Solstice Backup 4.2 Administration Guide

2550 Garcia Avenue
Mountain View, CA 94043
U.S.A.

Part No: 802-6104-10
Revision A, April 1996

SunSoft
A Sun Microsystems, Inc. Business

Contents

Preface

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

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

Note – In this document, the terms Solstice Backup, Backup, and NetWorker refer to the same product.

Audience

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

About this Guide

This guide assumes that Backup has already been installed on your server and clients. If the software is not installed, refer to the *Solstice Backup 4.2 Installation and Maintenance Guide* for installation instructions.

Caution – You must install the Backup software before you can proceed with the information in this guide.

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

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

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

Customer Support

SunSoft and SMCC support programs are designed to meet customers' complete system needs. The programs support a wide range of needs.

For detailed information about our services, support policies, and software subscriptions, contact your local customer support office for programs and their availability.

Electronic Support Access

In addition, Legato has a number of Technical Bulletins available electronically. To obtain a list of available Technical Bulletins, send e-mail to request@legato.com with a subject line of "send bulletins index."

You can download Sun's Technical Bulletins and binary patches from our World Wide Web site (http://www.sun.com/solstice/).

Technical Bulletins are also available from FaxWorker at (415) 812-6156.

What Typographic Changes Mean

The following table describes the typographic changes used in this book.

Table 0-1

Typeface or Symbol	Meaning	Example
`AaBbCc123`	The names of commands, files, and directories; on-screen computer output	Edit your `.login` file. Use `ls -a` to list all files. `machine_name% You have mail.`
`AaBbCc123`	What you type, contrasted with on-screen computer output	`machine_name%` **`su`** `Password:`
AaBbCc123	Command-line placeholder: replace with a real name or value	To delete a file, type `rm` *filename*.
AaBbCc123	Book titles, new words or terms, or words to be emphasized	Read Chapter 6 in the *Solstice Backup 4.2 User Guide*. These are called *class* options. You *must* be root to do this.

Shell Prompts in Command Examples

The following table shows the default system prompt and superuser prompt for the C shell, Bourne shell, and Korn shell.

Table 0-2

Shell	Prompt
C shell prompt	`machine_name%`
C shell superuser prompt	`machine_name#`
Bourne shell and Korn shell prompt	`$`
Bourne shell and Korn shell superuser prompt	`#`

Introduction 1

Computer systems should be regularly backed up to protect against the loss of valuable data. In a networked environment, where the number of files grows as systems are added, the need to protect data is crucial.

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

Overview

The following information introduces you to a few basic Backup software terms and concepts.

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

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

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

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

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

There are many other Backup terms and concepts you will become familiar with as you read this manual and use Backup.

Solstice Backup Product Family

In addition to Backup products for UNIX-based servers and TCP/IP networks, SunSoft offers other software products that provide additional network backup options.

- Backup for Windows NT™ provides support for Windows-based servers and TCP/IP networks and can back up a wide variety of UNIX machines, NetWare™ clients, and PC desktops.
- Backup for NetWare backs up and recovers data on NetWare networks and is fully interoperable with UNIX systems.
- Backup ClientPak™ software allows a wide variety of client platforms on heterogeneous networks to back up to and recover data from Backup servers.
- The Backup Jukebox Software Module allows you to fully realize the automatic backup capabilities of Backup.
- Backup database modules reliably back up the most widely used databases, including Oracle, Sybase, and Informix.
- The Backup Archive Application provides you with a sophisticated method for archiving valuable client data.
- The Backup Hierarchical Storage Management (HSM) Application provides a way to automatically move data between different media types and locations, thereby conserving network management resources.

- The Backup Simple Network Management Protocol (SNMP) Module allows communication of Backup event notifications to SNMP-compliant network management stations.

All Backup server products support the latest backup devices that provide the highest performance and most reliable solutions for your network backup and recovery requirements. Contact SunSoft or your local customer support contact for information about these products.

Using the Backup Programs

The Backup software package consists of three programs:

- *Backup Administrator* – used for managing and configuring your Backup environment
- *Backup Backup* – used for backing up data
- *Backup Recover* – used for recovering files

Users displayed on the Administrators list in the Server window have permission to use and maintain Administrator configurations. Users not included on the list can start and open the Backup Administrator window. However, any attempt to modify settings will result in the display of a message similar to the following.

