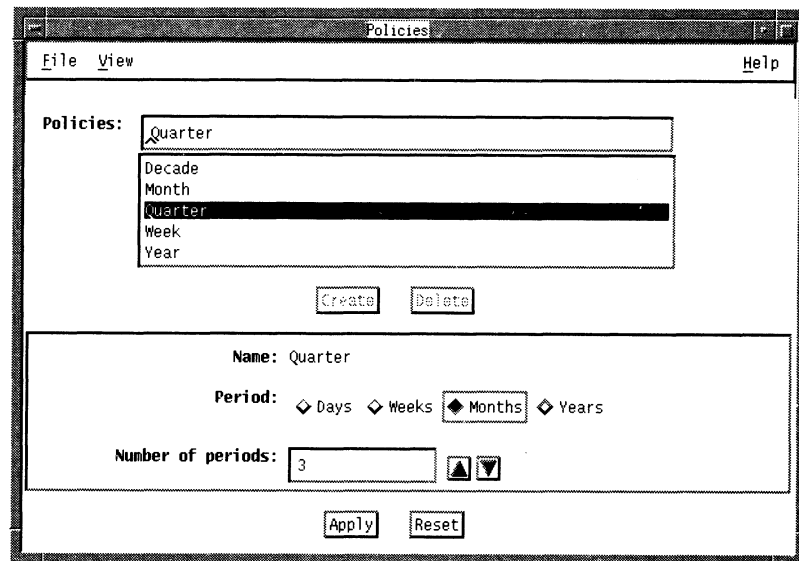
If you create new policies for automatic index management using the Policies window, the policies appear as selections for both the Browse and Retention policies in the Clients window. Every client requires both a Browse and Retention policy. When you select Browse and Retention policies in the Clients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

Periodically, NetWorker compares the backup date of the entries in the indexes with the time period specified by the policies. If a file entry has passed the time period specified by the browse policy, it is automatically removed from the file index, thus freeing up space on the disk. When a media entry has passed the time period specified by the retention policy it will be marked as recyclable in the media index. The entry still needs to be removed from the media index.

To create policies for automatic index management, open the Policies window by selecting Policies from the Administration pull-down menu.



The Policies window is shown below:



The Policies scrolling list displays the names of all policies known to the current NetWorker server. If you have not created any new policies, the scrolling list will display the following five pre-configured policies: Decade, Month, Quarter, Week, Year.

When a policy is highlighted in the scrolling list, its name, period, and number of periods appear in the lower section of the window.

To create a new policy, follow these steps:

1. Press 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. Press the Apply button to finish the process. After you press Apply, the newly created policy appears in the Policies window scrolling list and also in the Clients window as both a Browse policy and a Retention policy.

If you make a mistake, press the Reset button to reset the window to where it was before the last Apply.

Enter any name for the policy that you want, but you may want 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 on-line file index for three months.

Follow these steps to delete a policy:

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

Tip 

You are not allowed to delete policies that are currently in use or the pre-configured policies. The pre-configured policies are Quarter and Year.

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. Press the Apply button to change the policy.

Important 

A policy cannot be changed so that the client Browse policy exceeds its Retention policy. This means that an entry must be

removed from the file index before it can be marked recyclable or removed from the media index.

Adding and Deleting Backup Devices


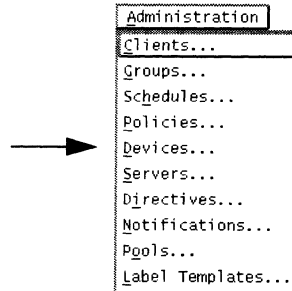
3

With some versions, the NetWorker installation script requires that you enter at least one device *pathname* for a NetWorker backup device. You may add more devices, since 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 it is full and then uses the backup volume in another device until it is full, and so on.

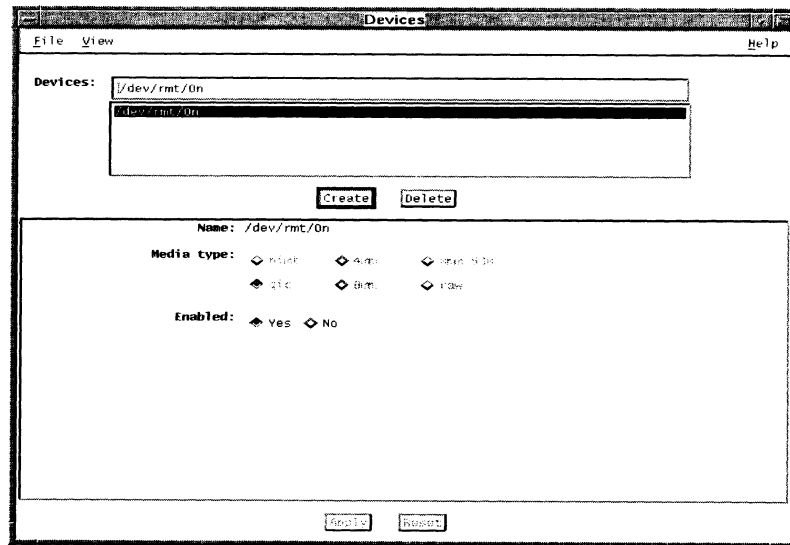
NetWorker supports the following types of backup devices and their corresponding backup media:

- half-inch magnetic tape drives
- quarter-inch cartridge tape drives
- 4mm (DAT) drives
- 8mm tape drives
- 8mm 5 Gigabyte tape drives
- raw disk drives
- the CLARiiON™ tape array (contains 4mm DAT drives)

To view current devices or to add a new device, select Devices from the Administration pull-down menu.



The Devices window appears, similar to the one shown below.



The device *pathnames* displayed in the Devices scrolling list represent the backup devices connected to the NetWorker server. A pathname may indicate either a stand-alone device or

a CLARiiON tape array. These devices can be used for backing up and recovering client files.

The Name field displays the *pathname* of the currently selected device or accepts a *pathname* of a new device.

The Media type of the currently selected device is highlighted. The Media type choices are also used to select the type of new devices. The choices represent half-inch magnetic tape drive (himt), quarter-inch cartridge tape drive (qic), four millimeter tape drive (4mm), eight millimeter tape drive (8mm), eight millimeter five gigabyte tape drive (8mm 5GB), and disk. These buttons are inactive until you press 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.

To add a device, you provide the following information:

- ❑ the *pathname* of the device

The *pathname* of a no-rewind tape device follows the syntax */dev/rmt/un*, where *u* is the unit number; *n* indicates no-rewind.

Important ⇨

A no-rewind tape backup device is required 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 will be overwritten by the next backup.

- ❑ 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. Select Devices from the Administration pull-down menu.

2. Press the Create button.
3. Enter a new device *pathname* into the Name field.
4. Select one Media type button for the type of device.
5. Press Apply.

Name:

Media type: hint 4mm 8mm 5GB
 qic 8mm raw

3

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

1. Select the device you want to delete in the Devices scrolling list.
2. Press the Delete button.

An “ok to delete?” confirmation box appears.

3. Select Ok to delete the device.

Any additions or deletions you make will appear in the Devices window.

Selecting Server Options

This section contains information about server options and includes: changing server parallelism, using concurrent devices, and adding or changing NetWorker administrators.

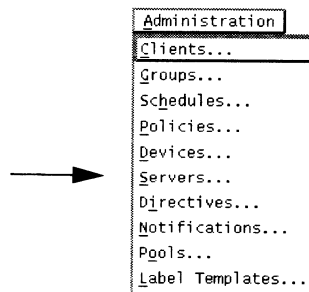
Changing Server Parallelism

A NetWorker server may back up files from many clients in parallel and mix the files onto the backup volume. This feature keeps a steady stream of files supplied to the server backup devices so that they operate at high speed.

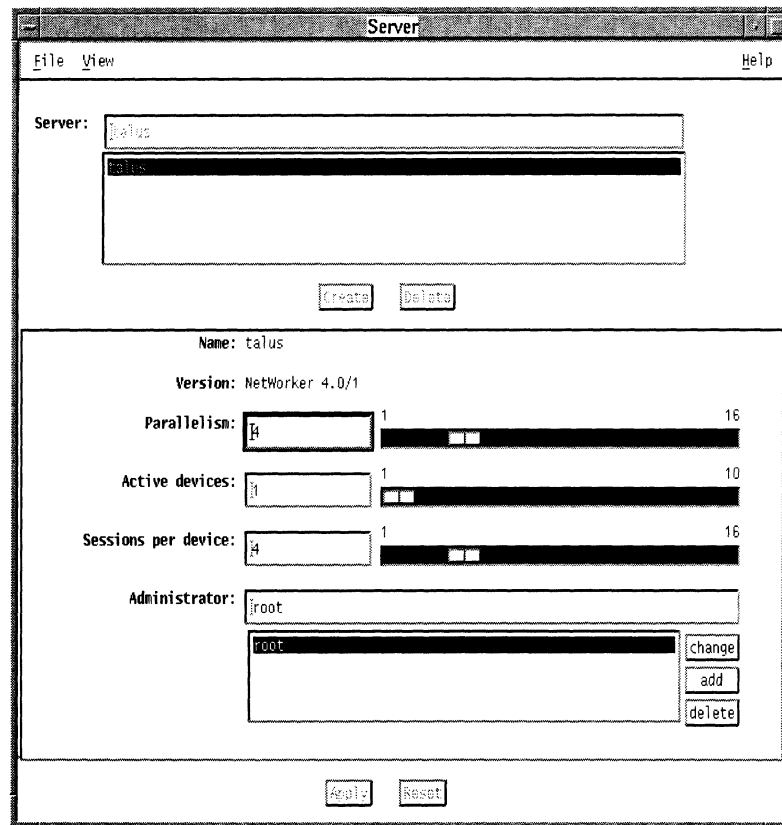
To optimize the efficiency of the server, you can control how many clients can back up their files at one time by changing the value in the Parallelism field in the Server window. You may want to adjust the number of clients allowed to back up at a time if:

- the server is taking too long to back up all the clients (increase the parallelism)
- the server is so busy backing up clients it is not available for any other tasks (decrease the parallelism)
- the server runs out of swap space or memory

To open the Server window, select Servers from the Administration pull-down menu.



The Server window is shown below:



3

The Name field displays the name of the current NetWorker server.

The Version field displays the version of the software currently installed on the NetWorker server.

The Parallelism value is controlled by a sliding bar. You can set the parallelism to any value between 1 and 16, by using the mouse to slide the bar to the left and right.

The Administrator scrolling list displays the names of the users or user groups who have permission to change the configuration of the NetWorker server. For more information, see “Adding or Changing Administrators” in this chapter.

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

1. Select Servers from the Administration pull-down 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. Press the Apply button.

When a client finishes a backup, another client may start a backup, but the parallelism value will not be exceeded. For example, if there are 8 NetWorker clients and parallelism is set to 4, after the first client finishes a backup, the fifth client will start, so that a total of 4 clients are backing up at one time. Any 4 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 1.

Using Concurrent Devices

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

To back up to concurrent devices, you enter the maximum number of Active devices and the desired number of Sessions per device in the Server window, as shown below.

Name: talus

Version: NetWorker 4.0/1

Parallelism: 1 16

Active devices: 1 10

Sessions per device: 1 16

Administrator:

3

The Active devices and Sessions per device values are controlled by a sliding bar. You can set the number of active devices to any value between 1 and 16. You can set the number of save sessions per device to any number between 1 and 16. A save session is generated when a client starts a backup. A single client can generate multiple overlapping save sessions.

The Active devices value sets the maximum number of media devices NetWorker will use concurrently for backups. An active device is one that is being used by the NetWorker server for backups or recovers

The Sessions per device value is the target number of save sessions accepted by an active device. Sessions per device also determines which device processes the next save session. If a device is receiving the target number of save sessions, then the server sends the next session to the next under-utilized 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.

Together with Parallelism, the Active devices and Sessions per device values control the backup traffic from the NetWorker server to the media devices.

To configure the NetWorker server to back up to concurrent devices:

1. Select Servers from the Administration pull-down menu.
2. Set the Parallelism value. See “Changing Server Parallelism” in this chapter.
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. Press the Apply button.

Tip 

To distribute your client backups evenly across the available backup devices, simply divide the number entered in the Parallelism field by the number in the Active devices field. Then enter that number into the Sessions per device field. For example, if the number you have entered in the Parallelism field is eight, and you have two entered in the Active devices field, the number you will enter in the Sessions per device field will be four.

Adding or Changing Administrators

NetWorker is shipped with a pre-configured settings so that only *root* has permission to change its system configuration. You can add, delete, or change the permission to give other users or user groups system administration privileges. When using

NetWorker for the first time, run it as *root* to add the user or user group names to the Administrator list.


The system administration permissions are modified in the Server window.

To change the setting, follow these steps:

1. Open the Server window by selecting Servers from the Administration pull-down menu.
2. Enter the new user name or NIS (network information service) netgroup name into the Administrator field. 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. Press the add button. (The Change button lets you change the currently selected name.)



4. Press the Apply button to apply the changes.

NOTE: If you get the error message "user *user_name* needs to be on administrator's list," it means you do not have permission to make configuration changes.

To delete a name from the Administrator list:

1. Select the name in the scrolling list.

2. Press the Delete button.
3. Press the Apply button.

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

Using Directives

3

In a typical directory tree, there are many files that do not need backing up. For example, *core* files which are generated when a process terminates abnormally, or object files (".o" files) which are easily re-generated from source. You probably do not want to waste backup volumes or network bandwidth backing up these files.

Or, suppose your NetWorker client is a fast machine, but your network is slow or busy. You may want to *compress* the client's files in order to use less backup volumes and network bandwidth backing up its files.

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

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" will skip 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 pre-configured directives: "Default," "Default with compression," "DOS standard directives," and "NetWare standard directives." Each

set of directives covers 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.

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

Default with compression – 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.

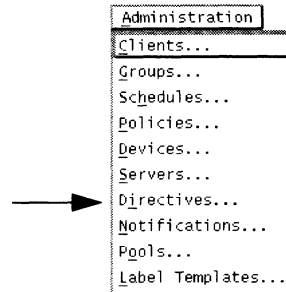
DOS standard directives – use when you are backing up your DOS clients.

NetWare standard directives – use when you are backing up your NetWare clients.

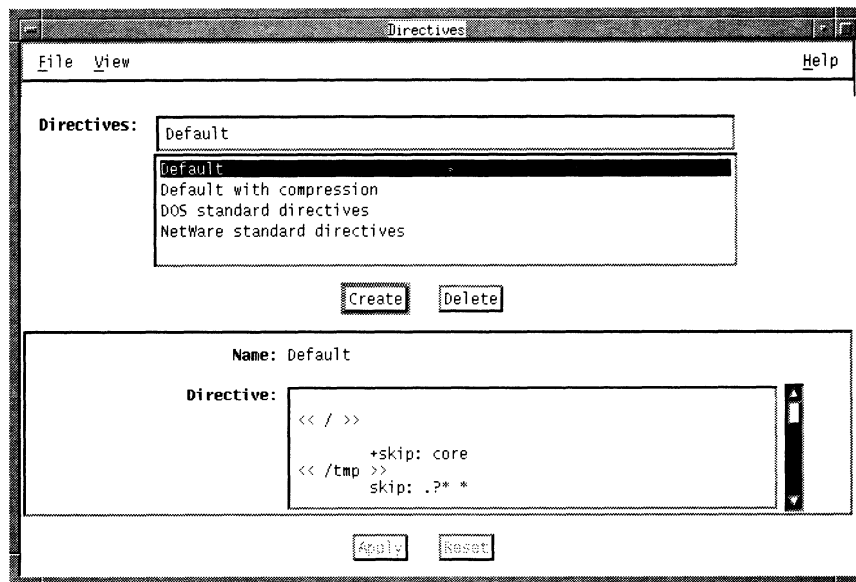
Tip 

NetWorker will work just fine if you use the pre-configured directives without ever changing them. If you are an advanced user, and want to add custom directives, please read this section and refer to the *nsr* man page.

To see the Directives window, select Directives from the Administration pull-down menu.



The Directives window is shown below:

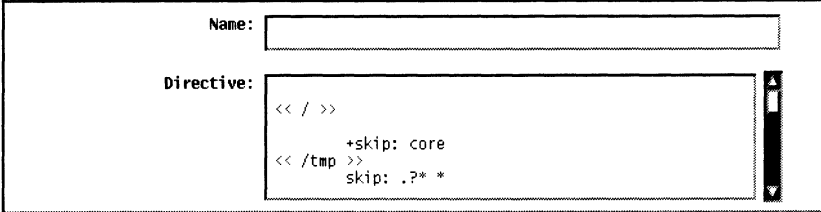


The Directives scrolling list displays the names of the directives known to the NetWorker server. The Directive display shows you the contents of the currently selected set of directives. Each

directive in the list is preceded by the directory to which it applies (for example, `<< /tmp >>`).

To add a set of directives or a single directive instruction, press the Create button.

The bottom half of the window changes, allowing you to change the name and contents of the directive.



Name:

Directive:

```
<< / >>
+skip: core
<< /tmp >>
skip: .* *
```

3

You can type directly into the Directive display – it is a text window. The correct syntax for directives is explained in more detail in the `nsr` man page. After you add a directive or set of directives, press the Apply button to apply the change. To change a directive, press the Change button and edit the window.

Also, if you would like more information about directives, and the algorithm used to apply directives when both server and clients have directives, refer to the `nsr` man page.

Setting Event Notification

As system administrator, you may be interested in knowing about events taking place in the NetWorker environment. You

will see messages on the system console (*/dev/console*) or have them mailed to you electronically.


There are three kinds of notification that are of particular interest to you:

- did the network-wide backup complete for all the clients scheduled for the backup?
- is NetWorker requesting any backup media to be mounted?
- is it time to reduce the size of the on-line index?

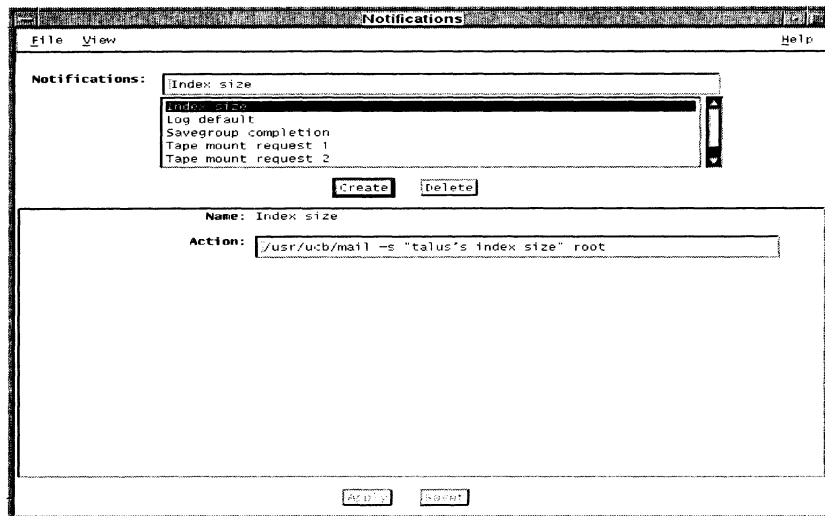
NetWorker is shipped with seven types of notification: *Registration, Savegroup completion, Log default, Index size, Tape mount request 1, Tape mount request 2, and Tape mount request 3.*

Important ⇨ These notices alert you regarding important NetWorker events. If a group of clients did not complete a nightly backup, for example, you will be informed of this by the savegroup completion notices in your electronic mail.

Open the Notifications window by selecting Notifications from the Administration pull-down menu.



The Notifications window is shown below:



3

The Notifications scrolling list displays the names of all the notifications known to the NetWorker server. Information about the currently selected notification is displayed in the lower half of the window.

The Name field displays the name of the currently selected event notification.

The Action field displays what 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.

Log Default

The log default notification uses a UNIX facility called **syslog** to log and distribute notification about *all* NetWorker events. These events include requests for backup volume mounts, index size notices, and savegroup completion notices. How this information is distributed depends on how you have configured **syslog**. When NetWorker was installed, it created entries for logging and contacting operators. You can customize these entries. Refer to the **syslog** man page for information on configuring the distribution of log information.

3

Name:	Log default
Action:	<code>/usr/bin/logger -p daemon.notice</code>

Index Size

NetWorker checks the size of its on-line 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*.

Name:	Index size
Action:	<code>/bin/mail -s "atlas's index size" root</code>

The example above notifies you when the index for the client *atlas* is getting large. If you want the message to be mailed to someone other than *root*, you can edit Action and substitute *root* with a different user login name or mailing list, for example. If you see

this message, see Chapter 4, "Managing the On-line Indexes," for information on reducing the size of the indexes.

Savegroup Completion

When NetWorker finishes backing up a group of clients, it sends a completion message via electronic mail to *root*.

Name: Savegroup completion
Action: <code>/bin/mail -s "atlas's savegroup completion" root</code>

3

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 Main window. If no one fills the request, NetWorker sends another request after fifteen minutes. NetWorker sends a third request after another thirty-seven minutes, if no one fills the request.

The first mount request has a blank Action field, so the request will appear only in the Pending display of the NetWorker Main window. The second mount request sends an alert to the logger, and the third request sends electronic mail messages to *root*.

Name: Tape mount request 3
Action: <code>/bin/mail -s "atlas's tape mount request" root</code>

Configuring the Server's Clients

Before a client can communicate with the NetWorker server, it must have access to the NetWorker client software. Clients may have the software installed locally on their disks, or may NFS-mount the software over the network. Refer to the NetWorker release notice for installation details. After the software has been successfully installed on the client, you may configure the NetWorker server for that client.


If you try to add more clients to your NetWorker server beyond the number you have purchased, you will receive the error message "Too many clients – maximum is *number*."

NetWorker clients can manually back up and recover files from a NetWorker server and can browse the on-line 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.

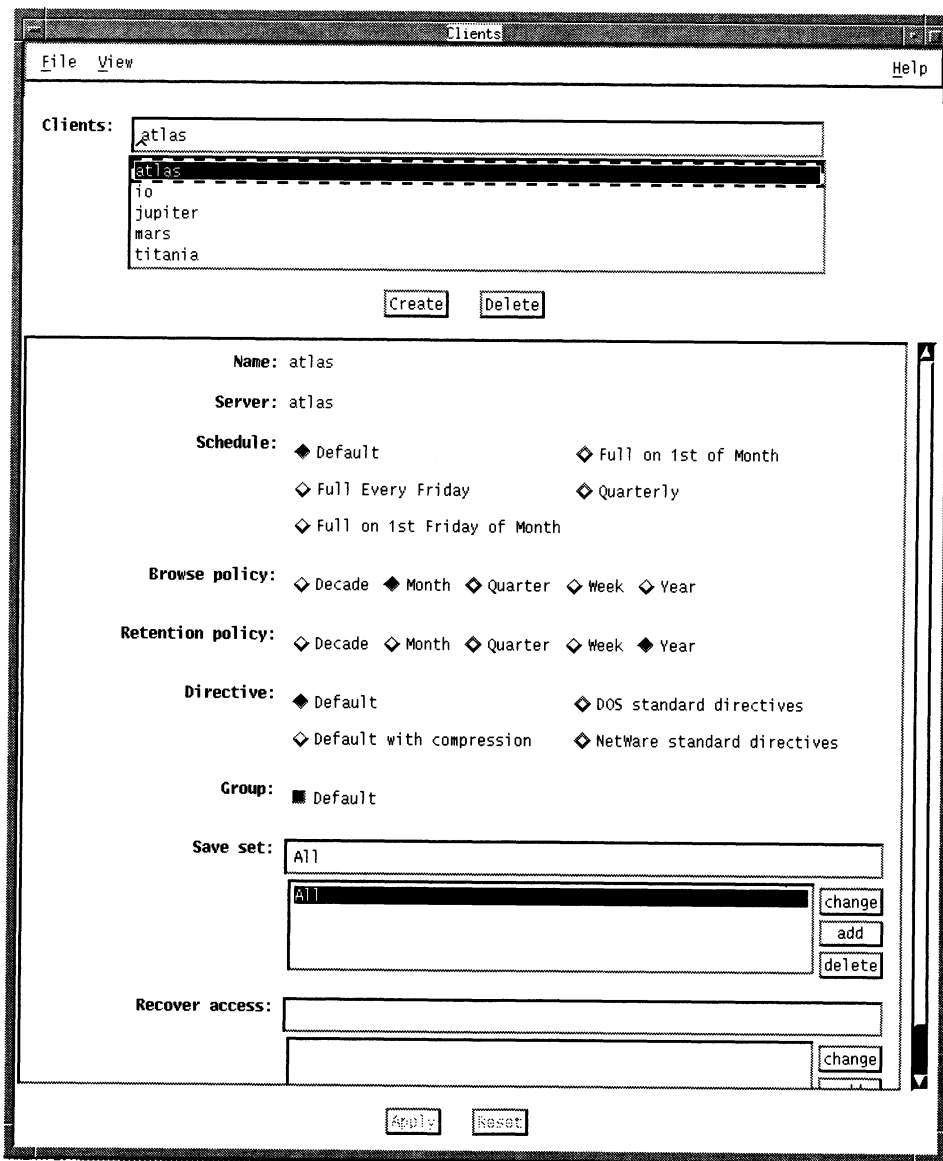
When you configure a client, you choose the following:

- a NetWorker server, if you have more than one
- a backup schedule
- a browse and a retention policy
- a directive
- a backup group (none, one, or many)
- which client files you want NetWorker to back up, if not all of them
- what other machines, if any, are allowed to recover the client's files

The client configuration tasks are completed in the Clients window. Select Clients from the Administration pull-down menu to open the Clients window.