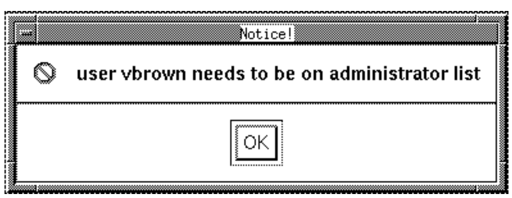

▼ Running the Backup Administrator Program

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

If the Admin icon already appears on your screen, simply double-click to open it. The icon includes a graphic and the name of the selected Backup server.

Backup Administrator icon

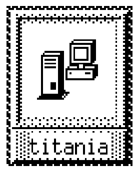

To easily identify unique servers in a multi-server network, Backup includes the current server name in the icon. To change the icon to display the Admin caption, use X11 resources to change the resource *iconServerName*, following these steps:

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

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

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

```
# nwadmin &
```

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

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

The Backup Administrator window appears, as shown.

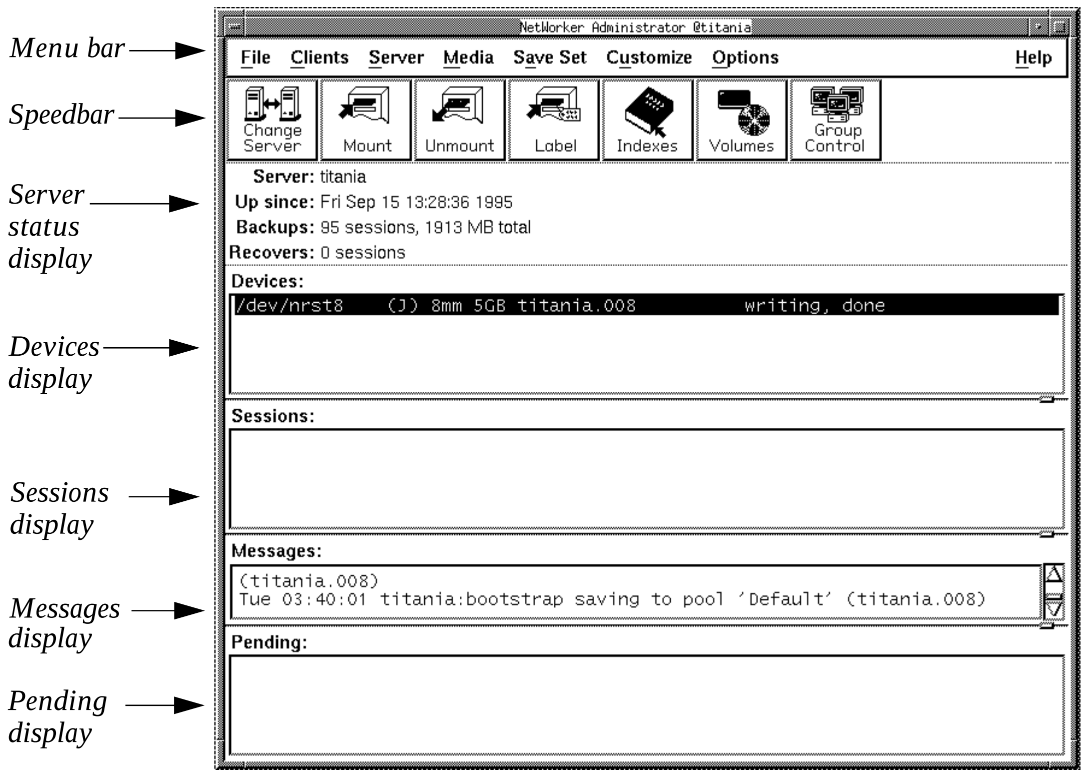

The Backup Administrator window has a menu bar, speedbar, and five displays.

- Menu bar – contains the File, Clients, Server, Media, Save Set, Customize, Options, and Help menus.

- Speedbar – provides convenient shortcut buttons for performing common Backup tasks. Throughout the manual, where appropriate, speedbar buttons are provided as an alternative to menu commands. Click a speedbar button to open its associated window.
  The following describes the speedbar buttons and their functions.

Change Server button – opens the Change Server dialog box; use this button instead of selecting the Change Server command from the Server menu.

Mount button – mounts a backup volume in a backup device; use this button instead of selecting the Mount command from the Media menu.

Unmount button – unmounts a backup volume in the backup device; use this button instead of selecting the Unmount command from the Media menu.

Label button – labels a volume in the backup device; use this button instead of selecting the Label command from the Media menu.

Indexes button – opens the Indexes window; use this button instead of selecting the Indexes command from the Clients pull-down menu.

Volumes button – opens the Volume Management window; use this button instead of selecting the Volumes command from the Media menu.

Group Control button – opens the Group Control window; use this button instead of selecting the Group Control command from the Server menu.

- Server status display – shows how long the current server has been online, the number of backups and recovers the server completed since it has been online, and the amount of data the backups contain.
- Devices display – shows the devices known to the server, backup volumes currently mounted, and the device status.
- Sessions display – provides current information about which clients are backing up, browsing, or recovering files.
- Messages display – provides a scrollable history of the server activity during the last 24 hours.

Note – The daily information shown in the Messages display is also added to the Backup `/nsr/logs/daemon.log` file. You can search the contents of this log file for messages no longer viewable in the Messages display.

- Pending display – displays messages requiring attention.
  When an event requiring operator intervention occurs, the program icon changes to alert you that Backup Administrator is in a “pending” mode.

Administrator “pending” icon

Running the Backup Program

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

If the Backup icon is already displayed on your screen, double-click it to open the window.

Backup icon

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

```
# nwbackup &
```

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

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

Refer to the *Solstice Backup 4.2 User Guide* for detailed information about the Backup program.

Running the Backup Recover Program

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

If the Recover icon is already displayed on your screen, double-click it to open the window.

Backup Recover icon

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

```
# nwrecover &
```

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

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

Refer to the *Solstice Backup 4.2 User Guide* for detailed information about the Backup Recover program.

▼ Using Backup Menus

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

There are two choices for navigating the Backup menus: the point and select method with a mouse or keyboard shortcuts.

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

1. **Hold down the [◆] key, if you are using a SPARC system, or the [**Alt**] key.**

2. **Press the key that corresponds to the underlined letter of the desired menu.**

3. **Release both keys, then press the key that corresponds to the underlined letter in the desired command name.**

File Menu

The File menu in the main Backup Administrator window contains the Exit command for exiting Backup.

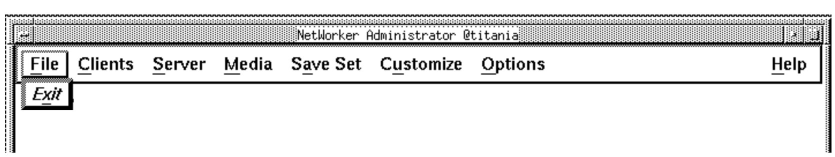

Other windows opened through Backup Administrator pull-down menus have different versions of the File menu containing the Save, Print, and Exit commands.

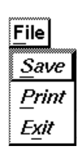

The Save and Print commands are only available when a Backup window is in the Tabular view. See “Displaying a Tabular View” on page 16 for more information.

▼ Printing Window Contents

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

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

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

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

1. **Select the Tabular command from the View menu. The content of the Backup window changes to a tabular view.**

2. **Select the Print command from the File menu. The Print Options dialog box appears.**

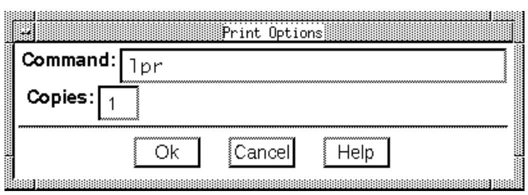

   The Command field displays the environment variable PRINTER, if set; otherwise the default, “lpr”, is displayed. You can modify the existing print command or enter a completely different one.