**3**

The Clients window is shown below:



3

The Clients scrolling list displays all the NetWorker clients configured for the NetWorker server.

The Name field displays the *hostname* of the currently selected client.

The Server field displays the *hostname* of the currently selected NetWorker server.

The Schedule choices display all the schedules known to the NetWorker server. NetWorker is shipped with five pre-configured schedules for your convenience. Any other schedule you create using the Schedules window will also appear here as a choice.

The Browse policy and Retention policy choices display the policies for automatic index management known to the NetWorker server. NetWorker is shipped with five pre-configured policies for your convenience. Any other policy you create using the Policies window will also appear here as a choice.

The Directive choices display the directives recognized by the NetWorker server. There are four sets of pre-configured directives shipped with NetWorker, and any new ones you create using the Directives window will also appear here.

The Group choices display all the backup groups known to the NetWorker server.

The Save set scrolling list displays the filesystems NetWorker backs up for the client. The pre-configured setting "All" is a key word that means "back up *all* locally-mounted filesystems for this client."

The Recover access scrolling list displays the *hostnames* or *Netgroups* of the machines that are allowed to browse and recover the client's files. NetWorker clients are pre-configured so that only the client itself can browse and recover its files.

To configure a new client, all you need to do is provide the client's *hostname* and make choices in the Clients window. To

modify an existing client, select its name in the Clients scrolling list and make changes to the window.

Adding a New Client

To configure a new NetWorker client, follow these steps:

1. Open the Clients window.
2. Press 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.

Next, 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:

1. Select All in the Save set scrolling list.
2. Press the Delete button.
3. Enter the *pathname* for the filesystem that you want backed up into the Save set field. For example:

Save set:

usr/etc	change
	add
	delete

4. Delete All from the Save set scrolling list.

The *pathname* will appear in the Save set scrolling list.

Important ⇨ This feature allows you to schedule filesystems to be backed up at different times. For example, if a client has a large amount of data, you may want to schedule its filesystems separately for backup. For more information, see "Scheduling a Large Filesystem for a Client" in this chapter.

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

If you want to give other machines access to the client's files, follow these steps:

1. Enter the *hostname* of the machine or netgroup name to which you are giving access to the client's files in the Recover access field.

Recover access:

<input type="text"/>	change
<input type="text"/>	add
<input type="text"/>	delete

2. Press the add button next to the Recover access scrolling list.

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

Next, choose a backup schedule, browse policy, retention policy, and directive for the client:

1. Select a schedule from the Schedules choices. Pre-configured schedules provided by NetWorker, and any new schedules you created with the Schedules window will appear here.
2. Select both a browse and a retention policy for the client indexes. You may use the pre-configured policies of Quarter for the Browse policy and Year for the Retention policy.
3. Select a set of directives for the client. The four pre-configured sets of directives will appear as Directive choices.
4. Press the Apply button to apply the changes and add the new client to the server.
5. Your new client will appear in the Clients scrolling list of the window.

If you make a mistake, press the Reset button to reset the window to your last apply.

Selecting Policies for a Client

After you have created the policies in the Policies window, you select them for a client as a Browse policy or as a Retention policy.

To select policies for a client, open the Clients window by selecting Clients from the Administration pull-down menu.

Notice the Browse policy and Retention policy selections in the Clients window. The five pre-configured policies and the custom

policies that you have created appear as buttons for both Browse and Retention policy choices.

Browse policy: ◊ Decade ◊ Month ◆ Quarter ◊ Week ◊ Year
Retention policy: ◊ Decade ◊ Month ◊ Quarter ◊ Week ◆ Year

Press the button for the desired Browse or Retention policy. The pre-configured selections are Quarter for the Browse policy and Year for the Retention policy.

Important ⇨ Remember, a client's Browse policy may not exceed the Retention policy. An error message appears if you attempt to assign a Browse policy to a client that is longer than the Retention policy assigned to that client. For example, one month for the Browse policy and one week for the Retention policy is unacceptable because a month is greater than a week.

3


Scheduling Large Client Filesystems

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

You can, however, schedule the client *filesystems* to be backed up at different times. By separating the client's filesystems into different backup groups, you will still be able to 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 displays the filesystems that NetWorker backs up for the client. All the

filesystems are backed up, as long as the keyword All is left in the list.



To separate the filesystems, you add the client *twice* to the server list of clients to back up. In one case, you configure the client to back up half its filesystems, and in the other case, the other half of its filesystems. These repeated clients can also have different schedules and belong to different backup groups for more efficient backups.

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

```
% df -lk
Filesystem          kbytes  used   avail  capacity  Mounted on
/dev/dsk/root        15087  9898   3680   73%       /
/dev/dsk/usr         361474 313232 12094  96%       /usr
/dev/dsk/usr_opt_X11 559860 362532 141342 72%       /usr/src
/dev/dsk/export      818627 695502 41262  94%       /export
/dev/dsk/tmp         28181  10713  14649  42%       /var
/dev/dsk/udd_jupiter 559860 435366 68508  86%       /udd/jupiter
```

The largest filesystem is `/export`, with 818627 kbytes of space.

Next, estimate how long it would take to back up this filesystem:

```
818627 kbytes / 300 kbytes/second = 2728.76 seconds
2728.76 seconds / 60 seconds/minute = 45.48 minutes
```

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

For example, this client's filesystems will be separated as follows:

```

/dev/dsk/root          15087  9898   3680  73%  /
/dev/dsk/usr           361474 313232 12094  96%  /usr
/dev/dsk/usr_opt_X11  559860 362532 141342 72%  /usr/src
    
```

will be backed up according to a "Monday Full" schedule, and


```

/dev/dsk/export       818627 695502 41262 94%  /export
/dev/dsk/tmp           28181  10713  14649 42%  /var
/dev/dsk/udd_jupiter559860 435366 68508 86%  /udd/jupiter
    
```

will be backed up according to a "Wednesday Full" schedule. Each full backup should take no more than six hours to complete. (The NetWorker server has schedules called "Monday Full" and "Wednesday Full" that were previously created.)

Next, use the Clients window to configure the client:

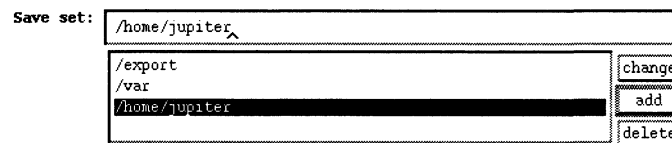
1. Open the Clients window.
2. Select the client whose filesystems you are separates in the scrolling list, or press Create to create a new client.
3. Configure the client by choosing a backup group, directive, policy, and the "Monday Full" schedule.
4. Delete "All" from the Save set list and add each filesystem you want to back up on the "Monday Full" schedule, one at a time:



5. Press the Apply button.

Now follow the same procedure for the remaining filesystems you want to back up on the "Wednesday Full" schedule.

1. Press the Create button to add the client a second time with different filesystems scheduled.
2. Configure the client by choosing the backup group, directive, and policy; 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, one at a time:



Save set: /home/jupiter_x

/export	change
/var	add
/home/jupiter	delete

4. Press the Apply button.

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

Important ⇨ 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 will *never* be 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. Press the Delete button. NetWorker will ask for a confirmation to delete the client.

Deleting a client means it cannot back up its files to the NetWorker server or recover its files. All the backup history for the client still remains in the file index and media index, until you specifically remove it. See "Removing a Backup Volume" in Chapter 4 for instructions on how to remove the backup history for a client from the indexes.

3

Securing the NetWorker Environment

The NetWorker server validates client access by checking the list of clients known to the server. The server will only accept connections that are initiated from a secure port on the client. Secure ports can only be opened by *root*, so NetWorker programs run **setuid** to *root*. 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 Recover access list in the Clients window.

Once a connection has been established, the client programs **save**, **savefs**, and **recover** set their effective uid to the uid of the user who initiated the program so that all local filesystem and system call access is done as that user. This prevents users from recovering or backing up files to which they should not have access. The exception to this rule is that the user name *operator* and users in the group *operator* have filesystem access privileges of *root*. This allows the administrator to set up a login or group

for the operators who will initiate backups and recovers on behalf of other users, without giving the operators *root* access to client machines. For more information about security, refer to the `nsr` man page.

Restricting Client Access

NetWorker is pre-configured to 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 Recover access list in the Clients window. See "Configuring the Server's Clients" in this chapter for more information.

Access control for the client programs can be further tightened by turning off the set-uid bit. This will restrict the client machines' use of the `save`, `savefs`, and `recover` programs to *root*. To allow access by *root* and *operator*, but not by other users, change the group ownership of these programs to *operator*, and set the mode bits to allow execution by owner and group, but not by others.

The `savegroup` command initiates the `savefs` command on each client machine in a backup group by sending a remote command request to the `nsrexecd` daemon.

The `nsrexecd` daemon runs on NetWorker clients. This daemon, which starts whenever a client boots and enters run level 3, provides a secure and restrictive way for NetWorker to start automatic backups on clients.

Chapter 4

Managing the On-line Indexes

This chapter describes NetWorker's powerful index management features for manually managing the on-line indexes. You can also configure NetWorker to automatically manage the indexes using the Policies window. (See "Creating Index Policies" in Chapter 3 for more information on automatic index management.)

This chapter provides you with the following concepts of indexes and index management:

- ❑ a description of on-line indexes
- ❑ a description of save sets
- ❑ a description of the four actions that can be performed on an on-line index

Also included in this chapter are the following instructions on manually managing the indexes and backup volumes:

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

On-line Indexes

NetWorker maintains two types of indexes: a *file index* and a *media index*. The file index stores information about the files

backed up by NetWorker, and the media index maps the backed-up files to backup volumes. These indexes are used by NetWorker to locate the files requested for recovery.

Each entry in the file index typically includes the following information for a backed-up file: filename, number of blocks, access permissions, number of links, owner, group, size, last modified time, and backup time. The file index 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, the file may be recovered using the NetWorker Recover window.

The media index is usually much smaller than the file index. 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 file and media indexes require disk space. The size of an index is proportional to the number of entries stored in it. With NetWorker Policies, Indexes, and Volumes windows, you create policies for automatic index management, monitor the contents of the indexes, select entries for removal, and mark backup volumes as recyclable.

4

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.

To conserve both index and backup volume space, NetWorker provides a variety of backup levels. For example, a *full* backup contains all files in a given filesystem, while an *incremental* backup contains only the files that have changed since the previous backup. Since directories often contain some files that

do not change together with files that do change, incremental backups conserve both index and volume space. (See "Understanding the Backup Levels" in Chapter 3 for more information on backup levels.)

To recover a complete directory, some files will be needed from the incremental backup and some from the full backup. Without the underlying full backup, the directory could not be completely recovered. Without the incremental backup, only the old version of the directory could be recovered. The incremental backup depends on the full backup. Since both the full and incremental backups are needed to recover the complete directory, NetWorker checks these dependencies when removing save sets from the indexes either manually or automatically. A save set is not removed until all of its dependent save sets have been removed.

The diagram below shows the relationship between incremental and full backups.



Important ⇔ Incremental backups are dependent upon the previous full backup. A save set is not removed until all of its dependent save sets have been removed.

Index Actions

There are four actions that can be performed on an index: *inserting entries, browsing, removing entries, and reclaiming space.*

Inserting entries in an index occurs during a backup. The entries are inserted into the free space in the indexes. If there is no free space in an index, more space will be acquired from the filesystem to hold the new entries.

Browsing the file index occurs when you are using the Recover window to locate a file. Browsing the media index occurs when you are using the Volumes or Indexes windows to view the contents of a backup volume. Browsing neither increases nor decreases the size of an index. When browsing an index, you are simply looking through the index for information concerning your saved files or the contents of your backup volumes.

Removing entries frees space in the index. The free space will be used when new entries are inserted. Entries are automatically removed based on the length of time they have been stored in the index using the Browse and Retention policies. They may also be removed manually by selecting *Remove oldest cycle* in the Indexes window or *Remove* on the Volume menu in the Volumes window.

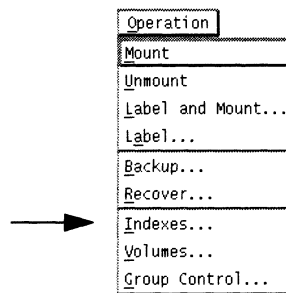
Reclaiming space returns the empty space, created when entries are removed from the index, to the filesystem. It occurs when you use the *Reclaim space* command in the Indexes window.

Managing the Indexes Manually


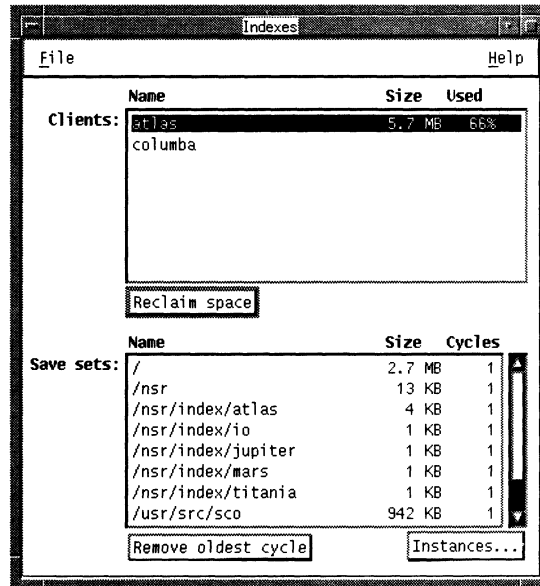
Using the Policies and Clients windows, you configure NetWorker to automatically maintain your on-line indexes at a steady size, containing only the most current backups. (For instructions on creating policies that automatically manage your indexes, see the section entitled "Creating Index Policies" in Chapter 3.)

Using the Indexes and Volumes windows, you can manually remove entries for save sets from the file index and remove entries for backup volumes from the media index.

To manually manage your file indexes, open the Indexes window by selecting Indexes from the Operation pull-down menu.



The Indexes window appears, as shown below.



The Indexes window has two scrolling lists: Clients and Save sets.

The Clients scrolling list provides the names of the client machines configured for the current NetWorker server and the following information about each client file index:

- Name – NetWorker client names
- Size – the current allocated size of the client file index
- Used – the percentage of this space actually used by the entries in the file index

Selecting the Reclaim space button decreases the size of the highlighted client file index by compressing the space in the index left by removed entries, thus freeing up disk space.

The Save sets scrolling list displays the save sets contained in the file index for the highlighted 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*.

The Save sets scrolling list provides the following information about a client save set:

- Name – client save set names
- Size – An estimate of the amount of index space used by the save set group in the client file index.
- Cycles – The number of cycles of 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.

The Remove oldest cycle button removes the oldest full-to-full cycle of the highlighted save set group.

The Reclaim space and Remove oldest cycle buttons are used for the index operations described below.

The Reclaim space button in the Indexes window displays a dialog box that starts an index operation. The Remove oldest cycle button displays a confirmation box that starts an index operation. Since the index can perform only one operation at a time, the Reclaim index space dialog box has a message line that reflects the current action on an index. For example, if an index is busy removing the oldest cycle, the "Started removing oldest cycle" message appears in the Reclaim Index Space dialog box.

These messages appear before you start the index operation. The messages for the index operations are listed below:

- Removing oldest cycle may take considerable time
- Reclaiming space may take considerable time

These messages appear after you press the Apply button. Day of the week is Monday – Sunday and hh:mm uses the 24 hour clock for hours and minutes. The Apply button is disabled while the index is busy.

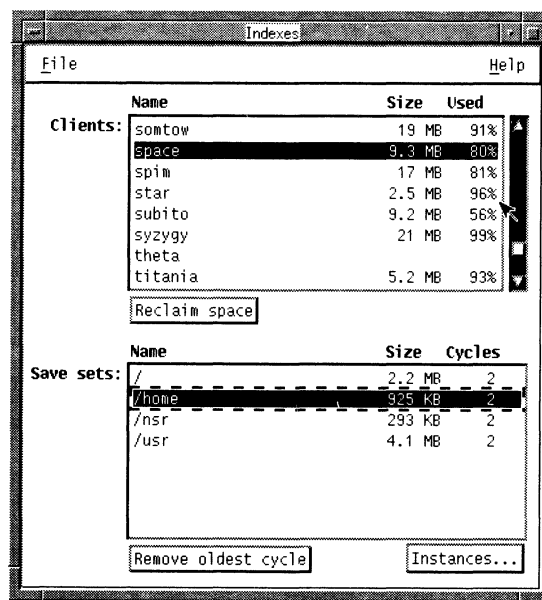
- Started removing oldest cycle at *day of week hh:min*
- Started reclaiming space at *day of week hh:min*
- Index is currently busy

This message applies to the Remove oldest cycle operation only. It appears when there are no entries for files in the client file index, and thus there are no save sets to remove.

- There are no save sets to purge for *client_name*

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, select the client and the save set in the scrolling lists in the Indexes window. For example:

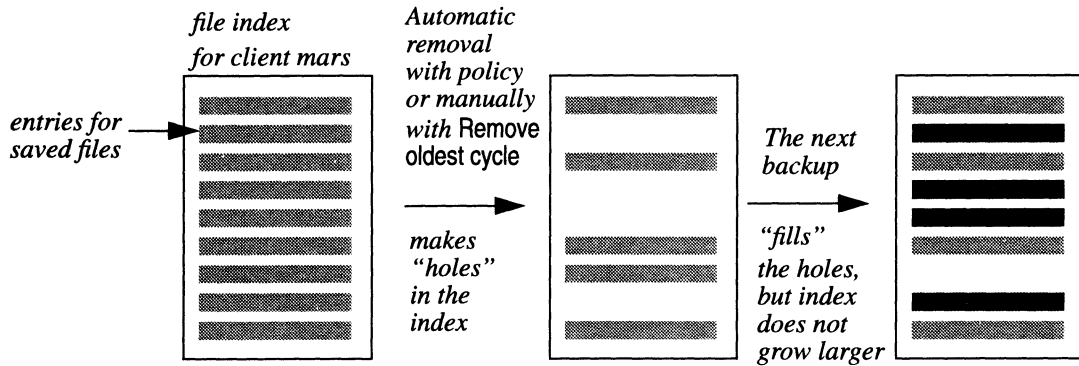


Press the Remove oldest cycle button to remove the oldest cycle of the save sets highlighted in the Save sets scrolling list. A confirmation box appears, prompting you to confirm that you want to remove the oldest cycle of the selected save set for a client.

Important ⇔ If you do not highlight any save sets in the Save sets scrolling list, then NetWorker will remove *all* the oldest cycles of *all* the save sets displayed in the scrolling list for the client highlighted in the Clients scrolling list.

The diagram below illustrates 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.




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, you press the Reclaim space button.

Tip 


The last full cycle may not be removed using the Remove oldest cycle button.

Reclaiming Index Space

To reclaim the empty space created in the client file index when entries are automatically or manually removed, press the Reclaim space button.




The Reclaim Space dialog box appears, as shown below.




To start the Reclaim Space operation, press the Ok button. The Progress bar moves to show the progress of the index operation, as shown below. When an index is busy, the Ok button is disabled.

4



To dismiss the dialog box, press the Cancel button. This closes the window; however, the Reclaim space operation continues in the background. You can check the progress of the operation at any time by pressing the Reclaim space button to re-display the dialog box. When you display a dialog box for a busy index, the Progress bar is updated.

The diagram below illustrates what happens when you use the Reclaim space button. The index is re-written without the deleted entries, consequently it shrinks in size and opens up space for new entries. The space is returned to the filesystem.



Important ⇔ You may not simultaneously Reclaim space and Remove oldest cycle for one client. The index will be busy during both operations.

After either the Reclaim space or Remove oldest cycle operation has finished, the statistics in the Indexes window are updated to reflect the current state of the file index.

Viewing Save Set Details

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

The Instances window contains information on the number of files in a save set, its size, the date it was backed up, and the level of the save set: full, incr, and level.

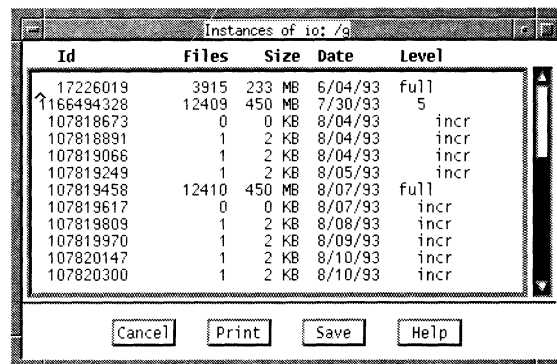
You can use the information in this window to determine how your resources are being used. For example, you may need to see how large a save set is so you can plan the amount of disk space you will need for the on-line indexes. Or you may need to

generate a report that details your NetWorker client backups to fulfill administrative requirements.

To open the Instances window follow these steps:

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. Press the Instances button.

The Instances window appears, as shown below.




Id	Files	Size	Date	Level
17226019	3915	233 MB	6/04/93	full
↑166494328	12409	450 MB	7/30/93	5
107818673	0	0 KB	8/04/93	incr
107818891	1	2 KB	8/04/93	incr
107819066	1	2 KB	8/04/93	incr
107819249	1	2 KB	8/05/93	incr
107819458	12410	450 MB	8/07/93	full
107819617	0	0 KB	8/07/93	incr
107819809	1	2 KB	8/08/93	incr
107819970	1	2 KB	8/09/93	incr
107820147	1	2 KB	8/10/93	incr
107820300	1	2 KB	8/10/93	incr

The title bar for the Instances window includes the name of the save set you are currently viewing. The Instances scrolling list contains the following information:

- ❑ Id – internal NetWorker id number for the save set
- ❑ Files – number of files in the save set
- ❑ Size – the size of the save set in megabytes or kilobytes
- ❑ Date – the date the save set was backed up
- ❑ Level – level of backup: full, incr, or level


The Instances window contains the Cancel, Print, Save, and Help buttons. Selecting the Help button will open the Save Set Instances Help window which contains information on related subjects; selecting the Cancel button closes the Indexes window.

To print the contents of the Instances window select the Print button. The Print Options dialog box appears, as shown below.



Enter the print command (for example, **lp -dlpr**) in the Command field and enter the number of copies you want in the Copies field, and press Ok to print.

To save the contents of the Instances window to a file press the Save button. The Save Options dialog box appears, as shown below.



The information is saved to a file named *networker.rpt*, which NetWorker automatically creates. However, you can save the data to a file of your choice by entering its name into the File field. You can use either the Append button to add the information or use the Overwrite button to replace the existing information in the file. Pressing Ok saves the information from the Instances window to the file.