3. **Enter the desired UNIX print command in the Command field.**

4. **Enter the number of copies in the Copies field.**

5. **Click Ok to print.**

▼ Saving Window Contents

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

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

1. **Select the Tabular command from the View menu. The contents of the window changes to a tabular view.**

2. **Select the Save command from the File menu. The Save Options dialog box appears.**

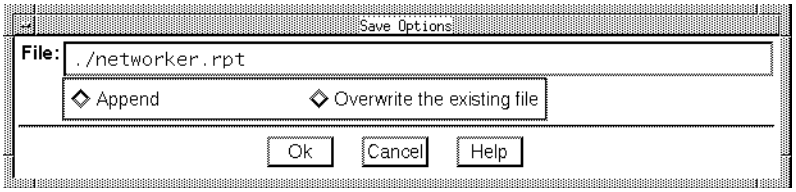

3. **Enter the name of the file in the File field.**
   The information is saved to a file specified by the environment variable REPORT, if set. Otherwise, the information is saved to *./Backup.rpt*. You can save the data to a different data file by entering its name into the File field.

4. **Specify an option for Backup to follow in case the file already exists. To update or append the existing file, click the Append button. To overwrite the existing file, click the Overwrite the existing file button.**

5. **Click Ok to save the contents of the window to the specified file.**

Exiting Backup

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

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

Clients Menu

The Clients menu contains commands for configuring your Backup clients and for monitoring and managing their indexes.

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

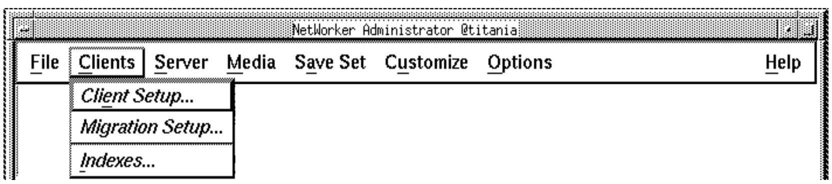

Server Menu

The Server menu contains commands for setting up, changing, and configuring the Backup server.

See Chapter 4, “Configuring and Monitoring the Server,” for a description of the commands available in the Server menu.

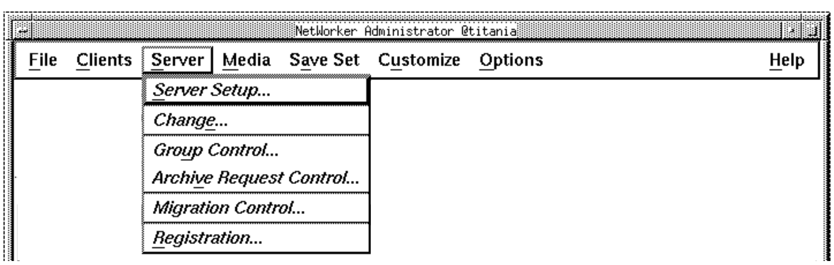

Media Menu

The Media menu contains the commands for using and managing the Backup server backup volumes. It also includes the commands for configuring and monitoring the devices used for backup and recovery, whether they are stand-alone backup devices or jukeboxes.

See Chapter 5, “Managing Media and Backup Devices,” for a description of the commands available in the Media menu.

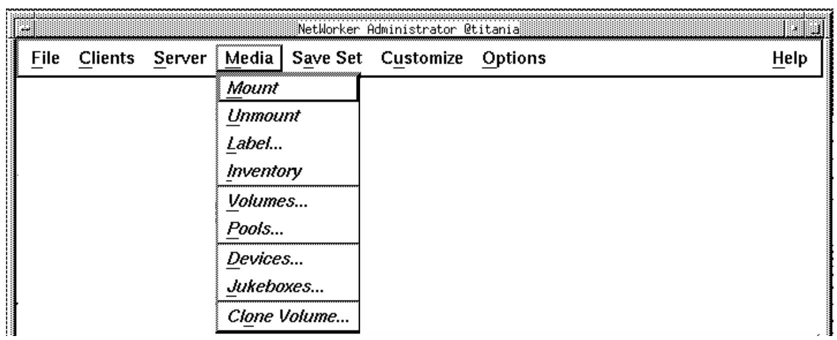

Save Set Menu

The Save Set menu contains commands for recovering entire save sets and for cloning, or making identical copies of save sets.

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

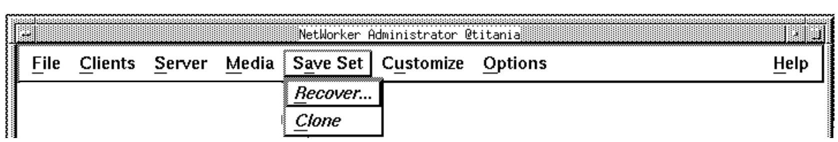

Customize Menu

The Customize menu contains selections for customizing and configuring your Backup product backups. The menu contains commands for configuring schedules, policies, directives, label templates, groups, and notifications.

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

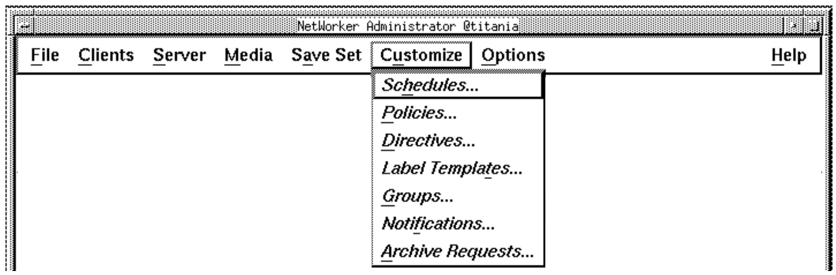

Options Menu

Use the Options menu to make selections for customizing the Backup interface. You can set the polling interval, choose whether or not to display the speedbar, and sound a tone when a pending message arrives. The Options menu includes the Polling Interval, Speedbar, and Tone commands.

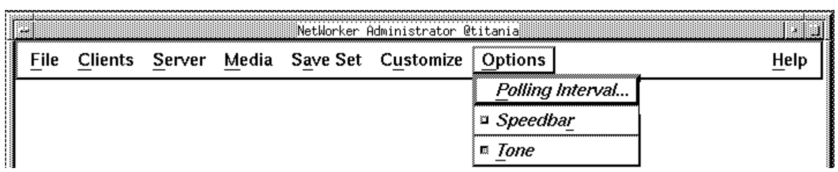

▼ Changing the Polling Interval

The Polling Interval command allows you to change how often, in seconds, the data in the Backup Administrator window is updated.

To change the polling interval, follow these steps:

1. Select Polling Interval from the Options menu. The Change Polling dialog box appears.

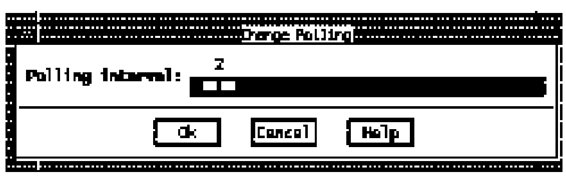

2. Use the mouse to slide the bar to the right or left to select the desired polling interval.

3. Click Ok to apply the change.

Displaying the Speedbar

Backup provides the convenience of a speedbar for the functions you perform most often. By default, the speedbar is enabled. Use the Speedbar toggle button in the Options menu to display or hide the speedbar.

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

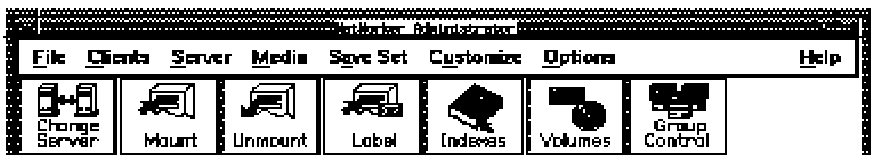

Setting the Tone Option

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

Help Menu

All Backup windows contain a Help menu that offers you a variety of choices for help on using Backup.

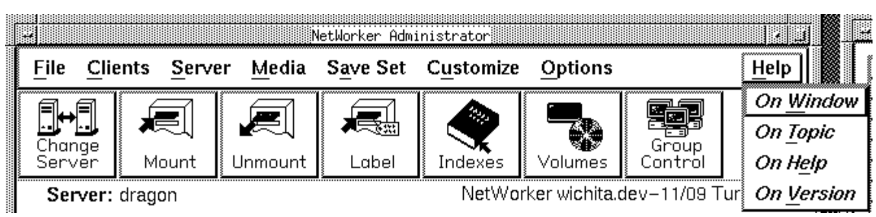

The Help menu contains four commands.

- On Window – displays a help screen with information about the currently-selected window.
- On Topic – displays a list of topics from which you select a topic.
- On Help – gives you information on how to use the Help facility.
- On Version – displays the version of Backup software running on the server.