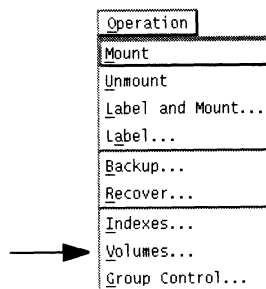
4

Managing the Volumes Manually

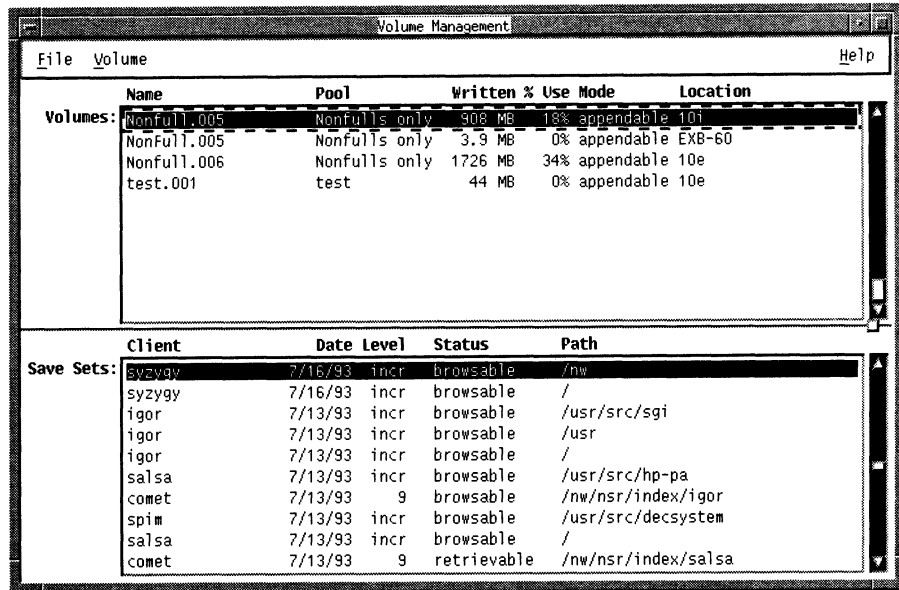
With automatic index management using policies, a backup volume is marked as recyclable when all of its save sets have passed the time period designated by the Retention policy. A recyclable volume is eligible for overwriting with new backups. Using the Volume Management window, you can manually mark entire backup volumes as recyclable, which also marks all the save sets on the volume as recyclable.

The Volume Management window displays the information stored in the media index (located on the server) about the backup volumes and their save sets.

To manually manage or to view information about your backup volumes, open the Volume Management window by selecting Volumes from the Operation pull-down menu.

4

The Volume Management window appears, as shown below.



The screenshot shows a window titled "Volume Management" with a menu bar containing "File", "Volume", and "Help". The window is divided into two main sections. The top section, labeled "Volumes:", contains a table with the following data:

Name	Pool	Written	% Use	Mode	Location
NonFull.005	NonFull's only	308 MB	18%	appendable	10i
NonFull.005	NonFull's only	3.9 MB	0%	appendable	EXB-60
NonFull.006	NonFull's only	1726 MB	34%	appendable	10e
test.001	test	44 MB	0%	appendable	10e

The bottom section, labeled "Save Sets:", contains a table with the following data:

Client	Date	Level	Status	Path
syzygy	7/16/93	incr	browsable	/nw
syzygy	7/16/93	incr	browsable	/
igor	7/13/93	incr	browsable	/usr/src/sgi
igor	7/13/93	incr	browsable	/usr
igor	7/13/93	incr	browsable	/
salsa	7/13/93	incr	browsable	/usr/src/hp-pa
comet	7/13/93	9	browsable	/nw/nsr/index/igor
spim	7/13/93	incr	browsable	/usr/src/decsystem
salsa	7/13/93	incr	browsable	/
comet	7/13/93	9	retrievable	/nw/nsr/index/salsa

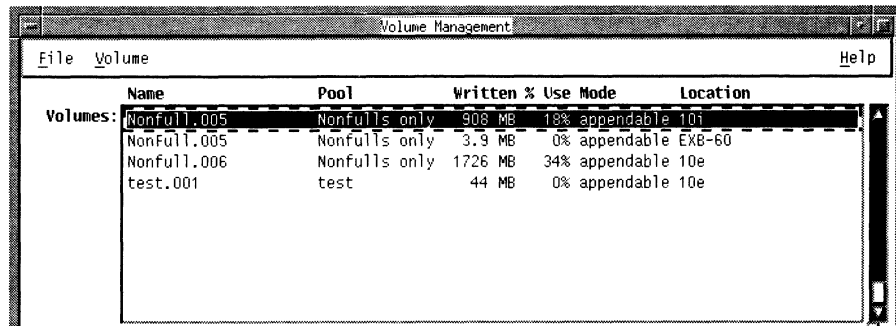
4

The Volume Management window displays information about the backup volumes and their save sets for the current NetWorker server. The information in the Volume Management window is taken from the media index.

The Volume Management window has a Volumes scrolling list and a Save sets scrolling list. 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

You can use the Volumes scrolling list to view the backup volumes for a server. The Volumes scrolling list is shown below.



The screenshot shows a window titled "Volume Management" with a menu bar containing "File", "Volume", and "Help". Below the menu bar is a table with the following columns: "Name", "Pool", "Written", "% Use", "Mode", and "Location". The table contains the following data:

Name	Pool	Written	% Use	Mode	Location
NonFull.005	NonFulls only	908 MB	18%	appendable	10i
NonFull.005	NonFulls only	3.9 MB	0%	appendable	EXB-60
NonFull.006	NonFulls only	1726 MB	34%	appendable	10e
test.001	test	44 MB	0%	appendable	10e

4

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

- The Name of the volume
- The name of the Pool to which the backup volume belongs
- The amount of data Written on the backup volume
- The percentage of the backup volume used compared to its total capacity (% Use)
- The Mode of the backup volume, either appendable, full, or recyclable. Appendable means that there is room for more data on the backup volume. Full means that there is no more room for data on the backup volume, and that the save sets have not yet passed the time period specified by the Browse and Retention policies. Recyclable means that all the save sets on the volume have passed both the Browse and Retention policies and are eligible to be overwritten.
- The Location of the volume as entered by the user. This serves as a reminder of where your volumes are physically located.

Viewing the Save Sets Scrolling List

You can view the save sets stored on a volume in the Save sets scrolling list. The Save sets scrolling list is shown below.

Client	Date	Level	Status	Path
syzygy	7/16/93	incr	browsable	/nw
syzygy	7/16/93	incr	browsable	/
igor	7/13/93	incr	browsable	/usr/src/sgi
igor	7/13/93	incr	browsable	/usr
igor	7/13/93	incr	browsable	/
salsa	7/13/93	incr	browsable	/usr/src/hp-pa
comet	7/13/93	9	browsable	/nw/nsr/index/igor
spim	7/13/93	incr	browsable	/usr/src/decsystem
salsa	7/13/93	incr	browsable	/
comet	7/13/93	9	retrievable	/nw/nsr/index/salsa

The Save sets scrolling list provides the following information about each save set stored on a volume:

- The name of the Client machine that created the save set
- The Date that the save set was created
- The Level of the backup that generated the save set. Backup levels apply only to scheduled backups. If no level is listed, the save set was generated by a manual backup initiated by a client system, or it is a save set from a backup volume that was created with NetWorker version 1.0, which did not support backup levels.
- The Status of the save set. Save set status may be either browsable, retrievable, or recyclable.

Browsable status means that 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 Recover window.

Retrievable status means that the entry for the save set has been removed from the client file index, either automatically due to the client Browse policy or manually using the Remove oldest cycle button in the Indexes window. In order to recover any files contained in this save set, the file index entries must be re-created using the **scanner** command. (See the **scanner** man page for information on using the **scanner** command.)

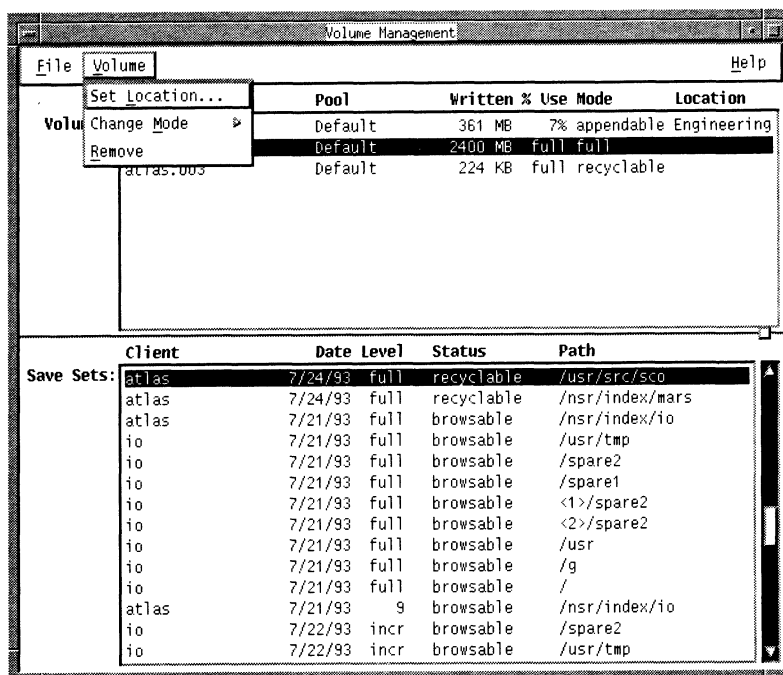
Recyclable status for a save set means that it has passed both the Browse and Retention policies time period and is eligible to be overwritten.

- ❑ The Path of the filesystem contained in the save set

Managing Backup Volumes

4

You can manually manage your backup volumes using the Volume pull-down menu. You can set a location for a volume, change the mode of a volume, or remove a volume from the media index with the Volume pull-down menu commands.

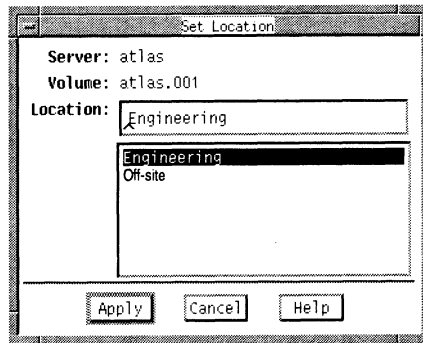


Setting the Location for Your Backup Volumes

As an optional reminder to yourself, you can set a location for a backup volume. This will help you find a backup volume when it is needed for recovering a file. A location can be a numbered bin, shelf, jukebox job pack, or any place that you physically store backup media.

To set the location for a backup volume:

1. Highlight a backup volume in the Volumes scrolling list.
2. Select Set Location from the Volume pull-down menu. The Set Location dialog box appears, as shown below.



3. Enter the location of the backup volume.
4. Press the Apply button to designate this location for the selected backup volume. The Volume Management window is updated to reflect the location for that backup volume.

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

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 to it during a backup. There are three possible modes for a backup volume:

- Appendable – There is room on the backup volume for more backups. NetWorker can *append* data to the backup volume and will use this backup volume for backups.

- ❑ Full – The capacity of the volume has been reached. A volume may also be in full mode if a “write error” is encountered during a backup. If this happens, NetWorker continues the backup on the next appendable volume or requests a new writable volume.
- ❑ Recyclable – The volume is available to be re-labeled and overwritten by future backups. You can manually mark an entire volume as recyclable by changing its mode. If a backup volume's mode is manually changed to recyclable, then all the save sets stored on the volume are also marked as recyclable.

NOTE: A backup volume is automatically recyclable if all the save sets on the volume are recyclable. Recyclable save sets have passed the time period designated by the Retention policy for the client.


To change the mode of a volume to recyclable so that it can be overwritten with new backups:

1. Highlight a volume that is in either full or appendable mode in the Volumes scrolling list. Notice that the save sets stored on the volume appear in the Save sets scrolling list.


Tip 

Look at the status of the save sets stored on the volume. If the status of all of the save sets is recyclable, then NetWorker automatically marks the volume as recyclable.

2. Select Change Mode from the Volume pull-down menu in the Volume Management window. Then select Recyclable from the Change Mode menu.




The following confirmation box appears:



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

To change the mode of a recyclable volume to appendable, highlight the volume in the Volume Management window and select Appendable from the Change Mode menu. Press the Ok button in the confirmation box.



Important ⇨

When you manually change the mode of a volume to recyclable in the Volumes Management window, all of its save sets are given recyclable status, regardless of their relationship to the Browse or Retention policies specified in the Clients window. If you are not certain that you want all of the save sets on a volume to be recyclable, then do not manually change the mode.

Removing a Backup Volume

Removing a backup volume occurs in two stages. First, you remove the entries from the appropriate file indexes. This changes the status of the browsable save sets to retrievable. You remove the file index entries when you do not need to browse and recover the files any longer using the Recover window.

Second, you remove the name of the backup volume from the media index, thus destroying all record of it. You remove a backup volume from the media index only if you will never need the data on the volume again, or if the volume has been physically damaged and is unusable. It is very rare that you will need to remove a backup volume from the media index.


Tip 

If a backup volume is lost or destroyed, you should remove it from both the on-line file and media indexes. Otherwise, NetWorker may request the backup volume to be mounted on the server backup device during a recover.

4


To remove the backup volume entries from the file index:

1. Highlight the desired volume in the Volumes scrolling list. Notice that the save sets stored on the volume appear in the Save sets scrolling list.
2. Select Remove from the Volume pull-down menu. The notice "Remove volume *name* on-line index entries?" appears, as shown below.



3. Press the **Ok** button to remove the entries for the files stored on the volume from the appropriate file indexes. The status of the save sets containing these files changes to retrievable in the Save sets scrolling list.
4. Press the **Cancel** button if you do not want to remove the entries for the volume files from the file indexes.

If you press **Ok** to remove the entries from the on-line indexes, the "Remove from media index" dialog box appears, as shown below.



1. Press the **Ok** button to remove the volume from the media index for the server. All record of the data on this backup volume is removed from the NetWorker indexes.
2. Press the **Cancel** button if you do not want to remove the volume from the server media index.

Important ⇨

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

Chapter 5

Recovering from a Disk Crash

The DG/UX operating system supports a number of advanced data recovery features and protections against system failure, including:

- ❑ Support for the CLARiiON™ disk-array storage system. This mass storage subsystem performs hardware disk mirroring, which allows you to replace failed disks dynamically without interrupting system activity. If your AViiON system includes a CLARiiON system, you may not have to recover any data with NetWorker after a disk failure.
- ❑ Software disk mirroring. This feature allows you to arrange for multiple identical images of a filesystem to be maintained on different disks. If a disk containing a software-mirrored filesystem fails, you will not have to restore any data to that filesystem with NetWorker.
- ❑ Dual-initiator configuration and failover. These features allow you to assign tape and disk drives to two hosts simultaneously, such that if one host fails, the devices can be accessed from the secondary host. If the NetWorker server is part of a dual-initiator configuration, it may be possible for network backups to proceed normally even though the NetWorker server is down.

This chapter discusses recovery from a disk crash that requires restoration of files from backup media with NetWorker. Before following the procedures described, first determine whether they are necessary. If your system includes a CLARiiON disk array, see its manual to make this determination. Software disk mirroring, dual-initiator configuration, and failover are described in

Chapter 7 of *Managing the DG/UX System* (093-701088). Use of dual-initiator configuration and failover with a CLARiiON disk-array storage system is discussed in *Physical Disk Failover in the DG/UX 5.4.2 Operating System* (012-004186) and *Achieving High Availability on the DG/UX System* (093-701133).

The procedures required to recover from a disk failure vary depending on what was on the disk. The failed disk may have contained one or more of the following:

- ❑ The DG/UX operating system. If this happens, you may have to re-install the DG/UX operating system from the release tape as explained in *Installing the DG/UX System* (093-701087).
- ❑ The NetWorker software. In this case, you may have to re-install NetWorker from its release tape as described in the NetWorker release notice and Chapter 5 of *Managing the DG/UX System*.
- ❑ The NetWorker index files. If this happens, you'll have to recover the index files before doing anything else with NetWorker. For this procedure, see "Recovering the NetWorker Index Files" in this chapter.
- ❑ One or more filesystems that NetWorker backs up. For this procedure, see "Recovering a Filesystem" in this chapter.

Before you can use NetWorker to recover any files, you must first replace (or repair) the disk and prepare it for use with the DG/UX system. You perform these procedures with the DG/UX **sysadm** program, as explained in Chapters 7 and 8 of *Managing the DG/UX System*. The next section describes information you may need to rebuild a disk so that NetWorker can recover its contents.

Preparing for a Crash

The ultimate disaster for a system is to lose all the files on a disk. Most sites back up their file servers daily in preparation for this event. If a system's primary disk suffers a crash, you can rebuild its filesystems with NetWorker, after you re-install the operating system (if necessary).

If the filesystem that contains the NetWorker index files (*/nsr*) is destroyed, the recovery procedure involves an extra step – you must recover the server's on-line 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 will need to replace it with another machine. You may do this as long as you:

- ❑ name the replacement server with the same *hostname* as the original NetWorker server
- ❑ re-install NetWorker using the same directory locations for the on-line indexes as in the original installation.

Important ⇨ Once you understand the procedure for a disaster recovery, make sure you have carefully thought of a disaster recovery plan for your site. If possible, you should test the 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 on-line indexes for those clients, including the indexes for the server itself. The server's index backup is named *bootstrap*, and is assigned an identification number called the *save set id* (ssid). This information is sent to a default printer, providing a hard copy for your files.

We recommend you take two additional precautionary steps to help you recover from a future crash:

- ❑ Keep a file containing hard copies of the *bootstrap* records. Place these daily sheets of paper in a three-ring binder or a file folder.
- ❑ Keep an up-to-date hard copy record of the disk layout, filesystem sizes, and mount points for the server and any clients that have hard disks.

File 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 on-line indexes, you will need the information on this piece of paper. Save these records in a binder or file.

The information sent to the printer looks similar to this:

```
August 20 03:30 1993 NetWorker bootstrap information Page 1
```

date	time	level	ssid	file	record	volume
8/19/93	2:29:08	9	1148868949	56	0	Nonfull.005
8/20/93	2:52:25	9	1148868985	77	0	Nonfull.001

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:

```
# saveindex -c server_name
```

To use the **saveindex** command, you must be *root* on the NetWorker server.

File the Disk Information

If a disk that is backed up by NetWorker fails, first you'll have to replace (or repair) the disk and restore the layout. Then NetWorker can restore the files from its indexes. To restore the layout, you will need a record of each disk's. You can get this information by piping the output of **df** and **admvdisk** to a printer. Do this on the NetWorker server and on each of its clients. Keep this information current and on file in a secure place.

The NetWorker indexes record backed-up filesystems based on their local mount points. Use **df** to find out the mount points, as illustrated below:

```

mars% df -kl
filesystem          kbytes  used  avail  capacity  mounted on
/dev/dsk/root        20000   17336  737    96%       /
/dev/dsk/usr         130000  115605 1866   98%       /usr
/dev/dsk/usr_opt_xdt 25000   21892  702    97%       /usr/opt/xdt
/dev/dsk/usr_opt_networker 20000  11385  6687   63%       /usr/opt/networker
/dev/dsk/usr_opt_X11 70000   63634  0      100%      /usr/opt/X11
/dev/dsk/tmp         8623    839    6958   11%       /tmp
/dev/dsk/var_opt_networker 5000   374    4147   8%        /var/opt/networker
/dev/dsk/mail_new    23499   12745  8491   60%       /var/mail
/dev/dsk/srv         122880  56157  54873  51%       /srv
/dev/dsk/srv_swap    173244  30914  125638 20%       /srv/swap
/dev/dsk/relimon     5000    835    3686   18%       /var/opt/relimon
/dev/dsk/admin_SAVE  8500    2630   5056   34%       /local/admin_SAVE
/dev/pdisk/7         579066  579066 0       100%      /pdd/tcpip/cdrom
/dev/dsk/commtg3     519634  434591 34954   93%       /udd/commtg3

```

After replacing a failed disk and creating empty filesystems (left column), mount the filesystems on the same mount point names (right column) as on the failed disk.

Mount point names are also shown in the file */etc/fstab*. As an alternative to using **df**, you can save printed copies of each system's */etc/fstab* file.

DG/UX filesystems reside inside *virtual disks*, which reside on physical disks. A virtual disk name is the same as the filesystem name shown in the **df** output, left column, minus the */dev/dsk*

node. Thus, virtual disk *root* corresponds to `/dev/dsk/root` shown in `df` output.

A virtual disk consists of one or more partitions, which may reside on the same or different physical disks. When re-creating the virtual disks on the replacement disk, you'll need to know how large to make them. The NetWorker restore operation will fail for any filesystem that is too small to contain the backed-up files.

Use `admvdisk` to find out the sizes of a system's virtual disks and which physical disks they reside on, as illustrated below:

```

mars# admvdisk -o list -R
Name          Volume  Size      Type
BAD_AREA      V       1575      partition on cied(0,2)
<no name>    <no name> 639975    physical for cied(0,2)

admin_SAVE    V       17000     partition on cied(0,0)
<no name>    <no name> 639975    physical for cied(0,0)

commtg3       V       1039268   aggregation of 2 pieces
<no name>    <no name> 814599    partition on sd(dgsc(0,7),0,0)
<no name>    <no name> 3933040   physical for sd(dgsc(0,7),0,0)
<no name>    <no name> 224669    partition on cied(0,0)
<no name>    <no name> 639975    physical for cied(0,0)

mail_new      V       46998     partition on cied(0,1)
<no name>    <no name> 639975    physical for cied(0,1)

relimon       V       10000     partition on cied(0,2)
<no name>    <no name> 639975    physical for cied(0,2)

root          V       40000     partition on sd(cisc(0),6,0)
<no name>    <no name> 1015812   physical for sd(cisc(0),6,0)

srv           V       245760    partition on cied(0,1)
<no name>    <no name> 639975    physical for cied(0,1)

srv_swap      V       346488    partition on cied(0,1)
<no name>    <no name> 639975    physical for cied(0,1)

swap         V       135000    partition on cied(0,0)
<no name>    <no name> 639975    physical for cied(0,0)

```

swap_2	V	300000	partition on cied(0,2)
<no name>		639975	physical for cied(0,2)
tmp	V	17246	partition on cied(0,0)
<no name>		639975	physical for cied(0,0)
usr	V	260000	aggregation of 2 pieces
<no name>		240000	partition on sd(cisc(0),6,0)
<no name>		1015812	physical for sd(cisc(0),6,0)
<no name>		20000	partition on sd(cisc(0),6,0)
<no name>		1015812	physical for sd(cisc(0),6,0)
usr_opt_X11	V	140000	partition on sd(cisc(0),6,0)
<no name>		1015812	physical for sd(cisc(0),6,0)
usr_opt_networker	V	40000	partition on sd(cisc(0),6,0)
<no name>		1015812	physical for sd(cisc(0),6,0)
usr_opt_xdt	V	50000	partition on sd(cisc(0),6,0)
<no name>		1015812	physical for sd(cisc(0),6,0)
var_opt_networker	V	10000	partition on cied(0,0)
<no name>		639975	physical for cied(0,0)

If any part of a virtual disk becomes inaccessible, the entire virtual disk becomes inaccessible and will have to be restored. For example, a piece of *commtg3* (see above) resides on the physical disk *cied(0,0)*. If this disk fails, all of *commtg3* becomes inaccessible. You must recreate a virtual disk of at least 1039268 blocks and restore the contents from backup.

When recreating virtual disks after a disk failure, it is not necessary to recreate the original piece divisions or to use the same virtual disk name. For example, instead of creating *commtg3* as a 2-piece 1039268-block virtual disk spanning two different physical disks, as originally, you could recreate it as a single-piece 1039268-block virtual disk named *commtg3_new*. Provided you mount the reconstructed virtual disk as previously (on mount point */udd/commtg3*), NetWorker will restore the filesystem correctly.

Creation and management of disks and filesystems are explained in *Managing the DG/UX System*, Chapters 7 and 8.

Recovering the NetWorker Index Files

This section addresses the case where the filesystem containing NetWorker's index files is lost in a disk failure. (By default, index files are stored on virtual disk *var_opt_networker*, mounted on */var/opt/networker*; */nsr* is a symbolic link to this location.) The indexes contain information NetWorker needs to recover local and client filesystems. Thus, you'll have to recover them first.

You recover the server's indexes from the backup media by typing **recoverindex** and supplying the *bootstrap* (save set identification number) when prompted. If you followed the procedure recommended in the previous section, you will have a piece of paper with the name of the backup media you need and the *bootstrap* ssid.