View Menu

The View pull-down menus located in many of the Backup Administrator windows contain the Tabular, Details, and Search commands. Use these commands to view the contents of a window in tabular format or to see more details in a current window.

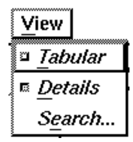

▼ Displaying a Tabular View

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

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

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

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

1. **Select the Client Setup command from the Clients menu. The Clients window appears, as shown.**

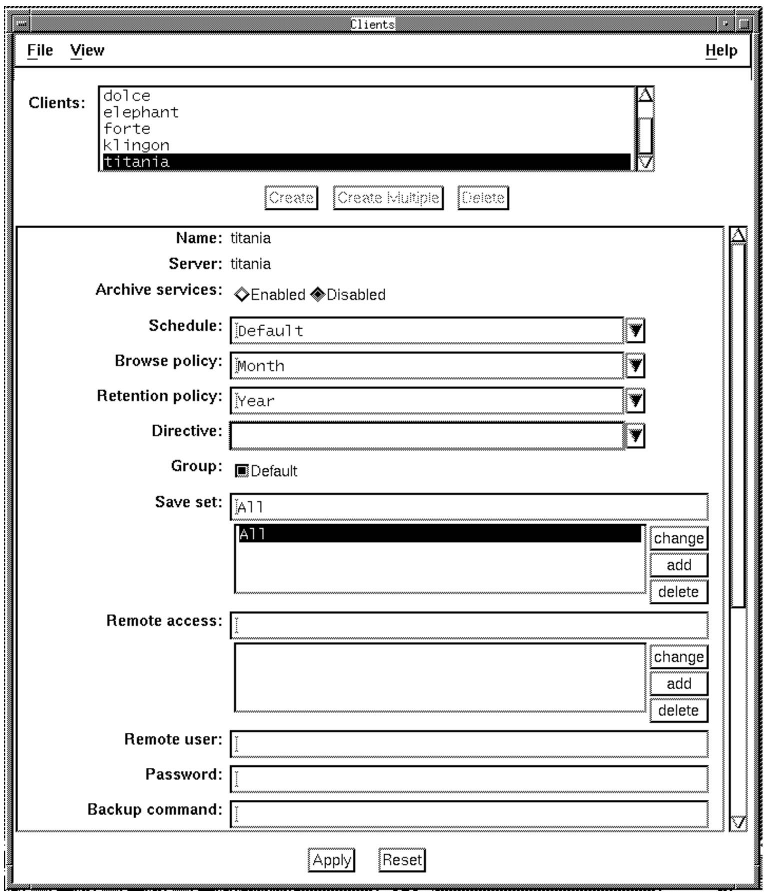

2. **Select the Tabular toggle button in the View menu. The Clients window changes to a tabular view, as shown.**

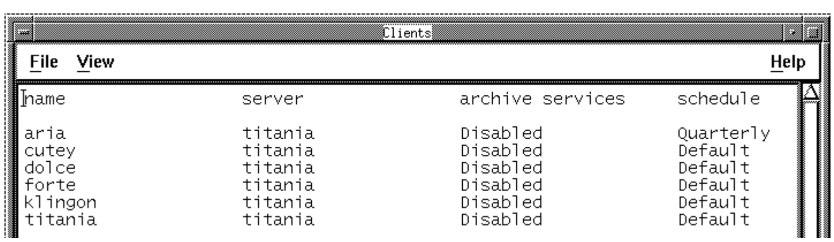

Use the scroll bars or resize the window to see all of the client configuration information.

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

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

▼ Displaying Details

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

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

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

1. Select the Groups command in the Customize menu. The Groups window appears, as shown.

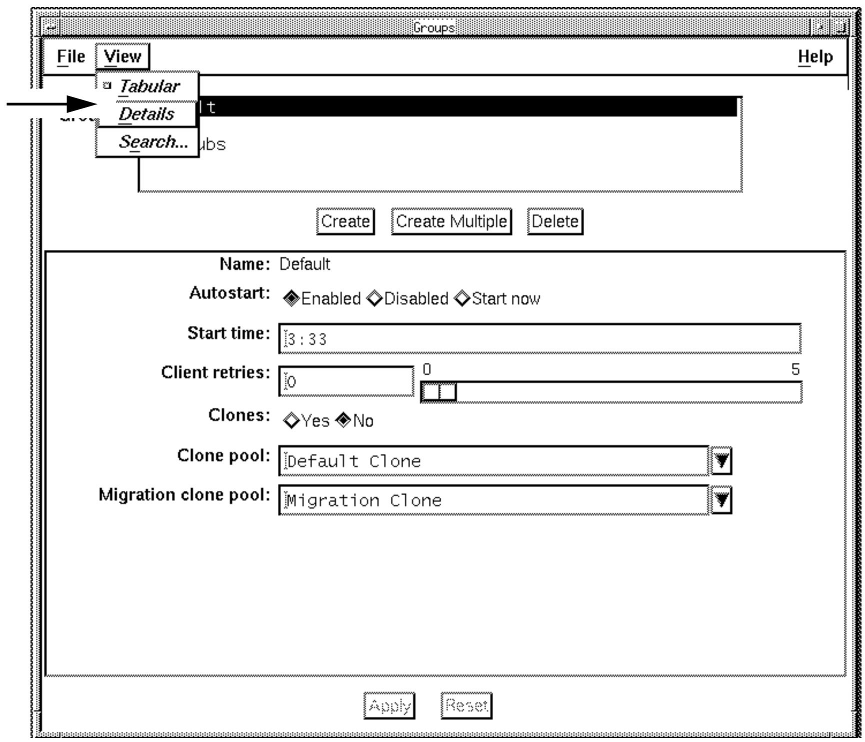

2. Select the Details toggle button from the View menu.

The Groups window changes to show you more detailed information, as shown.

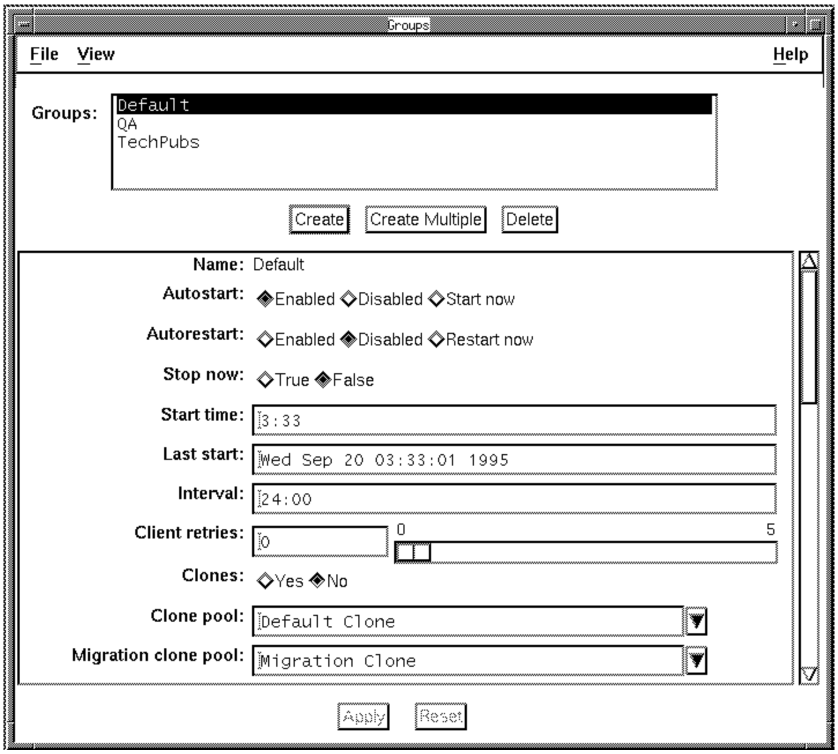

The Groups window displays additional fields with more selections and information. To view the entire contents of the window, use your mouse to drag the scroll bar to the bottom of the window.

▼ Searching for Resources

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

The following example shows how to search for a Backup client after opening the Clients window:

1. **Select Search from the View menu. The Search dialog box appears, as shown.**

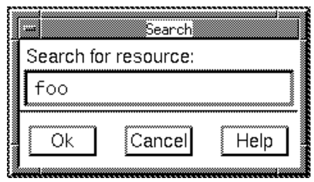

2. **Enter all or part of the name of the client you want to locate in the Search for resource field.**

Note – The search is not case-sensitive. You can enter the search string without having to recall exactly how a Backup resource is entered on your system.