For example, ssid 1148869870 below is the most recent *bootstrap* backup:

August 20 03:30 1993 NetWorker bootstrap information Page 1

date	time	level	ssid	file	record	volume
8/08/93	7:44:38	full	1148869706	55	0	Full .008
8/09/93	6:12:09	9	1148869754	48	0	Nonfull .001
8/10/93	6:14:23	9	1148869808	63	0	Nonfull .001
8/11/93	6:29:58	9	1148869870	88	0	Nonfull .001

If you do not have this piece of paper, you can still recover the indexes by finding the ssid using the **scanner** command. (See "Finding the Bootstrap Save Set ID" in this chapter.)

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 media:

1. Find the printout of the *bootstrap* save set id information. You will need it for the next two steps.
2. Retrieve the backup media that contains the most recent backup named *bootstrap*, and load it into the server's device.
3. Use the **recoverindex** command to extract the contents of the *bootstrap* backup. For example:

```
mars# recoverindex  
nsr/bin/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(8) and nsr_crash(8)  
man pages for more details.
```

```
What is the name of the tape drive you plan on using  
[dev/rmt/0]?
```

```
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/0. When you have done this, press  
<RETURN>:
```

```
Scanning dev/nrst8 for save set 1148869870; this may  
take a while...
```

```
scanner: scanning 8mm 5GB tape space.006 on dev/rmt/0
```

```
uasmm -r nsr/res/nsr.res
```

```
uasmm -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/nrst8, write
protected
```

The bootstrap entry in the on-line index for mars has been recovered. The complete index is now being reconstructed from the various partial indexes which were saved during the normal saves for this server.

```
# 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 how the shell prompt appears once *bootstrap* is recovered. You can use NetWorker commands such as **nsrwatch** to watch the progress of the server or **networker** to bring up the NetWorker Main window during the recovery of the index.

Replacing the `/nsr/res` Directory

The **recoverindex** command will also recover 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, can not be overwritten or relocated; instead, the recovered `/nsr/res` directory is renamed `/nsr/res.R`.

To complete the recovery of the `/nsr/res` directory, you will need to shut down NetWorker, move the recovered `/nsr/res` directory into its original location, and then restart NetWorker.

Complete these steps after **recoverindex** has finished and this final message appears:

```
nsrindexasm: The on-line index is now fully recovered.
```

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

```
# nsr_shutdown
```

2. Save the original `/nsr/res` directory, and move the recovered version into the correct location.

```
# cd /nsr
# mv res res.orig
# mv res.R res
```

- Restart the NetWorker server. When it restarts, it will use the recovered configuration data.

```
# /usr/sbin/init.d/rc.networker start
```

- Once you have verified that the NetWorker configuration is correct, you can remove the */nsr/res.orig* directory.

```
# rm -r /nsr/res.orig
```

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:

- Place the most recent media used for backups in the server device (for example, in */dev/rmt/0n*).
- Read the contents of the backup media with the **scanner** command as illustrated below:

```
# scanner /dev/rmt/0n
scanner: scanning mars.2 on /dev/rmt/0n

client symbolic name  save time      size  files ssid  S
venus /                04/10/93 03:12:40 4326980 599 17475596 E
mars /                 04/10/93 03:06:11 4478460 621 17475595 E
venus /usr             04/10/93 03:13:58 6918900 965 17475597 E
venus /home/venus     04/10/93 03:14:14 9661320 825 17475598 E
mars bootstrap         04/10/93 03:33:51 26607144 2 17850937 E
```

In this example, the *bootstrap* save set id is 17850937. Once you find the most recent *bootstrap* save set id, you can use the **recoverindex** command to recover the server's index.

Recovering a Filesystem

Once the operating system, NetWorker programs, and NetWorker indexes are available, and after logical disks and filesystems have been created on the replacement disk and mounted, you are ready to recover the filesystems. You should recover one filesystem at a time, for two reasons:

- ❑ NetWorker will add all the files in one filesystem but will stop at a mount point.
- ❑ You may run out of swap space if you add too many filesystems at once to your list of data to recover.

For example, to recover a filesystem on host *venus* mounted as */pdd/venus*, type these commands:

```
venus# recover
recover> add /pdd/venus
recover> recover
```

```
/pdd/venus is being recovered into its original
location
Requesting 33023 files, this may take a while...
```

Repeat this procedure for each filesystem to be recovered. To monitor the process, either open a Main window or use the **nsrwatch** command. NetWorker will send a message telling you when a new tape needs to be loaded.

When you have restored all filesystems, exit the **recover** program:

```
recover> quit
```

Summary

First determine whether NetWorker is needed to recover any files. If the failed disk was in a CLARiiON disk array, you may be able to simply insert a replacement disk. If the filesystems on a failed disk were being mirrored on a different disk, use the software mirror synchronize procedure rather than NetWorker to recover the filesystems.

If the system disk is lost, re-install DG/UX on a replacement disk from the DG/UX release media, using the same *hostname* and the same mount point names for any filesystems to be recovered with NetWorker.

If the disk containing the NetWorker server binaries is lost, re-install NetWorker from its release media. On each NetWorker client, either extract the binaries from the server or temporarily NFS-mount them.

Replace or repair the damaged disk, format and register it, recreate virtual disks and filesystems, and mount them using the same mount point names.

If the filesystem containing the indexes (*/nsr*) is destroyed on a NetWorker server, use **recoverindex** to recover the NetWorker server's indexes.

Once the operating system and NetWorker are working, the disk has been replaced and rebuilt, and the NetWorker indexes are accessible, recover the lost filesystems one at a time.

Appendix A

Theory of Operations

This appendix is for the system administrator or operator who is interested in more information about NetWorker and how it works. You should be familiar with UNIX administrative concepts, the X Window System interface, and the **nsradmin** command.

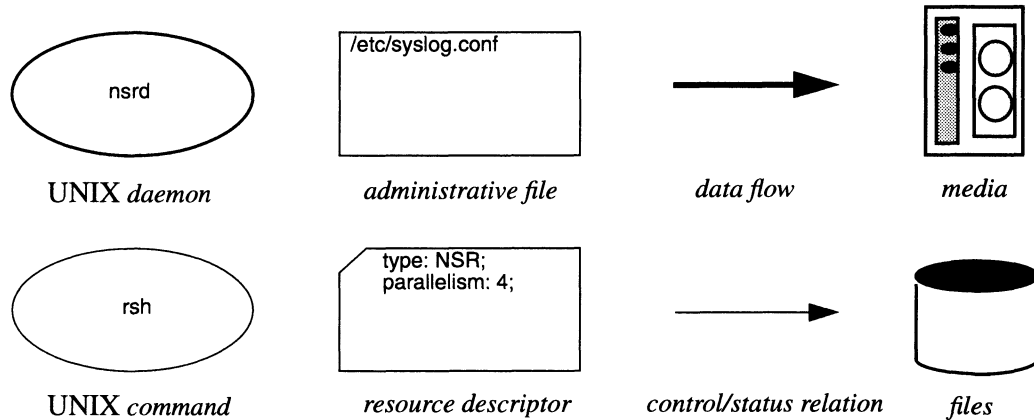
Introduction

The NetWorker man pages provide complete and detailed information that will help you administer NetWorker. The man pages list the commands in alphabetical order, so you may verify the commands 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 will reference the various NetWorker man pages that contain details concerning each section. The man page equivalent to this appendix is **nsr**.

You should read this appendix while sitting at your workstation so that you may read the UNIX and NetWorker man pages, execute NetWorker commands (especially **nsradmin**) and verify the pictures via UNIX administrative commands like **ps**).

The basic picture elements used in this appendix are shown below:



NetWorker Media Format

NetWorker's media format is designed to meet the following requirements:

- ❑ 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.
- ❑ 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.
- ❑ Provides special handling for certain types of files via the Application Specific Module (ASM) architecture.
- ❑ Ability to track media through media labels.

No other existing media format supports these facilities.

NetWorker's media format is fully documented in the **mm_data** man page.

Third parties are welcome to adopt the format.

The NetWorker System at Rest

There are four daemons which continue running on the NetWorker server machine even when it is at rest. The two internal server daemons, **nsrmmd** and **nsrmddb** 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**.

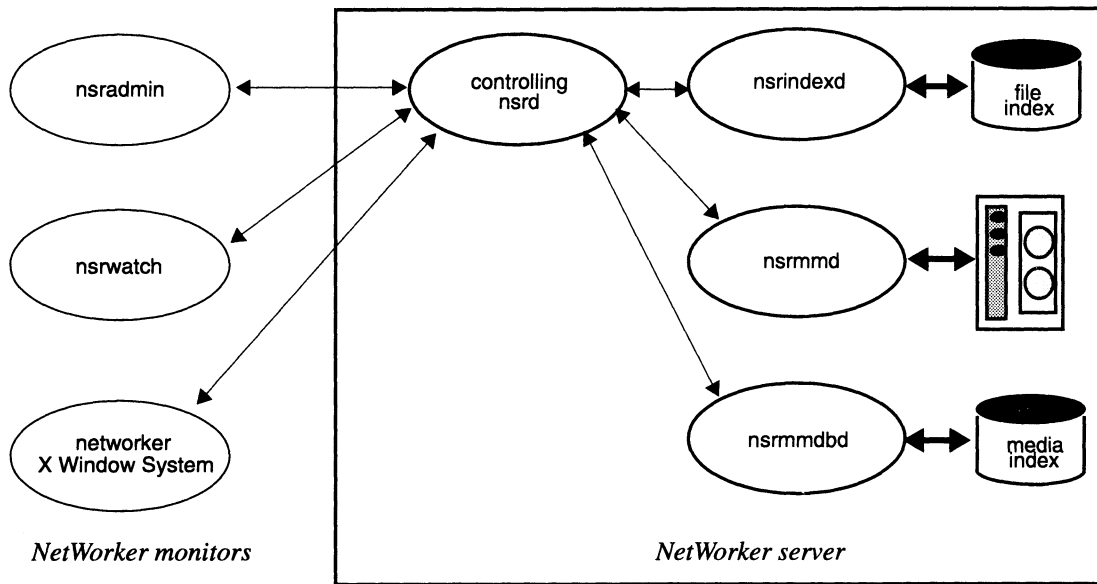
The controlling **nsrd** is the root of all NetWorker functionality. It receives status information from the internal daemons and reports the status to programs such as **nsrwatch** and **networker**. The **nsrd** daemon allows the browsing and editing of NetWorker resource descriptors by clients using the **nsradmin** program. 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, NetWorker's 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 in order to make progress.

A

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 than 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 below:



The relevant man pages are shown below:

Section	Name	Comments
8	nsradmin	controls all aspects of NetWorker server administration
8	nsrwatch	character-based monitor
8	networker	X Window System interface for NetWorker system administration
8	nsrd	NetWorker server
8	nsrmmd	internal daemon interfaces to all backup devices
8	nsrmmdbd	internal daemon that manages the media index

Browsing and Modifying Resources

There are two ways to configure NetWorker – by using the X Window System interface and by editing the configuration information (called *resources*) from an ASCII terminal.


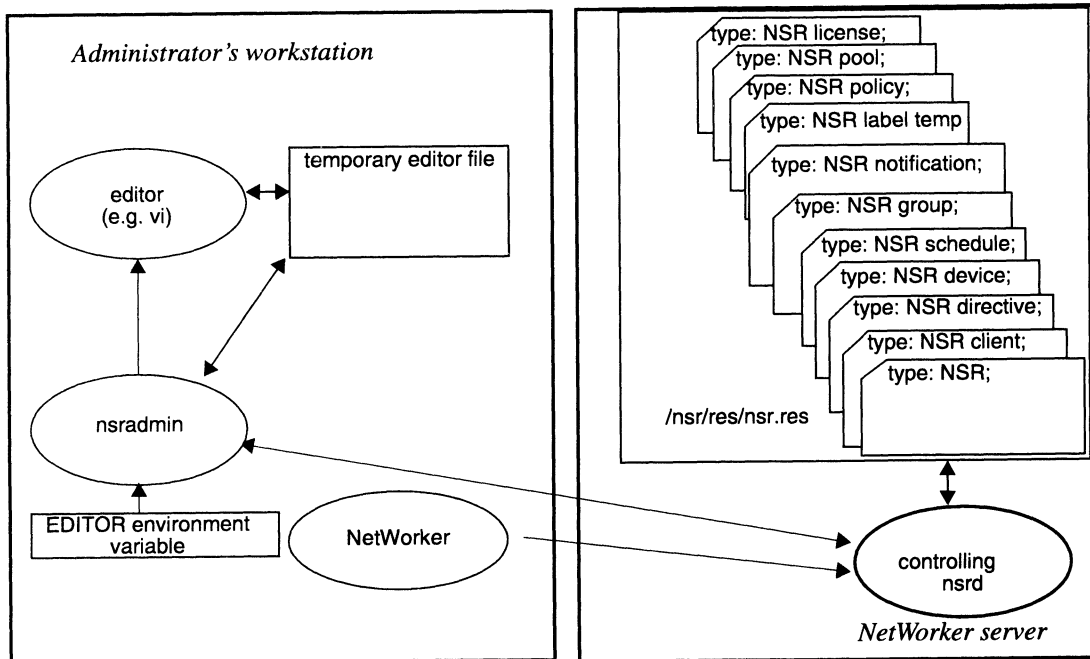
If you are not using the X Window System, you can browse and edit the NetWorker server's resources with the **nsradmin** command. You can either use your UNIX editor to browse and modify the NetWorker server's resource descriptors, or you can use the **curses** mode of **nsradmin**. The **curses** mode uses the **curses** library and */etc/termcap* to give you a visual display of the resources.

When using your editor, **nsradmin** copies the resource descriptors from the server into a temporary file on your machine. 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 machines 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.

A

The resource types that describe the NetWorker server configuration are shown below:



The relevant man pages are shown below:

Section	Name	Comments
8	networker	X Window System interface
8	nsradmin	invokes your editor on resource descriptor copies
1	setenv	how to make nsradmin choose an editor different from vi
5	nsr_resource	describes resource file format and types
8	nsrd	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.” The **nsr** man page documents the use of directives.

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.


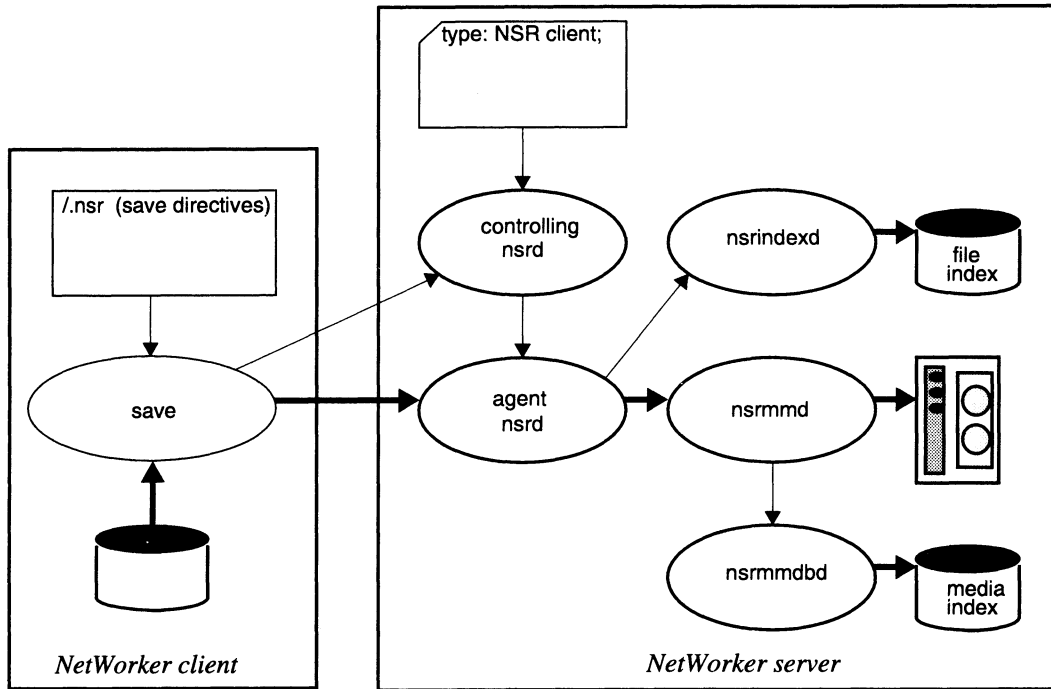
After session establishment with the server, **save** reads files and sends a save set to its agent **nsrd**. Upon completion, the client calls upon the server to commit the save set. To commit the save set, the agent **nsrd** instructs the **nsrmmmd** to commit the save set data to the backup media. 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” program. 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.

A

The illustration below describes a typical **save** session between client and server:



The relevant man pages are listed below:

Section	Name	Comments
5	nsr	documents save directives
5	nsr_client	documents the resource descriptor type "NSR client"
8	save	the command that saves files to the server
8	nsrd	NetWorker server daemon checks access control
8	nsrexecd	NetWorker client execution daemon
3	rcmd	discusses privileged ports
8	nsr	security section discuss policies and issues
8	nsr_data	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 UNIX user. Since file attributes are kept on-line, 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** command.

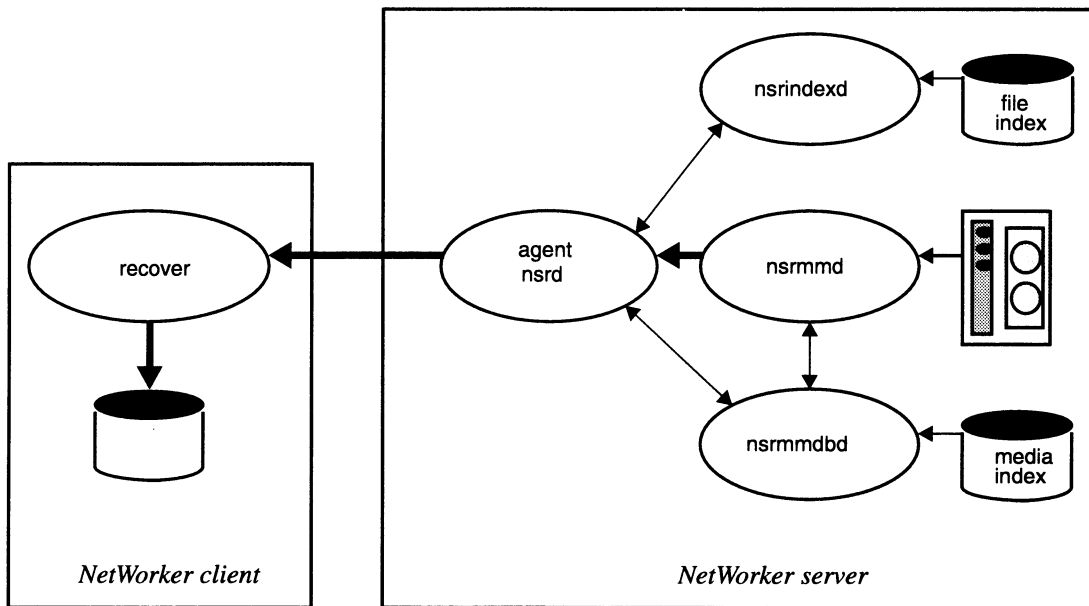
As users browse, they may build a recover list. Eventually the user may issue the **recover** command and the client submits the recovery list to the agent **nsrd** on the server. At this point the file and media indexes are used to determine the backup volume(s) and position(s) of the desired file(s).

The **nsrmmd** will request backup volumes and read them until the entire file recovery list is processed.

Clients browse the on-line indexes via remote procedure calls to their agent daemons. The device and **nsrmmd** are involved only when files are actually requested for recovery. Thus, NetWorker supports multiple concurrent browsers.



A typical **recover** session between client and server is illustrated below:



A

The relevant man page is shown below:

Section	Name	Comments
8	recover	documents browsing and recovery subcommands

Backing up the Whole Machine

While it is nice that the **save** command allows individuals to back up their directories, an unattended backup system has to back up whole machines according to predefined backup levels. The **savefs** command accomplishes this.

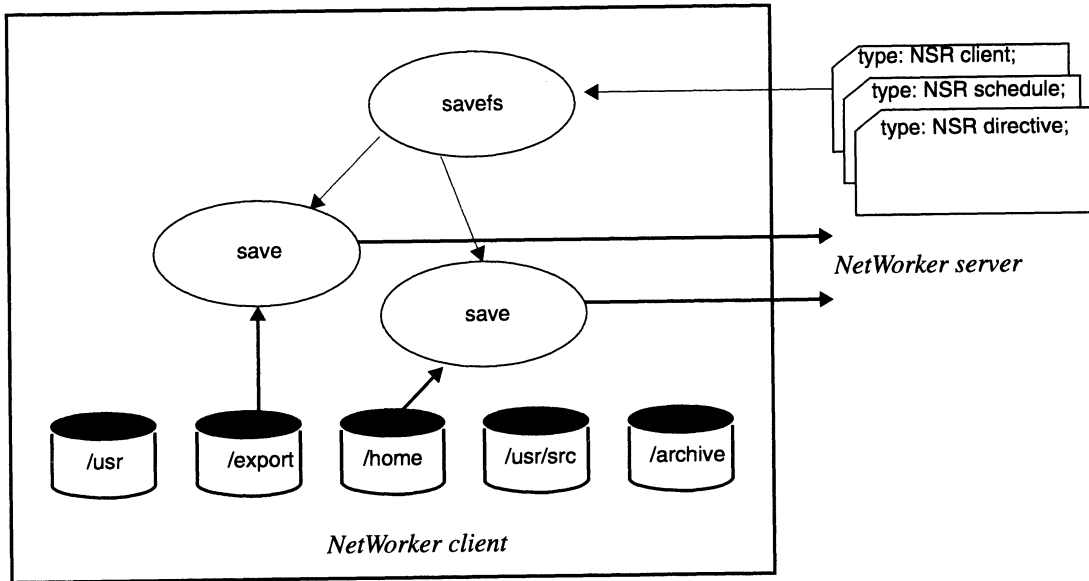
The **savefs** daemon determines which backup schedule to use for a client by looking for the schedule name in the "NSR client" resource. It then consults the appropriate "NSR schedule" resource to determine which backup level to use (full, incremental, or a level [1-9]). Given the save level and a filesystem, **savefs** saves only files within the filesystem which 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.

A

The `savefs` command is shown saving two filesystems in parallel in this illustration:



The relevant man pages are listed below:

Section	Name	Comments
8	savefs	saves the client's filesystems
5	nsr_client	determines the client configuration and which filesystems are backed up
5	nsr_directive	determines which save directives apply to the client
5	nsr_schedule	determines save level
8	save	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.


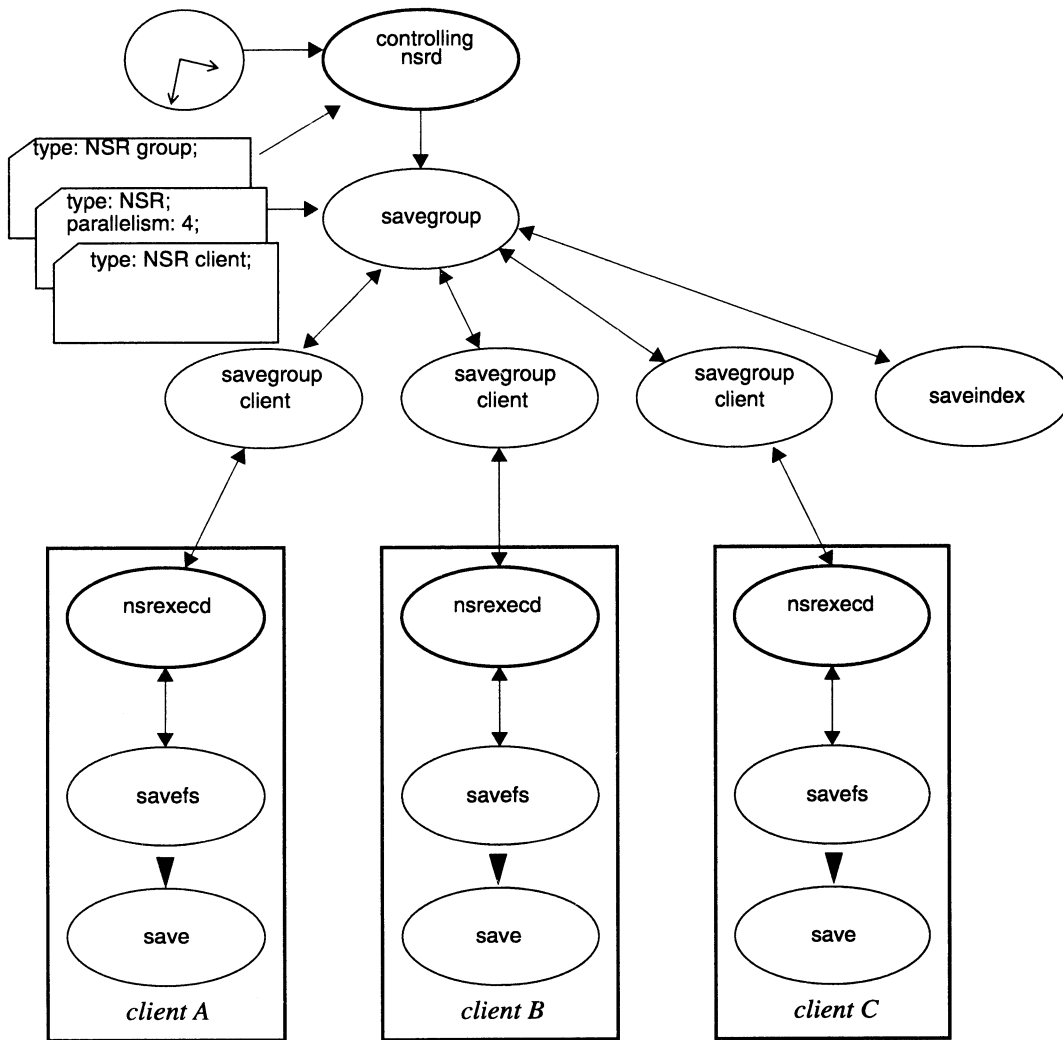
Upon completing all client saves, **savegroup** invokes the **saveindex** command to insure that the NetWorker server index is safely backed up.

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.

A

The NetWorker server **nsrd** starts the nightly save by invoking **savegroup**, as illustrated here:



The relevant man pages are listed below:


Section	Name	Comments
8	savegroup	conducts the nightly network saves
5	nsr_service	parallelism attribute controls how many clients save at once
5	nsr_group	selects the time of nightly saves
8	nsrexecd	remote execution system used to start savefs
8	rshd	remote execution system used to start savefs if nsrexecd is not present
8	savefs	saves each client's filesystems
8	saveindex	saves the server's index; detailed in a later section

Hints for Efficient Network Saves


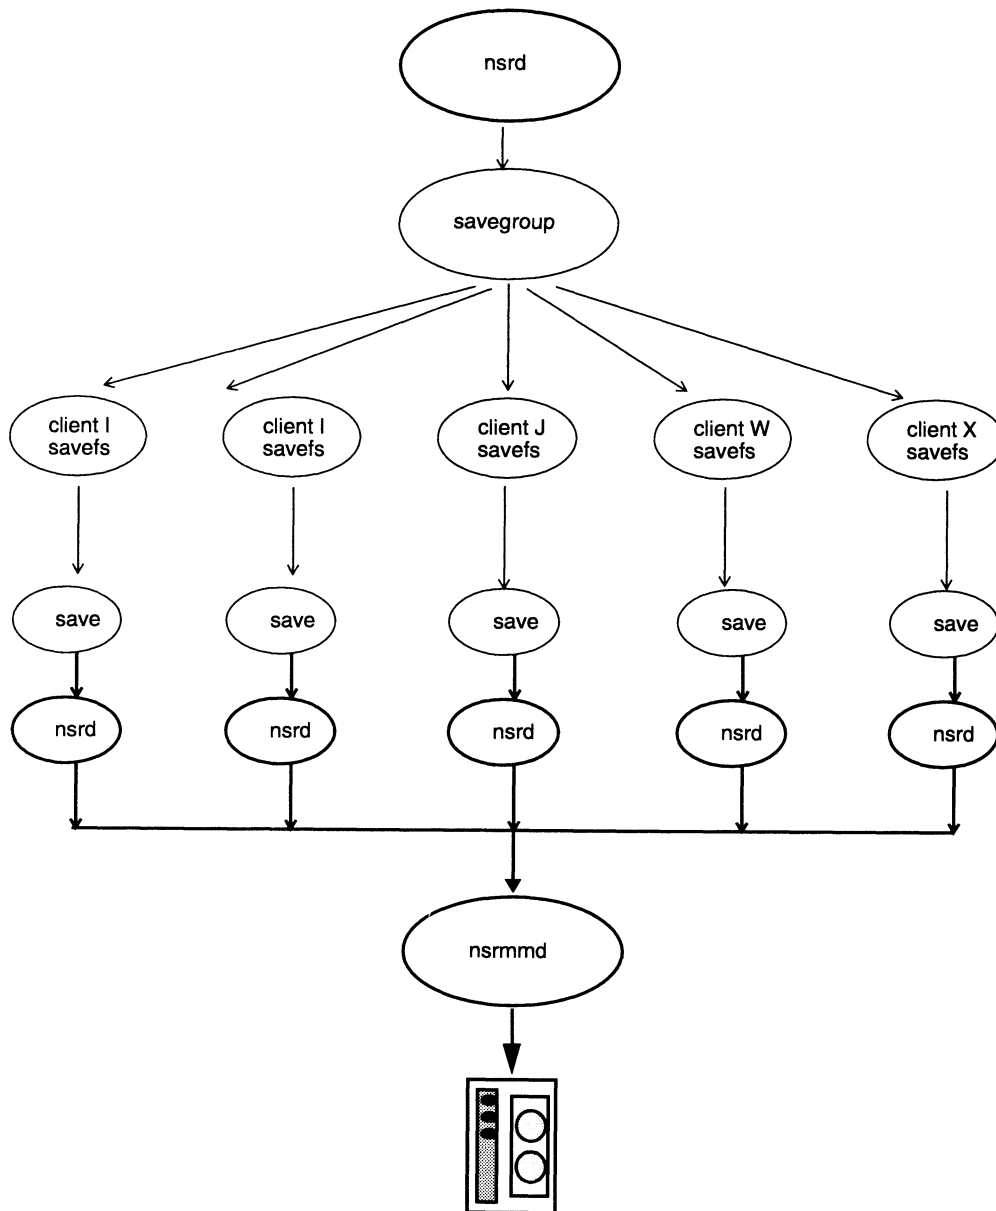
1. Make liberal use of incremental save levels. These are very efficient, since they take minimal backup media space and run very quickly.

NOTE: The daemon **saveindex** does a level 9 backup of the indexes to promote faster recovery after a disaster.

2. 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.



A nightly network backup is shown below. (Not shown are the server's controlling **nsrd** and **nsrmmdbd**, and the client's **nsrexecd**.)



Event Processing

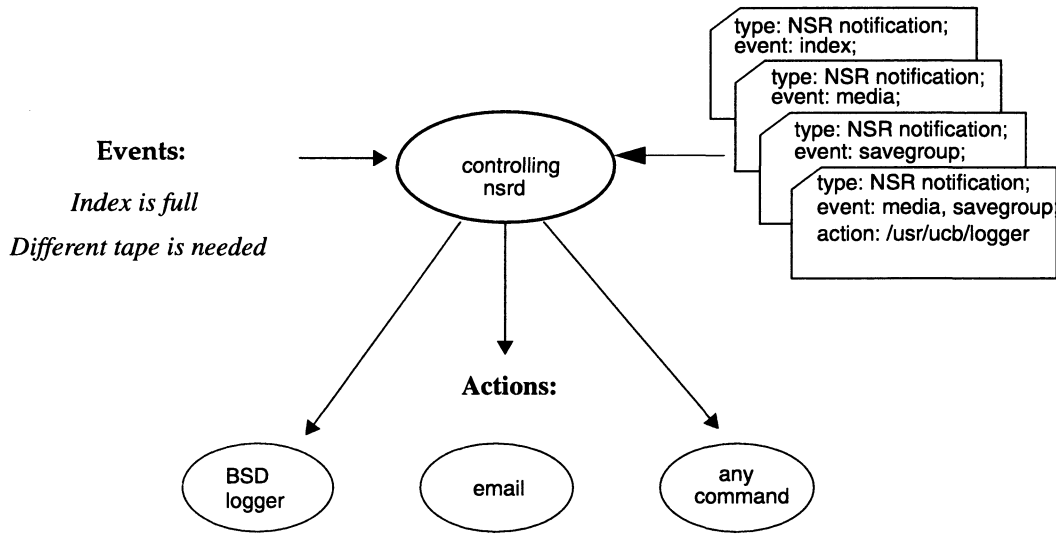
Besides reporting routine status to the controlling **nsrd**, the NetWorker subsystems are set to report major events. The administrator can cause any UNIX command to be invoked when an event occurs. He or she does this simply by modifying a "NSR notification" resource on the NetWorker server.

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 UNIX command. NetWorker defaults provide interfaces to the BSD **syslog** subsystem and electronic mail.

A

The notification subsystem is shown below:



The relevant man page is listed below:

Section	Name	Comments
5	nsr_notification	describes the notification process


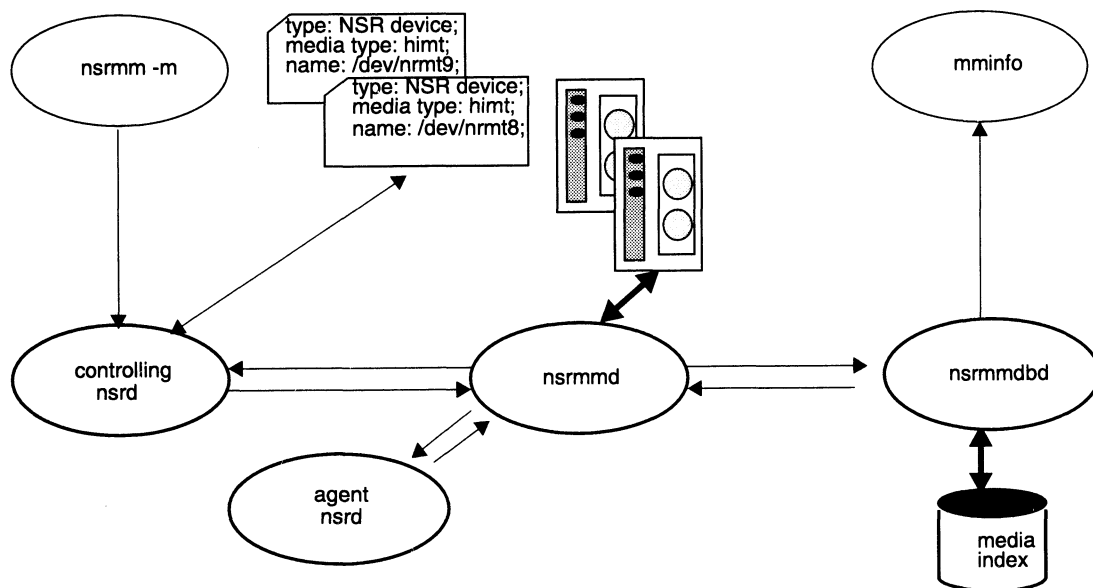
Media Management

There is a working relationship among **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 on-line. 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 file-mark is written.

During recovery, **nsrmm** 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.

Illustrated here is **nsrmm** and its relation to other processes:



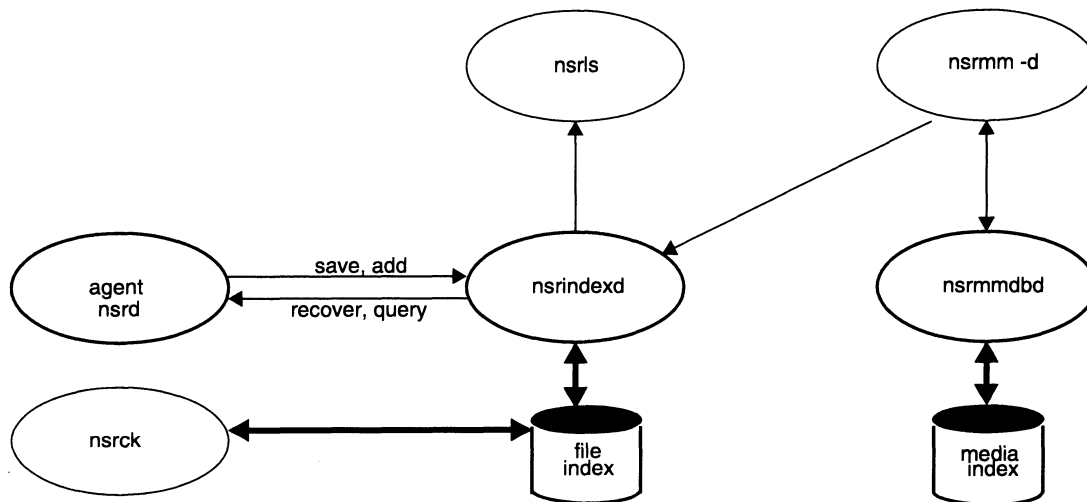
The relevant man pages are listed below:

Section	Name	Comments
8	nsrmm	media manager mounts, unmounts, deletes backup volumes
5	nsr_device	describes and names backup devices
8	nsrmm	writes and reads data to and from the backup devices
8	nsrmmdbd	manages the backup media library
8	mminfo	displays information about the backup media library

Index Management

The on-line index is built during saves and 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** is automatically invoked after failure (for example, an operating system crash) to guarantee the index database consistency before the NetWorker service is enabled.

The file index and its relationship to other processes is illustrated here:




The relevant man pages are listed below:

Section	Name	Comments
8	nsrmm	media manager deletes backup volumes
8	nsrindexd	manages the on-line index
8	nsrls	displays index usage statistics
8	nsrck	rebuilds indexes after a hard crash
5	nsr_layout	explains where all the NetWorker server's files are located


Disaster Recovery

The last act that **savegroup** performs each morning is to back up the NetWorker server's own index in such a manner that the index can be recovered without using itself. Besides saving the index, **saveindex** sends *bootstrap* recovery information to the line printer. Should the server's index ever need recovering, 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 away index information vital to the server itself and to print recovery parameters to the line printer. Should the server's indexes ever need recovering, these parameters are fed to the **recoverindex** command. Once the server's indexes are recovered, all other files may be recovered in the normal manner.



This illustration shows the progression of index information from saving to recovering:



The relevant man pages are listed below:

Section	Name	Comments
8	savegroup	typically invokes saveindex
8	saveindex	also covers recoverindex
8	scanner	used by recoverindex; reads raw NetWorker backup volumes
5	mm_data	describes backup media format
8	nsr_crash	more general information on recovering after a disk crash

Appendix B

Managing the NetWorker Environment

This appendix provides examples and suggestions for you to consider while you are thinking about setting up your NetWorker environment. It also offers background information to help you understand the logic behind NetWorker's backup schedule and index policy features.

NetWorker Backup Schedules

The NetWorker server backs up each client system across your 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 very simple or very sophisticated, depending on the needs of your environment. All clients can share the same schedule, or each client can have its own unique schedule. This section discusses some of the considerations you should keep in mind while determining which schedule best fits your situation.

Backup Levels

A backup schedule specifies what level of backup NetWorker will perform for a client on each day of a weekly or monthly period. NetWorker offers eleven different backup levels:

- ❑ Full – backs up all files, regardless of whether or not they have changed since the previous backup.

- ❑ Level 1 through level 9 – back up files that have been modified since a previous full backup or a backup of a lower numbered level. For example, a level 3 backs up all the files that have changed since the previous level 2, level 1, or full backup.
- ❑ Incremental – backs up all files that have changed since the previous backup of any level.

You can also skip a backup on a given day. You may want to schedule a “skip” backup on weekends or holidays when no one is available to load backup media.

NetWorker's on-screen calendar presents an easy method for setting up backups for each day of the month. You can designate a schedule and repeat it over a weekly or monthly period. For example, if you set up a full backup for one Friday, NetWorker automatically sets up a full backup for every Friday. Or, you can override the regularly scheduled backup level for a specific day.

There is no “correct” way to set up a backup schedule for a particular client or network of clients. The clients you need to back up will probably vary considerably – some will have a lot of critical data to back up, some may have a small amount of data that does not change very often. Consider the situation for each client, weigh the benefits of the different backup schedules, and then select the best schedule for each client.



Full Backups versus Incremental Backups

If your site has a small number of files, you may choose to perform a full backup every day, or perhaps once a week. This is a simple schedule to set up and execute, and it makes recovering from a disk crash easy – you simply need the last full backup volume.

The situations you should 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 will have to be monitoring the backup and change the media.
- ❑ Full backups cause the on-line indexes to grow more rapidly than do incremental or level backups.

You may 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.

Using Level Backups

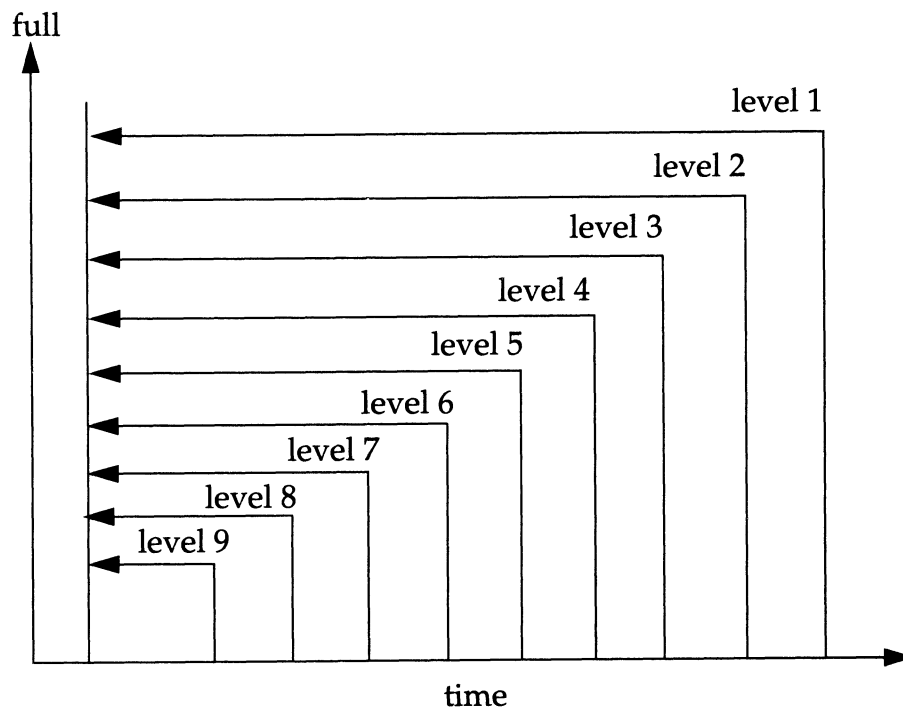
You can use level 1 through level 9 backups to moderate between the two extremes described above. Level 1 through level 9 backups allow you to set up a schedule for each client that balances your need for small, fast backups that do not take up too much index space and the need to recover quickly and easily from a disk crash.

A level backup serves as a checkpoint in your schedule since it collects into a single backup session all the files that have changed over many days or even weeks. Without a level backup, these files would be spread across tapes from many different

backup sessions. As a result, a level backup can simplify and speed file recovery.

To illustrate the effect of level 1 to level 9 backups, consider two examples. In the first example, a full backup takes place on the first day, followed by a level 9, level 8, level 7, and so on down to a level 1 backup over time.


A full backup followed by level 9 to level 1 is illustrated below:



The advantage of this schedule is that to recover from a disk crash, you only need two tapes: the one with the full backup, and the one with the last level backup. The disadvantage is that with each day, there are more changed files to back up, so the backups will take longer to complete.

The figure below illustrates a backup schedule that also starts out with a full, but the level backups which follow are in reverse order: starting with a level 1 on the first day following the full, on down to a level 9 backup. Each day, the backup will only back up the files which have changed on that day.

A full backup followed by level 1 to level 9 is illustrated below:



The advantage of this schedule is that each day's backup will be small and will complete in a short period of time. The disadvantage is that recovering from a disk crash will require the full backup tape and all of the level backup tapes up until the day of the disk crash.

Neither of these backup schedules is practical. They simply illustrate how level backups work. The real power of level

backups comes into play when you combine multiple levels along with fulls and incrementals.

A Typical Monthly Backup Schedule

Sites with even a few gigabytes of files to back up often choose a monthly schedule based on fulls, 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 while also making it relatively easy to recover in the event of a disk crash. This schedule offers several advantages. First, the level 5 backups simplify recovery. Assume that a disaster strikes on the 24th of the month. All the files that you need to recover an entire client system are located on tapes from just five backup sessions:

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

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

B

Illustrated below is level 5 and incremental backups after a full:



Backups Take Time

The amount of time you have to complete a backup on any given day will also influence the schedule that you decide to use. Thanks to flextime and around-the-world operations, many networks must be up and running for users from early in the morning until very late in the evening. And while NetWorker is able to back up live filesystems, most administrators want 100% of their network and systems capacity ready for users during work hours. What number of files can NetWorker back up in, for example, a four hour backup window? The answer depends on hardware, configuration, and backup schedules.

- ❑ Select a backup server with enough CPU power, memory, and bus bandwidth so that the backup server is not the bottleneck.
- ❑ Leave NetWorker's parallelism feature turned on. This feature causes multiple client systems to send their files to the backup server in parallel. This keeps a stream of files ready for the tape drive, so that it does not start and stop.

B

- ❑ Experiment with compressing files on the client systems to reduce the size of the data that has to be written to tape. Compression may speed your backup as long as the client systems are still able to supply files to the backup server fast enough to keep the tape drive streaming.
- ❑ Take advantage of NetWorker's ability to skip over specified files during the backup. For example, you could choose to skip over *core* files.
- ❑ Add a second backup device to your backup server. NetWorker can simultaneously back up to more than one device.

If your backup server is able to drive a single 8mm tape drive at an average of 400 KB/second (its maximum speed is 500 KB/second and some time is invariably lost loading the tape or rewinding, for example) you will be able to back up a maximum of 5.76 GB in four hours. If you have more than this amount of data to back up, then full backups will be limited to weekends and holidays when users will not be affected.

Where a backup requires loading and unloading of tapes, you must schedule it when an attendant is available. Many administrators find that an incremental backup of their network fits onto a single 4mm or 8mm tape, but they must schedule multi-tape full and level backups for specific nights or weekends when an operator is on duty to load additional tapes. If an operator will not be available over a holiday weekend 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.

Staggering the Backup Schedules

Networks with a large number of files can take a very long time, and require a lot of loading and unloading of tapes to complete a

full backup. There may not even be time in a night or an entire weekend to complete a full backup of all the systems across a very large network. An easy way to handle this problem is to stagger the clients' backup schedules. Rather than have every client system perform a full backup on Monday and incrementals the rest of the week, for example, you can schedule some clients to perform a full backup on Tuesday and others on Wednesday.


NetWorker goes one step further to smooth the backup load for very large client systems. With NetWorker you can assign a separate backup schedule to each filesystem. Each file system, in essence, is treated as if it is a completely separate client.

Convenience versus Security

NetWorker is a very sophisticated product. You may leave the same backup volume mounted in the server's backup device throughout a week or month, and when it becomes full, replace it with a new labeled backup volume. NetWorker tracks all the backups, no matter what day of the week or month, or what part of the backup schedule cycle is in effect. The same backup volume may contain full, level [1-9], or incremental backups, and 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 may lose a day of work. In the worst possible case, if the backup volume with the full backup is destroyed, users may lose all the work done since the last full backup. Therefore, some



administrators prefer to remove the backup volume used for a full backup 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, and the other volumes that contain the most recent level [1-9] and incremental backups.

Backup Schedule Examples

For your convenience, NetWorker is shipped with five pre-configured backup schedules. If these schedules fit your backup requirements, you can use them “out of the box.” Or, you may delete them and create new ones to accommodate your site-specific needs.

This section explains the logic behind each schedule. After understanding how they work, you may want to use them as examples to set up your own schedules.

Important ⇔

You are not allowed to change the name of an existing schedule. For example, if you want to *change* the schedule “Full Every Friday” to “Full Every Monday,” you must delete the “Full Every Friday” schedule and create a “Full Every Monday” schedule. You cannot change the existing schedule to complete full backups on Mondays instead of Fridays, and then edit its name.


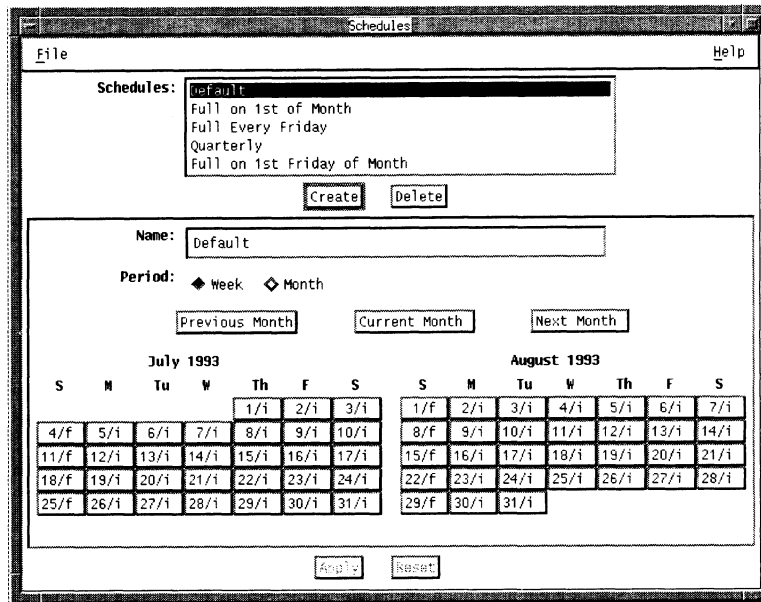
The most efficient way to best protect the systems from file loss *and* maintain control over the number of backup volumes is to follow full backups with level [1-9] and incremental backups.

Default

The Default schedule is the only one you may not delete. It is a weekly schedule, and completes a full backup every Sunday, followed by incremental backups all other days of the week.

This schedule is convenient if you want to pre-mount 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.

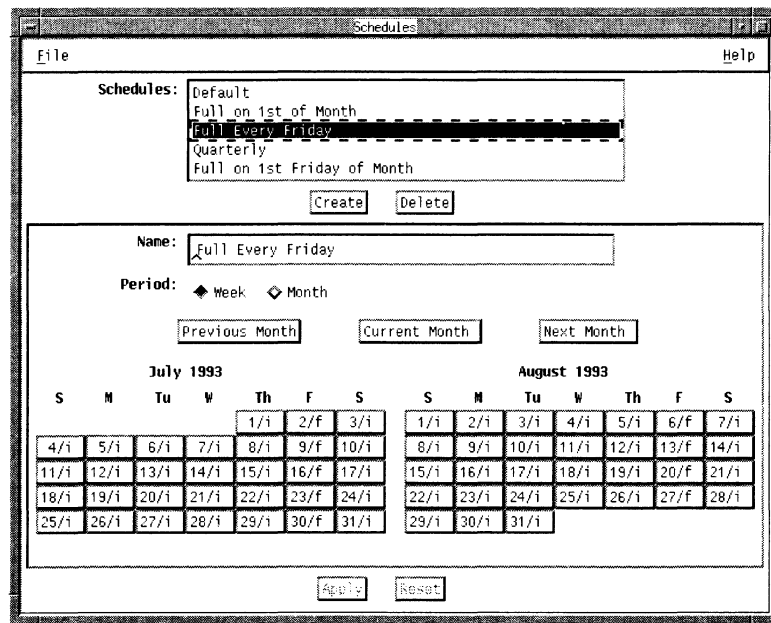
The Default schedule is shown below:



Each time you use the Schedules window to create a new weekly backup schedule, this pre-configured schedule will appear in the calendar as your starting point.

Full Every Friday

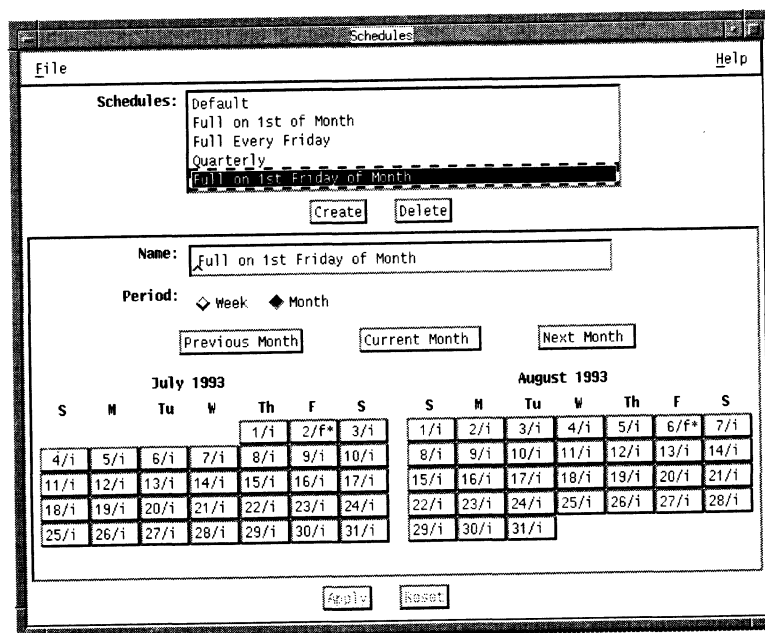
This weekly schedule completes a full backup every Friday, followed by incremental backups the other days of the week.



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 1st Friday of Month

This monthly schedule completes a full backup on the first Friday of the month, (not the first calendar day of the month). Incremental backups take place on all the other days.



The advantage of this schedule is that you complete a full backup only once a month. If you use this schedule, it would be a good idea to store the backup volume with the full backup in a safe place, and use other backup volumes for the incremental backups. It would also be a good idea to change backup volumes every few days for the incremental backups. If you allow all the incremental backups to be stored on one backup volume, and it is destroyed near the end of the month, you 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* (like Friday, in this example),

you must set the overrides in each month. (Notice the "f*" in the first Friday of each month.) This is because the first weekday (Monday through Friday) in a month may fall on any calendar day from 1 to 7.

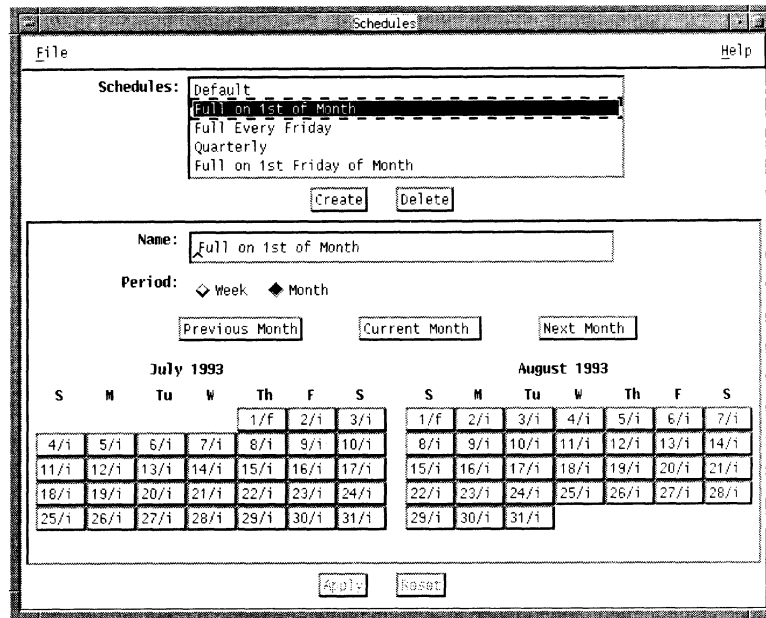
NOTE: After 1993, you will need to add the overrides for 1994. In other words, the overrides do not carry over from one year to the next.

Full on 1st 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 manually set any overrides.

**B**

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



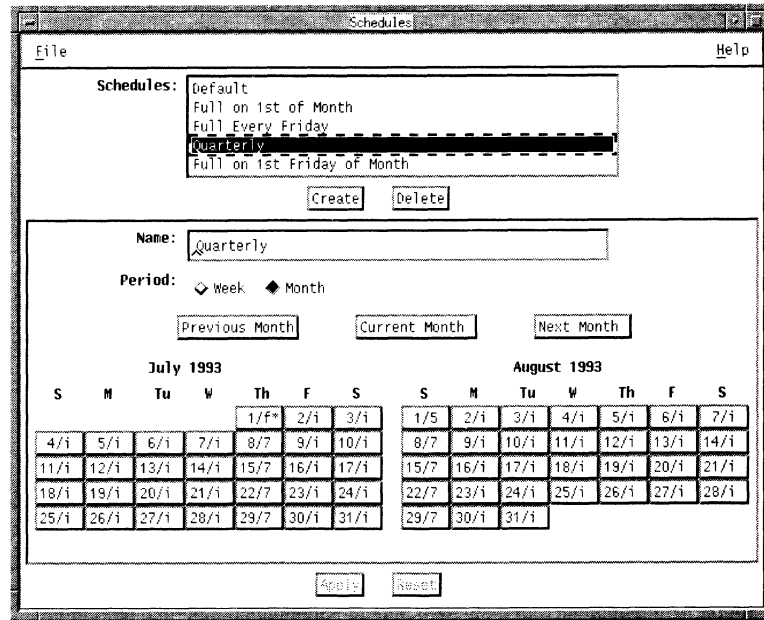
Quarterly

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.

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

B

data that has changed since the first day of the month. The daily changes are protected by incremental backups.



If you use this schedule, it is a good idea to segregate the backup 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. 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.


NetWorker Browse and Retention Policies

NetWorker maintains on-line indexes of all the files backed up for each client and an index of the files stored on each piece of media. NetWorker lets you set policies that automatically control how long the information is retained in these on-line indexes. This section explains NetWorker's browse and retention policies and the trade-off between providing faster, easier recovery for your users versus conserving disk space.

Browse Policy

One of Legato NetWorker's popular features is the ability of 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 on-line index (about 200 bytes each). Since disk space is limited, you need to establish a policy of how far back in time you will keep information about backed up files in the indexes.


The browse policy that you select specifies how long the entries for your files will 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.



How Browse Policies Work

To recover a complete directory or file system you often need to recover some files from incremental and level backups as well as from a full. The incremental backup is dependent on the level backups and, in turn, on the full. NetWorker will not delete the entries from any backups on which other backups depend. As a result, you may find that entries are deleted later than you might expect.

In the illustration shown below the browse policy is set to one week, which happens to equal one complete backup cycle.



B

NetWorker will not remove the first full backup from the on-line file index until all the incremental and level 5 backups that depend on it have expired. As a result, the full backup will actually stay in the on-line index for a period of time equal to the browse policy plus one full backup cycle.

The first full backup will not be removed from the on-line index in exactly one week, however, because there are incrementals and a level 5 backup that depend on the full which have not yet expired. Each incremental backup will be removed from the on-line index one week from the time it was completed. The level 5 backup will be removed one week after the last incremental that

depends on it is removed, and then the full backup will be removed at that same time.

The rule to remember is that a full backup will actually remain in the on-line 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 on-line and browsable in the file index you must also use for all incremental and level backups. With NetWorker you manage backup cycles (the period from one full backup to the next); you do not independently manage different levels of backups.

Reclaiming Disk Space

NetWorker automatically reclaims disk space that is freed up when entries are deleted from the on-line file indexes. However, the space is not returned immediately to your system.

NetWorker takes some time, processing power, and swap space in order to reclaim this space and to have this constantly taking place on your backup server would be inefficient. Instead NetWorker first reuses this space to store information about new files that are backed up. When the file index for a client reaches a point where less than 50% of its space is being used by files that have not reached the end of their browse period then NetWorker automatically invokes a process that returns the space to your system.

You may also reclaim disk space at any time by using the Reclaim space button in the Indexes window.

B

Recovering Files Removed from the Index

You can recover files whose entries have been removed from the on-line index because they have passed the Browse policy period as long as the files are still stored on a backup volume. The recover process will not be as easy as when the entries are still in the on-line index, however. Here are the basic steps for recovering a file whose entries have been removed from the on-line index:

- ❑ Use the Volumes window to find out the name of the backup volume that contains the save set.
- ❑ Use the **mminfo** command to determine the save set id. Use this syntax:

```
mminfo -v -s server -c client -N saveset volume_name
```
- ❑ Rebuild the file index entries for the save set using the **scanner -i -s save_set_id#** command at the system prompt. Enter the save set id number determined above for save set id#. Rebuilding the file index using the **scanner** command may take some time.
- ❑ Use the NetWorker Recover window to identify the needed file(s) and initiate the recovery.

Recovery is considerably easier if the file information is still in NetWorker's on-line index. That is why you want to set a browse policy that is long enough to cover most recovery requests.

Media Retention Policy

Your need to conserve disk space may lead you to establish a short browse period. NetWorker's 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.

Remember, 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's browse and retention policies combine to give you a hierarchy of recovery capability while keeping the disk space needed for the on-line 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.


B

Policy Examples

NetWorker is shipped with five pre-configured policies: Decade, Month, Quarter, Week, and Year. Use these policies to choose the length of time to retain the entries in both the file index and media index. Remember, the *retention* policy you select will affect the size of the *media* index, and control the length of time NetWorker will track the backup volumes and the data on each volume.

The *browse* policy will affect the size of the *file* index, and the length of time that NetWorker will retain entries for every file backed-up and visible in the Recover window. You must always choose a retention policy which is greater than or equal to the browse policy.

For more detailed information on browse and retention policies, refer to the section "Creating Index Policies" in Chapter 3.

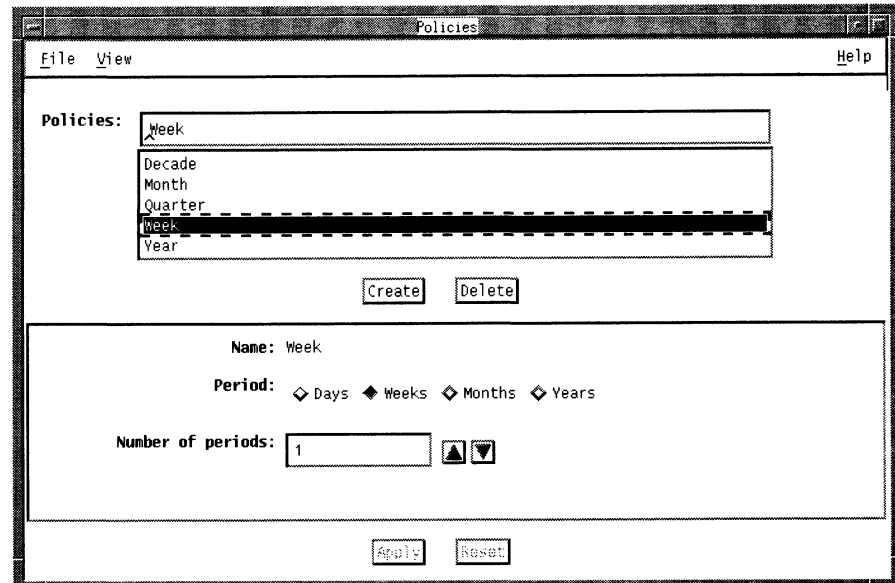


For example, if you choose Quarter for the retention policy for a client, and Month as the browse policy, the client will be able to browse all the file entries for backed-up files dating back to a month. Each month the oldest entries for the client's files will automatically be removed from the server's file index. However, the backup tapes which contain the data for the files will still be tracked by NetWorker in the media index.

Week Policy

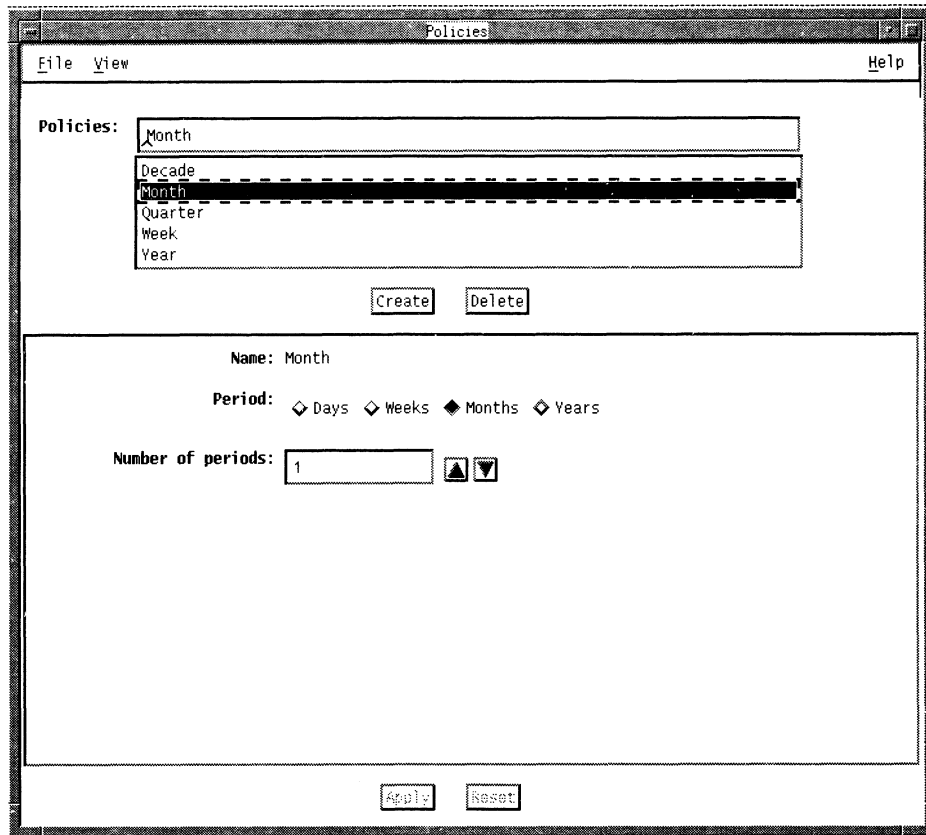
This policy maintains the file index entries or the media index entries for one week. If you use this as a browse policy, the users will only be able to view and mark files for a recovery which go back in time for a week. As a browse policy, it is useful when you have a limited amount of disk space and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

As a retention policy, Week means that your backup volumes will turn over quickly, and NetWorker will recycle through the tapes at a faster rate. Use this policy if you schedule weekly full backups, and need only to keep backup data for one backup cycle plus a week.

**B**


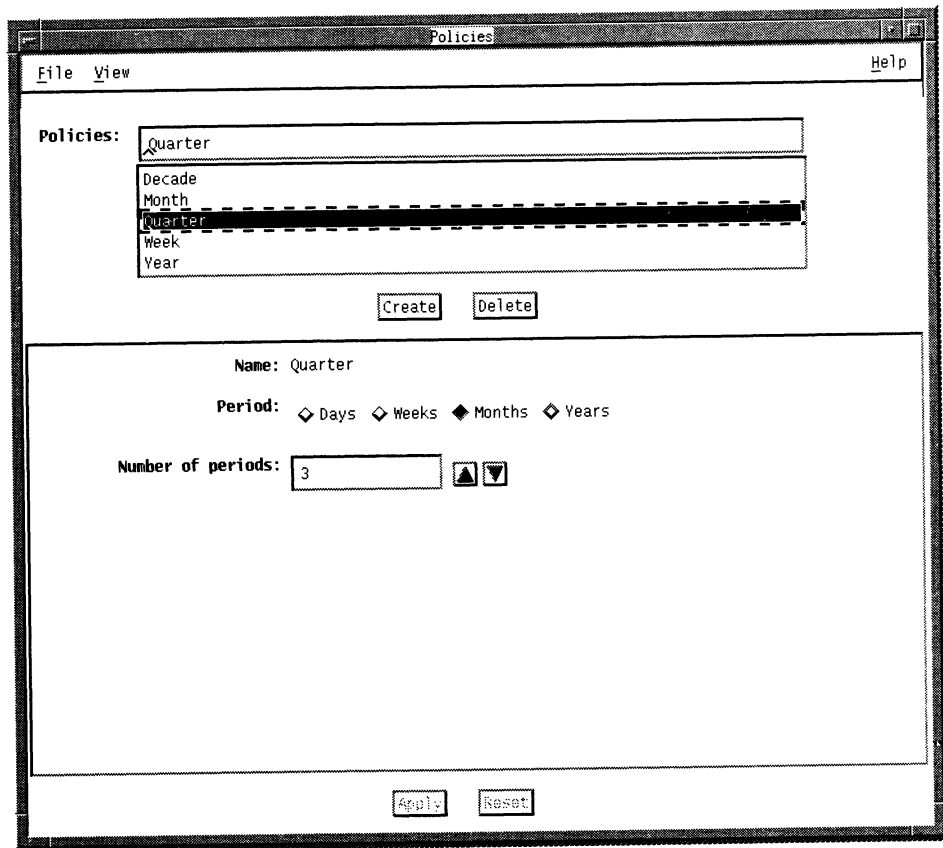
Month Policy

This browse policy allows users to view and recover versions of files dating back to at least a month. The Recover window would display versions for files backed-up for one full month plus a number of weeks. As a retention policy, NetWorker will maintain and track the backup volumes for one full backup cycle plus a month.



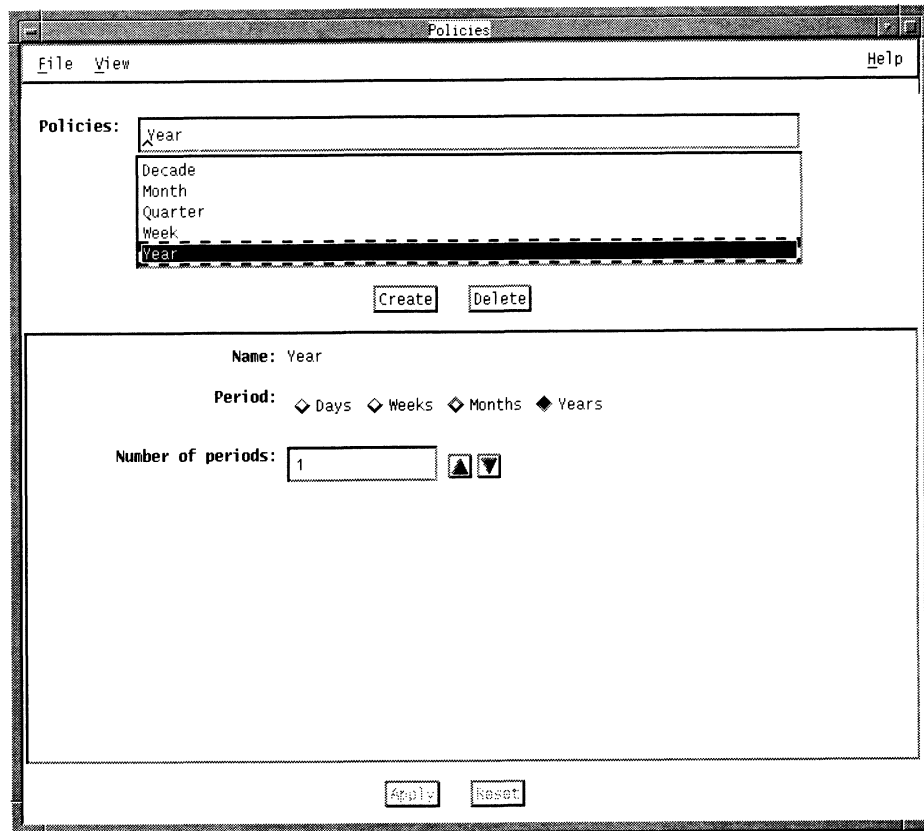
Quarter Policy

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

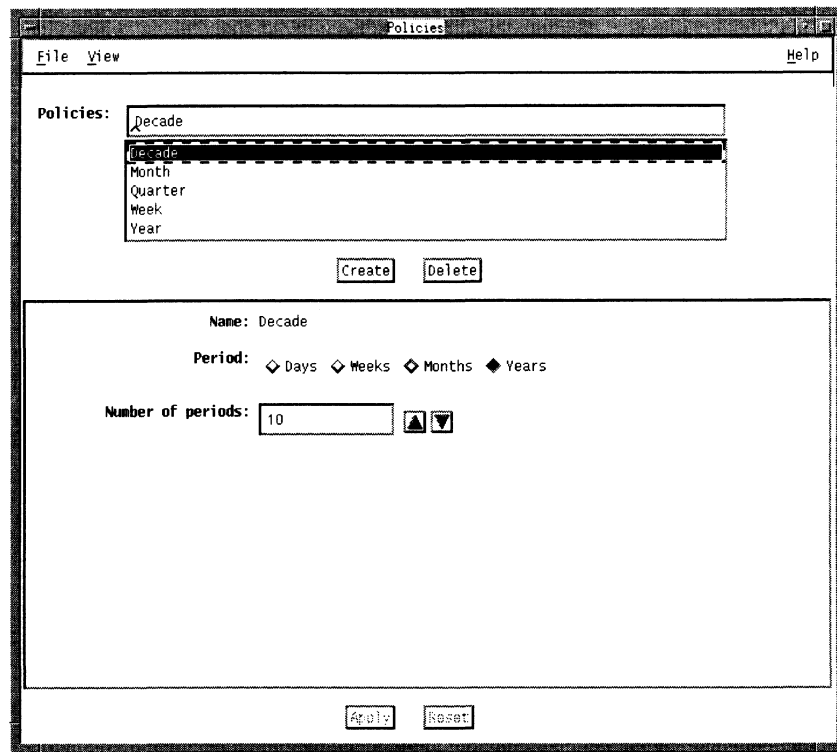
If you need to keep backed-up data on-line for several months, use the Year policy. For example, if your company requires ready access to information going back in time for at least three quarters, this is a good browse and retention policy. Realize, however, that NetWorker will require more disk space to maintain all the information on-line.

**B**

Decade Policy

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

Your NetWorker server will require *lots* of disk space for the on-line indexes if you choose Decade for your browse policy. Depending upon how much data you are backing up, ten years of file index entries could take up gigabytes of disk space.



It would make more sense to use Decade as the retention policy and use Quarter or Year as the browse policy. NetWorker can then track the backup volumes and the data on each one. You would always be able to retrieve data from an old backup volume 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.

Summary

There are no right or wrong answers to the questions of how to configure NetWorker. The challenge is to understand how to best take advantage of the power and flexibility that NetWorker offers for your specific environment. You should get started with NetWorker by using the pre-configured schedules and policies and then undertake small experiments. As your network of systems grows larger, as there are more and more files to back up, and as users see the advantages of NetWorker's fast file recovery you will need to continue to make adjustments. Fortunately, NetWorker was designed to change and adapt as your needs expand.

Guidelines for Choosing a Configuration

There are several factors that affect the NetWorker server configuration that best suits your backup and recover needs. The configuration consists of the hardware and software, which includes tape drives, 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.

Important ⇨ Please keep in mind that these are *guidelines*, and actual performances may vary.

The goal in selecting a configuration is to balance the different hardware and software limitations to achieve the overall data handling capabilities you require. Start by looking at the limits

of the major NetWorker configuration components: tape drives, network connection, 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.

<u>Drive</u>	<u>Data Transfer Rate</u>
Exabyte 8500	500 Kb/s
Exabyte 8505	650 Kb/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.

<u>Client/1-backup</u>	<u>Data Transfer Rate</u>
PC/DOS/SPX	80 Kb/s
PC/DOS/TCP	150 Kb/s
Sun SS2	200 Kb/s
IBM RS/6000	300 Kb/s



Network

Ethernet has an upper limit on bandwidth of about 1 Megabyte per second, but in practice, most networks can only handle about 500 Kbytes between a set of clients and a single server. Token ring has a lower maximum bandwidth (8 Mbits/s) but a higher utilization so the data transfer rates are approximately the same.

<u>Network</u>	<u>Rate</u>
Ethernet	500 Kb/s
Token Ring	500 Kb/s
FDDI	5 Mb/s

Server

The server must be able to handle the load of network packets, data movement, and tape drives in order to achieve the rates listed above. Most of the work on the server side is in data movement, context switching, and interrupt handling. The performance of all of these functions improves as the CPU speed increases. It takes approximately 20 MIPS to handle 500 Kb/s of data, although this tapers off at high CPU speeds because bus bandwidth and other bottlenecks begin to affect the data movement. Approximately 16 Mbytes of memory is required per 500 Kb/s of data rate handled by the server.

<u>Server CPU/Memory</u>	<u>Rate</u>
20 MIPS/16 Mbytes	500 Kb/s
50 MIPS/32 Mbytes	1000 Kb/s
100 MIPS/64 Mbytes	1500 Kb/s

Measuring Performance

If you are interested in measuring the performance of your NetWorker environment, you must take into consideration the server system and the client system.

The factors to consider for the server system are the speed of the backup device, the network speed, and the CPU speed. Factors to consider for the client system are filesystem traversing, generation of data, data on multiple disks, and CPU speed.

The next sections examine these factors in more detail.

Server Performance

This section provides examples on how to measure the performance of the server.

Backup Device Speed

Most tapes have step function in data rate. NetWorker uses 32K bytes/record. To measure tape speed, follow these steps:

1. Create a large file (at least 20 Mbytes) with non-zero data. For example:

```
# cat /dgux /dgux /dgux /dgux /dgux > big
```

2. Use the **dd** command to write the large file to tape four times and measure the time results:

```
/bin/time dd if=big of=/dev/nrst8 bs=32k conv=sync  
/bin/time dd if=big of=/dev/nrst8 bs=32k conv=sync  
/bin/time dd if=big of=/dev/nrst8 bs=32k conv=sync  
/bin/time dd if=big of=/dev/nrst8 bs=32k conv=sync
```

B

3. Divide the file's size by the average of the last three real times.
For example:

```
-rw-rw-r-- 1 root 20675420 Jan 7 11:04 big
95.2 real 13.0 user 11.9 sys
78.2 real 12.9 user 12.7 sys
78.0 real 12.8 user 12.5 sys
76.8 real 13.0 user 12.4 sys
Rate: 20190 Kbytes / 77.66 sec = 260 Kbytes/sec
```

This number gives you the rate of the tape speed.

CPU Speed

The CPU of a server 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 server. A larger MIPS rating means a faster machine.

Memory

The memory on the server limits the amount of data buffered between the NetWorker **save** command, agent daemon, and media management daemon.

Client Performance

This section provides examples on how to measure the performance of the NetWorker client.

Filesystem Traversing

To measure the filesystem traversing speed, follow the steps below:

1. Time the **uasm** command with the **-bi** option. For example:

```
# /bin/time uasm -bi /usr
13931 records 2667396 header bytes 350849148 data
bytes 124.9 real 10.8 user 34.0 sys
```

2. Divide the number of records by real time for rate per file. For example:

```
# 13931 records / 124.9 sec = 111.5 files/sec
```

Data Generation Rate

To measure the rate at which a client generates data for a backup, follow the steps below:

1. Time the **uasm** command with the **-si** option and re-direct the output to **/dev/null**. For example:


```
# /bin/time uasm -si /usr > /dev/null
```

2. Divide the number of bytes obtained (filesystem traversing) with the **uasm -bi** command by the real time generated by the **uasm -si** command. For example:

```
342626 Kbytes / 1199 = 286 Kbytes / sec
```

Data on Multiple Disks

Where filesystems are located on different physical disks, or on a CLARiiON disk array, NetWorker backs the the filesystems up in parallel. To measure parallel disk speeds, follow these steps:



1. Use the **df** or **du** command to find two directories of approximately the same size.
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 *kbytes/sec*) to obtain a combined rate.

This rate reflects the performance of NetWorker backing up data on multiple disks.

CPU Speed

The CPU of a client limits the following:

- The total data throughput to 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 machine.

**B**

Appendix C

Command Summary

This appendix contains a list of the most commonly-used NetWorker commands entered at the system prompt, and a table listing the NetWorker maintenance commands.


If you are not using the X Window System, or are using NetWorker from an ASCII terminal, the following commands are the ones you may use most often:

Use:	To:
nsrmm -l <i>volume_name</i>	label a backup volume with the name <i>volume_name</i>
nsrmm -m	mount the backup volume in the server drive
nsrmm -u	unmount the backup volume in the server drive
nsrmm -d -P <i>volume_name</i>	purge the backup volume named <i>volume_name</i>
nsrmm -d <i>volume_name</i>	delete the backup volume named <i>volume_name</i>
nsrmm -m -R	recycle the backup volume
mminfo	display information about the backup volumes
nsradmin	configure networker resources

nsrck - F	compact the index size after purging or deleting backup volumes
nsrexecd	NetWorker client execution daemons
nsrls	display information about the server's on-line indexes
nsrwatch	display the character-based NetWorker status monitor
savegroup -p <i>group_name</i>	see a preview of the backup of the group named <i>group_name</i>
savegroup <i>group_name</i>	start the backup of the clients in the group <i>group_name</i>

The table below lists the NetWorker maintenance commands:

Command	Description	Manual Page
compressasm	the NetWorker module for compressing and decompressing files	compressasm(8)
logasm	the NetWorker module for saving log files	logasm(8)
mailasm	the NetWorker module for saving and recovering mail files	mailasm(8)
nsrmm	the NetWorker media interface command	nsrmm(8)
mminfo	the NetWorker media information reporting command	mminfo(8)
mmlocate	reports the backup volume location information contained in the media database	mmlocate(8)
networker	a window-based display of the NetWorker server status	networker(8)
networker -x	starts NetWorker in troubleshooting mode, with hidden attributes displayed	networker(8)
nsr	guide to using NetWorker	nsr(8)
nsradmin	the NetWorker character-based program for system administration	nsradmin(8)
nsrck	the NetWorker check and repair program for the server indexes	nsrck(8)
nsr_crash	how to use NetWorker to recover from a disk crash	nsr_crash(8)



Command	Description	Manual Page
nsrd	the NetWorker server daemon	nsrd(8)
nsrexecd	the NetWorker client execution daemon	nsrexecd(8)
nsrim	the NetWorker index management program, usually invoked from the savegroup command	nsrim(8)
nsrindexasm	the module for saving and recovering NetWorker indexes	nsrindexasm(8)
nsrindexd	the NetWorker index daemon	nsrindexd(8)
nsrls	lists the statistics of NetWorker index files	nsrls(8)
nsrmmmd	the NetWorker media management daemon	nsrmmmd(8)
nsrmmdbasm	the module for saving and recovering NetWorker media databases	nsrmmdbasm(8)
nsrmmdbd	the NetWorker media index daemon	nsrmmdbd(8)
nsrwatch	displays the NetWorker server status from an ASCII terminal	nsrwatch(8)
nullasm	the NetWorker module for connecting save sets in the index	nullasm(8)
recover	the command to browse the on-line indexes and recover files from the NetWorker server	recover(8)
recoverindex	the command to recover the on-line indexes of a NetWorker server	saveindex(8)
save	the command to save files to the NetWorker server	save(8)

Command	Description	Manual Page
savefs	the command to save filesystems to the NetWorker server	savefs(8)
savegroup	the command to start saving the files of a group of NetWorker clients	savegroup(8)
saveindex	the command to save the on-line indexes of a NetWorker server	saveindex(8)
scanner	the command to read the contents of a volume to recover from NetWorker server crashes	scanner(8)
swapasm	the NetWorker module for saving and recovering swap files	swapasm(8)
tapeexercise	the command to exercise a tape drive in order to uncover problems	tapeexercise(8)
uasm	the NetWorker module for saving and recovering generic UNIX files	uasm(8)
xlateasm	the NetWorker module for encoding and decoding files	xlateasm(8)


Tip 

To print the NetWorker man pages, alphabetically by section, type:

```
pcat 'ls -1 /usr/catman/networker/man?/*' | col -bx |
lp -dpr
```

where **pr** is the name of an ASCII-capable printer. To print a single man page, such as **nsr(8)**, type:

```
pcat /usr/catman/networker/man8/nsr.8.Z | col -bx | lp
-dpr
```





C

Appendix D

Troubleshooting

This appendix contains troubleshooting information that addresses common questions concerning operating and configuring NetWorker.

Checking the NetWorker Daemons

If you have trouble starting NetWorker the daemons may not be running properly. To check the daemons, enter the following command:

```
# ps -ef | grep nsrd
```

If you discover that you need to start the NetWorker daemons enter the command below:

```
# /usr/sbin/init.d/rc.networker start
```

Using the networker -x Option

NetWorker can be invoked in an expert mode by starting it with the `-x` option. This displays hidden attributes that are used for advanced configuration and troubleshooting. To invoke NetWorker in expert mode, enter the following command at the system prompt:

```
mars% networker -x &
```

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 it will not be able to recover any files it backed up under its old name.

To change the name of a NetWorker client, you must first delete the old client name, then add the new client name, and rename the directory which contains the corresponding index.

Follow these steps:

1. Open the Clients window, and delete the old client name.
2. Create and configure the new client.

Use **mv** to rename the old client index directory to the new name.

1. Open a shell window.
2. Become *root* on the NetWorker server.
3. Change directories to */nsr/index*.
4. Use **mv** to rename the client's index. For example:

```
# mv old_client_name new_client_name
```

The following steps cause the database daemon **nsrmmdbd** to rename the client inside its database:

1. Make the old client name an alias of the new name.
2. Kill and restart the NetWorker daemons. To kill the daemons, use the **nsr_shutdown** command. To restart the daemons, use the **nsrd** command.

D

As soon as possible, complete a full backup of the renamed client's files.

Recover Access Issues

System administrators control client recover access by configuring the client. The Recover access list in the Clients window displays which machines can recover the client's files.

The following users have the ability to recover any files on any client:

root
operator
a member of the *operator* group

Other users can 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, will be owned by that user.

Previewing a Backup

Every time you add a new client to NetWorker, it is a good idea to check if NetWorker can 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.

For example:

```
# savegroup -p group_name
```

D

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, you will see an error message. Check for the following:

- ❑ Make sure **nsrexecd** 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 UNIX command **ps** on the client.
- ❑ If you do not use **nsrexecd** on a client, make sure the client can access the NetWorker software. Check that */usr/bin* is in the UNIX PATH variable for *root* on the server and on each client.

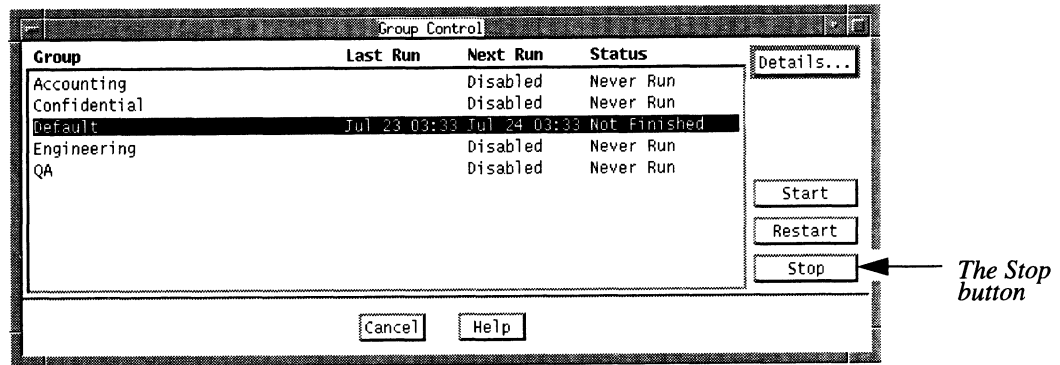
Be sure to edit the UNIX PATH for the appropriate shell for your system (for example, if you are using the *cshell*, edit the path of the *.cshrc* file; if you are using the *korn shell*, edit the path of the *.profile* file).

If you use a shell other than the *cshell*, you must enter the UNIX PATH in the Executable path field of the Clients window. You display this hidden attribute using the **networker -x** option.

Halting a Network Backup

To stop running a network-wide backup via the NetWorker X Windows interface, use the Stop button in the Group Control window.

D



The next network-wide backup will start as scheduled in the Start time field of the Groups window, or you may re-start the backup by pressing the Restart button in the Group Control window.


Backup Media Capacity

Occasionally you will find that NetWorker marks backup volumes as "full" when they are not really full. (The 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 may still be used for recoveries, but may not be used for any more backups.

If you see this "bad spot" behavior on many of the backup volumes, it may indicate the device needs cleaning or maintenance.

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



Savegroup Completion Messages

In the Notifications window, you configured NetWorker to mail the event notification about your savegroups. The Notifications window is pre-configured to mail the savegroup completion messages to *root*. Following are descriptions of error messages that may appear in the savegroup completion mail. Possible solutions are included.

Binding to Server Errors

NetWorker is designed to follow the client/server model. In a client/server model, servers provide services to the client through the Remote Procedure Call (RPC). These services live inside of long-lived UNIX processes, known as daemons.

For clients to find these services, the services must be registered with a registration service. When daemons start up they register themselves with the registration service. In UNIX the *portmapper* provides the registration service.

NetWorker servers provide a backup and recover service: they receive data from clients, store the data on backup media, and retrieve it on demand. If the NetWorker daemons are not running and a service is requested, for example, **save**, **recover**, **mminfo...**, then the following messages may appear in your savegroup completion mail:

"Server not available"

"RPC error, remote program is not registered (24015)"

These messages indicate the NetWorker daemons: *nsrd*, *nsrindexd*, *nsrmmmd*, *nsrmmdbd* are not running.

D

To restart the *nsr* daemons, enter **nsrd** at the system prompt:

```
mars# nsrd
```

Saving Remote Filesystems

You may receive the following error message in your savegroup completion mail when 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. Refer to the **nsrexecd** man page on how to re-configure **nsrexecd**.

When saving remote filesystems, files, or remote clients, NetWorker first tries to use **nsrexecd** and if that fails, invokes an **rsh** to that client as the **rcmd** user (the default **rcmd** user is *root*). If the client is supposed to run **nsrexecd**, re-start it. If not, make sure the NetWorker server can **rsh** as the **rcmd** user to the remote clients by adding an entry for *NetWorker_server rcmd_user* in the clients' */.rhosts* file(s). There are two areas where you may encounter problems with NetWorker and **rsh**:

1) 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** 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

D

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 will be 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 execute a command to read the *.<remote shell>rc* file.

Also, the remote shell has to allow an interactive session.

2) Rcmd user and *.rhosts*

The **rcmd user** is the user login name the NetWorker server will use 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 the following 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 will notify you that the file was changed during the backup in the Backup Status window and the savegroup completion mail. You can back up the file manually after it 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 the following steps:

1. Start NetWorker with the expert option by entering **networker -x** at the system prompt.
2. Open the Groups window and enter the name of the printer you are using to print the bootstrap in the Printer field.
3. Press Apply to save your changes.

Maximum Filename Length

The maximum size of a filename is the same as the UNIX *svid* limitation of 1024 characters.

D

Errors in Executing External ASMs

The following error message may appear during a backup:

```
client:/ save: external ASM '/xxxxxxasm' not found or not
executable
```

All external ASMs, (see **uasm** for a complete list of external ASMs) are separate programs, 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 program (typically **save** or **recover**). Rename or move the ASM following these conventions to correct this error.

Savegroup Completion Warning Messages

Occasionally the savegroup completion message includes one or more messages. These messages contain information that help the administrator understand why NetWorker performs certain tasks.

Below is one of the messages you might see:

```
quattro:/usr no cycles found in media db; doing full
save
```

In this example, the filesystem, */usr*, on the client *quattro* has no full saves listed in the media database. Therefore, despite the backup level pre-selected for that client's schedule, NetWorker will perform a full backup. This feature is important because it allows you to perform disaster recoveries for that client.

This message may also appear if the server and client's clocks are not synchronized. To avoid this, 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.

The following savegroup message may also appear:

```
NetWorker_server:index Saving server index because
server is not in an active group
```

If your server belongs to a group that is not enabled NetWorker will, to avoid a long recovery process, save the server's bootstrap information along with this group. As soon as possible you should enable the group your NetWorker server belongs to, or add the server to another group that is enabled.

Moving Indexes

Because the index databases are holey files, `cp` will create a file that consumes more disk space than the original file. To move indexes execute the following command in the `/nsr/index` directory:

```
mars# uasm -s -i <client index directory name> | (cd target_dir;
uasm -r)
```

Recovering Files from an Interrupted Backup

You will not be able to recover files from a backup terminated by killing the NetWorker daemons because the media index was not updated before the daemons died. Consequently, NetWorker will not know on which volume the requested file is located.

D

Determining the NetWorker Server

If you start NetWorker from a remotely mounted directory, you may receive the following message:

```
Using server server_name as server for client_name.
```

NetWorker looks for the system that is the fileserver of a remotely mounted directory and uses the NetWorker server assigned to that system as the backup server. To bypass this message, start NetWorker from a local filesystem.

Using nsrexecd

The **nsrexecd** daemon runs on NetWorker client machines. This daemon provides a secure and restrictive way for NetWorker to start automatic backups on clients. The **nsrexecd** daemon allows you to restrict access to a select set of NetWorker servers.

Security is increased by the use of a challenge/response scheme to ensure that only the NetWorker server is initiating connections, and not another program. Commands are restricted to those needed for saving data (the **save** and **savefs** commands are described in Appendix A, "Theory of Operations").

Should you need to reconfigure **nsrexecd** (for example, to allow a different NetWorker server to back up a client), edit the file */etc/networker.params* on the client. See **nsrexecd(8)** for a description of the command line configuration options. Put the arguments you want to pass to **nsrexecd** in the variable **nsrexecd_ARG**. After modifying **nsrexecd**, stop and restart the daemon with these commands:

```
# nsr_shutdown  
# /usr/sbin/init.d/rc.networker start
```

Glossary

This glossary contains terms and definitions found in this manual. Most of the terms are specific to the NetWorker products.

1-9	Intermediate backup levels. Each number represents a backup level. Lower levels back up more files
8mm	Choices in the Devices window; represents eight-millimeter magnetic tape.
archive	The process by which NetWorker backs up directories or files which have not been accessed for a period of time to an archive volume and then optionally deletes them to free up 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.
bootstrap	In NetWorker, the bootstrap consists of three files: the server's file index (<i>server.db</i>), the media index (<i>mmvolume</i>), and the server configuration information (<i>nsr.res</i>).

browse policy	The browse policy determines how long entries for your files will remain in the on-line file index.
client	A machine that accesses the NetWorker server to back up or recover files. Clients may be workstations, PCs, or file servers with gigabytes of data.
command line	The shell prompt, where you enter commands.
compressasm	The NetWorker directive used for compressing and decompressing files.
daemon	A long-lived program that implements a service. For example, <i>nsrd</i> 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 handle special files.
file index	A database of information maintained by NetWorker which tracks every file or filesystem backed up.
fileserver	A machine with disks that provides services to other machines on the network.
filesystem	1. A subtree of a UNIX file tree which is on a specific disk partition or other

	mount point. 2. The entire set of all UNIX files. 3. A method of storing files.
full (f)	A backup level in which all files are backed up, regardless of when they last changed.
group	A client or group of clients that starts backing up their 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 systems of different platforms that interact meaningfully across the network.
himt	A choice of media type in the Devices window; represents half-inch magnetic tape.
holey	A special ASM used to efficiently back up files that do not have all of their data blocks allocated.
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 machines from multiple vendors to communicate meaningfully.

job pack	A tray or carousel that holds several backup volumes.
level [1-9]	A backup level that backs up files that have changed since the last backup of any lower level.
machine	Any computer, including file or compute servers, diskfull workstations, or diskless workstations.
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.
middle (m)	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 to back up and recover filesystems.

NetWorker client	A machine that can access the backup and recover services from a NetWorker server.
NetWorker daemons	Daemons specific to the NetWorker environment.
NetWorker server	The machine on a network running the NetWorker software, containing the on-line indexes, and providing the backup and recover services to the clients on a network.
NetWorker resources	Components of NetWorker software configuration information, described by a list of attributes and values.
notice	A response to a NetWorker event.
nsrhost	The logical <i>hostname</i> of the machine that is the NetWorker server.
nsrwatch	A command that gives you a character-based display of NetWorker status information.
on-line indexes	The databases located on the server that contain all the information pertaining to the client backups and backup volumes.
operator	The person who monitors the server status, loads backup volumes into the server devices, and otherwise executes day-to-day tasks using NetWorker.
override	A backup level that takes place instead of the scheduled one.

pre-configured	Existing selections or configurations for different NetWorker features.
print	Send data to a printer.
qic	A choice in the Devices window; represents 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	The retention policy determines how long entries will be retained in the media index and thus be recoverable
recycle	Re-label a backup volume with the same name, thus removing all entries associated with the backup volume from the server's index.
resources	See <i>NetWorker resources</i>
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 machine on a network running the NetWorker software, containing the on-line index, and providing backup and

	recover services to the clients on a network.
shell prompt	The command line to which you enter UNIX commands.
skip (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.
super-user	A UNIX user with <i>root</i> privileges.
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 backup 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

- 1-9 backup levels 198
- 1-9 level backups 65
- 8mm tape cartridge
 - capacity 64
 - transfer rate 64
- 8mm tape drive 204
- A**
- Action field 121
- Active devices field 113
- Administration pull-down menu 24
- Administrator scrolling list 112
- administrators
 - adding 114
 - changing 115
 - deleting 115
- Append button 149
- appendable mode 51, 156
- Application Specific Modules 116
- Archive only choices 90
- automatic index management 100
- Autostart choices 44, 58
- B**
- backup
 - failed 49
 - files changed during 245
 - group 59
 - manual, network-wide 44
 - preview 240
 - scheduled, network-wide 43
 - setting up schedules 61
 - speeding 112
 - stopping a 240
- backup device speed 227
- backup group
 - changing time 60
 - configuring 56
 - creating 57
 - defined 56
 - deleting 61
 - details 48
 - pre-configured 56
 - restarting 48
 - starting, immediately 44, 47
 - stopping 48
- backup level
 - 1-9 198
 - defined 197
 - full 197
 - full versus incremental 199
 - incremental 198
 - setting up 198
- backup levels
 - 1-9 65
 - defined 66
 - full 65, 138
 - incremental 65, 138
 - override 63
 - skip 65
 - understanding 65
- backup media capacity 241
- backup schedule
 - convenience versus security 205
 - Default 207

- defined* 197
- Full Every Friday 208
- Full on 1st Friday of Month 209
- Full on 1st of Month 210
- Quarterly 211
- staggering 204
- time considerations 203
- typical 202
- backup schedules
 - changing 69
 - creating 68
 - deleting 67, 69
 - examples 69
 - pre-configured 63, 67
 - setting up 61
 - understanding 64
- Backup Status window 245
- backup volume
 - defined* 18
 - displaying information 50
 - finding the name 42
 - full 241
 - labeling 36, 39
 - location 155
 - managing 154
 - media format 176
 - mode 51, 152, 156
 - mounting 40
 - pre-mounting 41
 - removing 159
 - requested for recovery 41
 - unmounting 40
- bootstrap 163
- browsable status 153
- Browse policy 100
- browse policy 213
- button
 - Append 149
 - Details 48
 - Help 25
 - Instances 148
 - Overwrite 149
 - Reclaim space 53, 143
 - Remove oldest cycle 54, 142
 - resize 34
 - Restart 48
 - Separator 80
 - Start 47
 - Stop 48, 240
 - Update Server List 21

C

- Change Mode command 157
- Change Server command 21
- Change Server dialog box 21
- changing
 - backup group time 60
 - directives 119
 - mode 157
- choices
 - Archive only 90
 - Autostart 44, 58
 - Devices 91
 - Directive 127
 - Disabled 58
 - Enabled 58
 - Group 127
 - Media type 108
 - Month 63
 - Period 63
 - Pools 38
 - Schedule 127
 - Start now 45, 58

- Store file index entries 91
 - Week 63
- choosing server configurations 224
- client
 - CPU speed 230
 - data generation rate 229
 - performance 228
 - recover access 239
 - renaming 238
 - transfer rate 225
- client access, restricting 136
- client backups, optimizing 110
- Client retries sliding bar 59
- clients
 - access, restricting 136
 - adding 128
 - backups, optimizing 110
 - configuring 124
 - failed backup, re-trying 59
 - modifying existing 128
 - policies, selecting 130
 - removing 135
- Clients command 30
- Clients field 90
- Clients scrolling list 126, 142
- Clients window 31, 126, 130
- command
 - Change Mode 157
 - Change Server 21
 - Clients 30
 - Devices 107
 - Directives 118
 - Group Control 46
 - Groups 43
 - Indexes 52, 141
 - Label 37, 39
 - Label and Mount 37
 - Label Templates 77
 - mm** 233
 - mminfo** 231, 233
 - mmlocate** 233
 - Mount 41
 - mv** 238
 - networker** 33, 233
 - networker -x** 233, 237, 240
 - Normal 28
 - Notifications 120
 - nsr** 233
 - nsr_crash** 233
 - nsr_shutdown** 238
 - nsradmin** 231, 233
 - nsrck** 233
 - nsrck -F** 232
 - nsrd** 234, 238
 - nsrexecd** 136, 232, 240, 248
 - nsrim** 234
 - nsrindexd** 234
 - nsrls** 232, 234
 - nsrmm -d** 231
 - nsrmm -d -P** 231
 - nsrmm -l** 231
 - nsrmm -m** 231
 - nsrmm -m -R** 231
 - nsrmm -u** 231
 - nsrmmmd** 234
 - nsrmmdbd** 234, 238
 - nsrwatch** 232, 234
 - Policies 103
 - Pools 87
 - recover** 234
 - recoverindex** 234
 - Remove 159
 - rsh** 135
 - save** 234

- savefs** 235
 - savegroup** 232, 235
 - savegroup -p** 232, 239
 - saveindex** 235
 - scanner** 154, 160, 235
 - Schedules 62
 - Servers 110
 - Set Location 155
 - Tabular 27
 - tapeexercise** 235
 - Unmount 41
 - Volumes 51
 - commands (NetWorker command line)
 - nsrwatch** 173
 - recoverindex** 168, 169
 - saveindex** 164
 - scanner** 172
 - commands (UNIX)
 - df** 165
 - Complete Save Sets scrolling list 49
 - compressasm directive 68, 233
 - concurrent devices
 - Active devices 113
 - controlling backup traffic 114
 - defined* 112
 - distributing backups 114
 - Sessions per device 113
 - configuring
 - backup groups 56
 - clients 124
 - NetWorker servers 55
 - CPU speed 228
 - creating
 - backup groups 57, 59
 - backup schedules 68
 - clients 128
 - directives 117
 - index policies 100
 - label templates 77
 - volume pools 94
 - Current Month button 63
- D**
- daemons, *see commands*
 - NetWorker, starting* 237
 - Decade policy 223
 - deleting
 - backup groups 61
 - clients 135
 - devices 109
 - save sets 133
 - volume pools 87
 - dependent save sets 139
 - Details button 48
 - devices
 - adding 108
 - deleting 109
 - known to server 107
 - multiple 106
 - no-rewind 108
 - pathname* 108
 - supported by NetWorker 106
 - Devices choices 91
 - Devices command 107
 - Devices display 35
 - Devices scrolling list 107
 - Devices window 107
 - df** command 165
 - dialog box
 - Change Server 21
 - Label 39
 - Label and Mount 38

- Print Options 149
 - Save Options 149
 - Set Location 155
 - directive
 - compressasm 233
 - logasm 233
 - mailasm 233
 - nsrindexasm 234
 - nsrmmdbasm 234
 - nullasm 234
 - swapasm 235
 - uasm 235
 - xlateasm 235
 - Directive choices 127
 - Directive display 118
 - directives
 - changing 119
 - creating 118
 - defined* 116
 - pre-configured 116
 - using 116
 - Directives command 118
 - Directives scrolling list 118
 - Directives window 118
 - Disabled choice 58
 - disabling a network backup 58
 - display
 - Devices 35
 - Directive 118
 - Messages 35
 - Pending 35, 41
 - Server status 34
 - Sessions 35
 - displaying backup volume information 50
 - dkinfo** command 166
- E**
- electronic mail 54
 - Enabled choice 58
 - enabling a scheduled backup 58
 - event notification 119
 - exiting NetWorker 23
- F**
- failed backup 49
 - Failed Save Sets scrolling list 49
 - field
 - Action 121
 - Active devices 113
 - Clients 90
 - Fields 79
 - Next 80
 - Number of periods 105
 - Parallelism 111
 - Pools 89
 - Save sets 90
 - Server 127
 - Sessions per device 113
 - Start time 44, 58
 - Version 111
 - Volume Label 40
 - Volume name 39
 - Fields field 79
 - file
 - changed during backup 245
 - file index
 - defined* 100, 137
 - removing entries 100
 - filesystems
 - backing up specific 90
 - scheduling 131

finding a backup volume name 42
 full backup 65, 138
 full backup level 197
 full backup volumes 241
 full mode 157

G

getting help 20
 Group choices 127
 Group Control command 46
 Group Control window 45
 Group Status Details window 49
 Groups command 43
 Groups window 58

H

Help button 25
 Help pull-down menu 25
 hidden attributes 237

I

incremental backup level 65, 198
 index management
 automatic 100
 manual 141, 150
 index operations
 browsing 140
 inserting entries 140
 messages 143
 reclaiming space 140, 145
 reducing size 53

remove oldest cycle 144
 removing entries 140
 index size notification 122
 indexes
 file 100, 137
 managing 52, 137
 managing manually 140
 media 100, 138
 Indexes command 52, 141
 Indexes window 141
 Instances button 148
 Instances window 147

L

Label and Mount command 37
 Label and Mount dialog box 38
 Label command 37, 39
 Label dialog box 39
 label templates
 choosing 77
 creating 77
 defined 36, 75
 examples 80
 pre-configured 75
 Label Templates command 77
 Label Templates scrolling list 79
 Label Templates window 79
 labeling backup volume(s) 36, 39
 log default notification 122
 logasm directive 233

M

- mailasm directive 233
 - Main window 20, 34
 - managing backup volumes 154
 - managing indexes 52, 137
 - manual backup 44
 - manual index management 140, 150
 - media
 - format 176
 - index 100, 138
 - media request notification 123
 - Media type choices 108
 - Menu bar 34
 - Messages display 35
 - mm** command 233
 - mminfo** command 231, 233
 - mmlocate** command 233
 - mode
 - appendable 51, 152, 156
 - changing 157
 - defined* 51, 152
 - full 152, 157
 - recyclable 152, 157
 - Month choice 63
 - Month policy 220
 - Mount command 41
 - mounting backup volume(s) 40
 - multiple devices 106
 - multiple disks 229
- N**
- network transfer rate 226
 - NetWorker
 - browse policy 213
 - daily tasks 33
 - exiting 23
 - export mode 237
 - measuring performance 227
 - media format 176
 - notifications 119
 - operating 29
 - pre-configured policies 218
 - preparing for a disk crash 163
 - retention policy 216
 - securing environment 135
 - setting up environment 197
 - starting 19
 - supported devices 106
 - theory of operations 175
 - troubleshooting 237
 - networker** command 33, 233
 - NetWorker server
 - changing 21
 - choosing configurations 224
 - configuring 55
 - errors 242
 - parallelism 110
 - status 33
 - transfer rate 226
 - networker -x** command 233, 237, 240
 - Next field 80
 - Next Month button 63
 - Normal command 28
 - notification
 - about NetWorker events 119
 - backup media request 123
 - checking NetWorker notices 54
 - customizing 121
 - index size 122
 - log default 122

savegroup completion 123
 setting up 119
 Notifications command 120
 Notifications scrolling list 121
 Notifications window 121, 242
nsr command 233
nsr_crash command 233
nsr_shutdown command 238
nsradmin command 231, 233
nsrck command 233
nsrck -F command 232
nsrd command 234, 238
nsrexecd command 136, 232, 240,
 248
nsrim command 234
 nsrindexasm directive 234
nsrindexd command 234
nsrls command 232, 234
nsrmm -d command 231
nsrmm -d -P command 231
nsrmm -l command 231
nsrmm -m command 231
nsrmm -m -R command 231
nsrmm -u command 231
nsrmmmd command 234
 nsrmmdbasm directive 234
nsrmmdbd command 234, 238
nsrwatch command 173, 232, 234
 nullasm directive 234
 Number of periods field 105

O

open files, handling 50
 operating NetWorker 29
 Operation pull-down menu 23
 overriding a backup level 63

Overwrite button 149

P

Parallelism field 111
 parallelism, NetWorker server 110
 Pending display 35, 41
 Pending Save Sets scrolling list 49
 performance
 client 228
 measuring 227
 server 227
 Period choices 63
 policies
 Browse 100
 changing 105
 creating 104
 defined 100
 deleting 105
 naming 105
 pre-configured 104, 131
 recovering files 65
 Retention 100, 150
 selecting for clients 130
 Policies command 103
 Policies window 100
 policy
 Decade 223
 Month 220
 pre-configured 218
 Quarter 221
 Week 219
 Year 222
 Pools choices 38
 Pools command 87
 Pools field 89
 Pools window 88

- pre-configured
 - backup group 56
 - backup schedules 67
 - defined* 29
 - directives 116
 - label templates 75
 - pools 91
 - schedules 63
 - settings 29, 32
 - pre-mounting backup volumes 41
 - Previous Month button 63
 - Print Options dialog box 149
 - printing Instances window 149
 - pull-down menu
 - Administration 24
 - File 21
 - Help 25
 - Operation 23
 - View 27
 - Volume 154
- Q**
- Quarter policy 221
- R**
- Reclaim space button 53, 143
 - reclaiming disk space 53, 215
 - Recover access scrolling list 127, 129
 - recover** command 234
 - Recover window 138
 - recoverindex** command 168, 169, 234
 - recovering
 - backup volumes, requested 41
 - client access to 239
 - directories, complete 139
 - file policies 65
 - files 18
 - files, from index 216
 - recyclable mode 157
 - recyclable status 154
 - Remove command 159
 - Remove oldest cycle button 54, 142
 - removing a backup volume 159
 - removing a client 135
 - renaming a client 238
 - resize button 34
 - Restart button 48
 - restarting a backup group 48
 - Retention policy 100, 150, 254
 - retention policy 216
 - retrievable status 154
 - re-trying failed client backup 59
- S**
- save** command 234
 - Save Options dialog box 149
 - save session 113
 - save set
 - contents 52
 - defined* 18, 138
 - dependent 139
 - details 147
 - status 153
 - save set id 163
 - Save sets field 90
 - Save sets scrolling list 52, 127, 142, 153
 - savefs** command 235
 - savegroup
 - completion messages 242
 - error messages 243

- savegroup** command 232, 235
- savegroup completion notification 123
- savegroup -p** command 239
- saveindex** command 164, 235
- saving Instances window 149
- saving remote filesystems 243
- scanner** command 172, 235
- schedule
 - Default 207
 - Full Every Friday 208
 - Full on 1st Friday of Month 209
 - Full on 1st of Month 210
 - Quarterly 211
- Schedule choices 127
- scheduled backup 43
- Schedules command 62
- Schedules scrolling list 63
- Schedules window 63, 207
- scheduling filesystems 131
- scrolling list
 - Administrator 112
 - Clients 126, 142
 - Complete Save Sets 49
 - Devices 107
 - Directives 118
 - Failed Save Sets 49
 - Label Templates 79
 - Notifications 121
 - Pending Save Sets 49
 - Recover access 127, 129
 - Save sets 52, 127, 142
 - Schedules 63
 - Server 21
 - Volumes 51
- securing NetWorker environment 135
- Separator button 80
- Server field 127
- server performance 227
- Server scrolling list 21
- Server status display 34
- Server window 111, 113
- Servers command 110
- Sessions display 35
- Sessions per device field 113
- Set Location command 155
- Set Location dialog box 155
- setting up backup schedules 61
- setting up event notification 119
- setting up NetWorker environment 197
- skip backup level 65
- sort data 87
- speed
 - backup device 227
 - CPU 228
- speed key 23
- staggering backup levels 204
- Start button 47
- Start now choice 45, 58
- Start time field 44, 58
- starting
 - backups immediately 47
 - backups, immediately 44
 - NetWorker 19
- starting NetWorker daemons 237
- Stop button 48, 240
- stopping a backup 240
- stopping a backup group 48
- Store file index entries choices 91
- swapasm directive 235

T

Tabular command 27
 tabular view 26
 tape drives
 supported 225
 transfer rates 225
tapeexercise command 235
 theory of operations 175
 transfer rate
 clients 225
 network 226
 NetWorker server 226
 tape drives 225
 troubleshooting NetWorker 237
 typical backup schedule 202

U

uasm directive 235
 Unmount command 41
 Update Server List button 21

V

Version 111
 View pull-down menu 27
 Volume Management window 51, 151
 Volume name field 39, 40
 volume pools
 changing 97
 creating 94
 defined 85
 deleting 87, 89
 duplicates 96
 examples 98
 pre-configured 91
 specific clients 90

 specific filesystems 90
 Volume pull-down menu 154
 Volumes command 51
 Volumes scrolling list 51, 152

W

Week choice 63
 Week policy 219
 window 245
 Backup Status 245
 Clients 31, 126, 130
 Devices 107
 Directives 118
 Group Control 45
 Group Status Details 49
 Groups 58
 Indexes 141
 Instances 148
 Label Templates 79
 Main 20, 34
 Notifications 121, 242
 Policies 100
 Pools 88
 Recover 138
 Schedules 63, 207
 Server 111, 113
 Tabular 26
 Volume Management 51, 151

X

xlateasm directive 235

Y

Year policy 222

TIPS ORDERING PROCEDURES

TO ORDER

1. An order can be placed with the TIPS group in two ways:
 - a) **MAIL ORDER** - Use the order form on the opposite page and fill in all requested information. Be sure to include shipping charges and local sales tax. If applicable, write in your tax exempt number in the space provided on the order form.

Send your order form with payment to:

Data General Corporation
ATTN: Educational Services/TIPS G155
4400 Computer Drive
Westboro, MA 01581-9973

- b) **TELEPHONE** - Call TIPS at (508) 870-1600 for all orders that will be charged by credit card or paid for by purchase orders over \$50.00. Operators are available from 8:30 AM to 5:00 PM EST.

METHOD OF PAYMENT

2. As a customer, you have several payment options:
 - a) **Purchase Order** - Minimum of \$50. If ordering by mail, a hard copy of the purchase order must accompany order.
 - b) **Check or Money Order** - Make payable to Data General Corporation.
 - c) **Credit Card** - A minimum order of \$20 is required for Mastercard or Visa orders.

SHIPPING

3. To determine the charge for UPS shipping and handling, check the total quantity of units in your order and refer to the following chart:

Total Quantity	Shipping & Handling Charge
1-4 Units	\$5.00
5-10 Units	\$8.00
11-40 Units	\$10.00
41-200 Units	\$30.00
Over 200 Units	\$100.00

If overnight or second day shipment is desired, this information should be indicated on the order form. A separate charge will be determined at time of shipment and added to your bill.

VOLUME DISCOUNTS

4. The TIPS discount schedule is based upon the total value of the order.

Order Amount	Discount
\$1-\$149.99	0%
\$150-\$499.99	10%
Over \$500	20%

TERMS AND CONDITIONS

5. Read the TIPS terms and conditions on the reverse side of the order form carefully. These must be adhered to at all times.

DELIVERY

6. Allow at least two weeks for delivery.

RETURNS

7. Items ordered through the TIPS catalog may not be returned for credit.
8. Order discrepancies must be reported within 15 days of shipment date. Contact your TIPS Administrator at (508) 870-1600 to notify the TIPS department of any problems.

INTERNATIONAL ORDERS

9. Customers outside of the United States must obtain documentation from their local Data General Subsidiary or Representative. Any TIPS orders received by Data General U.S. Headquarters will be forwarded to the appropriate DG Subsidiary or Representative for processing.

DATA GENERAL CORPORATION TECHNICAL INFORMATION AND PUBLICATIONS SERVICE TERMS AND CONDITIONS

Data General Corporation ("DGC") provides its Technical Information and Publications Service (TIPS) solely in accordance with the following terms and conditions and more specifically to the Customer signing the Educational Services TIPS Order Form. These terms and conditions apply to all orders, telephone, telex, or mail. By accepting these products the Customer accepts and agrees to be bound by these terms and conditions.

1. CUSTOMER CERTIFICATION

Customer hereby certifies that it is the owner or lessee of the DGC equipment and/or licensee/sub-licensee of the software which is the subject matter of the publication(s) ordered hereunder.

2. TAXES

Customer shall be responsible for all taxes, including taxes paid or payable by DGC for products or services supplied under this Agreement, exclusive of taxes based on DGC's net income, unless Customer provides written proof of exemption.

3. DATA AND PROPRIETARY RIGHTS

Portions of the publications and materials supplied under this Agreement are proprietary and will be so marked. Customer shall abide by such markings. DGC retains for itself exclusively all proprietary rights (including manufacturing rights) in and to all designs, engineering details and other data pertaining to the products described in such publication. Licensed software materials are provided pursuant to the terms and conditions of the Program License Agreement (PLA) between the Customer and DGC and such PLA is made a part of and incorporated into this Agreement by reference. A copyright notice on any data by itself does not constitute or evidence a publication or public disclosure.

4. LIMITED MEDIA WARRANTY

DGC warrants the CLI Macros media, provided by DGC to the Customer under this Agreement, against physical defects for a period of ninety (90) days from the date of shipment by DGC. DGC will replace defective media at no charge to you, provided it is returned postage prepaid to DGC within the ninety (90) day warranty period. This shall be your exclusive remedy and DGC's sole obligation and liability for defective media. This limited media warranty does not apply if the media has been damaged by accident, abuse or misuse.

5. DISCLAIMER OF WARRANTY

EXCEPT FOR THE LIMITED MEDIA WARRANTY NOTED ABOVE, DGC MAKES NO WARRANTIES, EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, WARRANTIES OF MERCHANTABILITY AND FITNESS FOR PARTICULAR PURPOSE ON ANY OF THE PUBLICATIONS, CLI MACROS OR MATERIALS SUPPLIED HEREUNDER.

6. LIMITATION OF LIABILITY

A. CUSTOMER AGREES THAT DGC'S LIABILITY, IF ANY, FOR DAMAGES, INCLUDING BUT NOT LIMITED TO LIABILITY ARISING OUT OF CONTRACT, NEGLIGENCE, STRICT LIABILITY IN TORT OR WARRANTY SHALL NOT EXCEED THE CHARGES PAID BY CUSTOMER FOR THE PARTICULAR PUBLICATION OR CLI MACRO INVOLVED. THIS LIMITATION OF LIABILITY SHALL NOT APPLY TO CLAIMS FOR PERSONAL INJURY CAUSED SOLELY BY DGC'S NEGLIGENCE. OTHER THAN THE CHARGES REFERENCED HEREIN, IN NO EVENT SHALL DGC BE LIABLE FOR ANY INCIDENTAL, INDIRECT, SPECIAL OR CONSEQUENTIAL DAMAGES WHATSOEVER, INCLUDING BUT NOT LIMITED TO LOST PROFITS AND DAMAGES RESULTING FROM LOSS OF USE, OR LOST DATA, OR DELIVERY DELAYS, EVEN IF DGC HAS BEEN ADVISED, KNEW OR SHOULD HAVE KNOWN OF THE POSSIBILITY THEREOF; OR FOR ANY CLAIM BY ANY THIRD PARTY.

B. ANY ACTION AGAINST DGC MUST BE COMMENCED WITHIN ONE (1) YEAR AFTER THE CAUSE OF ACTION ACCRUES.

7. GENERAL

A valid contract binding upon DGC will come into being only at the time of DGC's acceptance of the referenced Educational Services Order Form. Such contract is governed by the laws of the Commonwealth of Massachusetts, excluding its conflict of law rules. Such contract is not assignable. These terms and conditions constitute the entire agreement between the parties with respect to the subject matter hereof and supersedes all prior oral or written communications, agreements and understandings. These terms and conditions shall prevail notwithstanding any different, conflicting or additional terms and conditions which may appear on any order submitted by Customer. DGC hereby rejects all such different, conflicting, or additional terms.

8. IMPORTANT NOTICE REGARDING AOS/VS INTERNALS SERIES (ORDER #1865 & #1875)

Customer understands that information and material presented in the AOS/VS Internals Series documents may be specific to a particular revision of the product. Consequently user programs or systems based on this information and material may be revision-locked and may not function properly with prior or future revisions of the product. Therefore, Data General makes no representations as to the utility of this information and material beyond the current revision level which is the subject of the manual. Any use thereof by you or your company is at your own risk. Data General disclaims any liability arising from any such use and I and my company (Customer) hold Data General completely harmless therefrom.

Legato NetWorker
Administrator's Guide

069-100495-01

Cut here and insert in binder spine pocket