---

3. **Click Ok to begin the search.**

If the resource is found, it is highlighted in the scrolling list. If the resource is not found, a dialog box similar to the following appears.

Generic Buttons

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

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

- Reset button – click to cancel changes and display the current settings.

▼ Create Multiple Button

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

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

1. **Open the Clients window by selecting Client Setup from the Clients menu.**
2. **Click the Create Multiple button.**

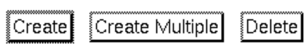

3. **Enter each client name in the Name field, one at a time. Click the add button after entering each name.**
4. **Select the common configurations that apply to all the new clients.**
5. **Click the Apply button to apply the configurations to all the new clients.**

▼ Print and Save Buttons

Several Backup windows contain Print and Save buttons for printing and saving the contents of the window. The following example demonstrates how to print and save the contents of the Instances dialog box.

1. **Open the Indexes window by clicking the Indexes speedbar button or by selecting Indexes from the Clients pull-down menu.**
2. **Select a save set and click the Instances button in the Indexes window to open the Instances dialog box.**

3. **Click the Print button in the Instances dialog box. The Print Options dialog box appears, as shown.**

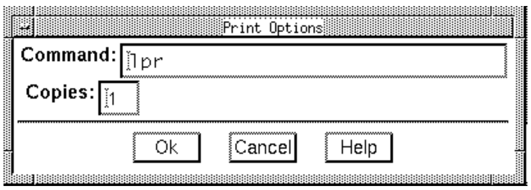

4. **Enter the print command (for example, lpr) in the Command field.**

5. **Enter the number of copies in the Copies field.**

6. **Click Ok to print.**

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

1. **Open the Indexes window by selecting Indexes from the Clients pull-down menu.**

2. **Click the Instances button in the Indexes window to open the Instances dialog box.**

3. **Click the Save button in the Instances dialog box. The Save Options dialog box appears, as shown.**

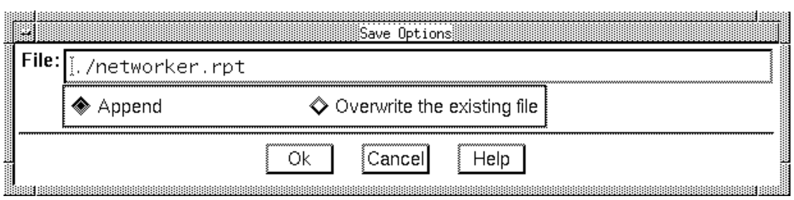

   The information is saved to a file specified by the environment variable REPORT, if set; otherwise, the information is saved to *./Backup.rpt*. You can save the data to a file of your choice by entering its name in the File field.

4. **Specify an option for Backup to follow in case the file already exists. To update or append the existing file, click the Append button. To overwrite the existing file, click the Overwrite the existing file button.**

5. **Click Ok to save the information from the Instances dialog box to the specified file.**

1

Summary

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

Chapter 2, “Getting Started,” provides an opportunity to test drive Backup. Detailed instructions lead you through the tasks required for a first-time backup of your files using Backup.

Getting Started 2

This chapter leads you through a test drive of the Backup software and introduces you to the benefits of the powerful features Backup provides.

This test drive takes you through the following tasks:

- starting Backup
- reviewing the preconfigured settings
- labeling and mounting a backup volume
- performing an unscheduled backup and recovery
- viewing the Schedules window for information on scheduling backups
- viewing the Clients window for information about your Backup clients
- enabling a backup in the Groups window
- performing a scheduled backup immediately using the Group Control window
- verifying the scheduled backup in the Group Control window

If you have an evaluation copy of the Backup software, it is valid for 30 days from the time you install it. This provides time to learn about Backup's features and to decide which Backup configuration best benefits your network environment.

Starting Backup

To start the Backup Administrator program, enter the `nwadmin` command, shown below, at the system prompt or select the Backup Administrator icon, if already displayed.

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

If you do not specify a server when you start the Backup Administrator program and you have more than one server on your network, Backup chooses the currently-selected server. For more information on the location of the server, see the `nsr` man page.

Note – To ensure that you do not have difficulty starting Backup, make sure your DISPLAY variable is set correctly, and the X Window System interface is running.

The Backup Administrator window appears, as shown.

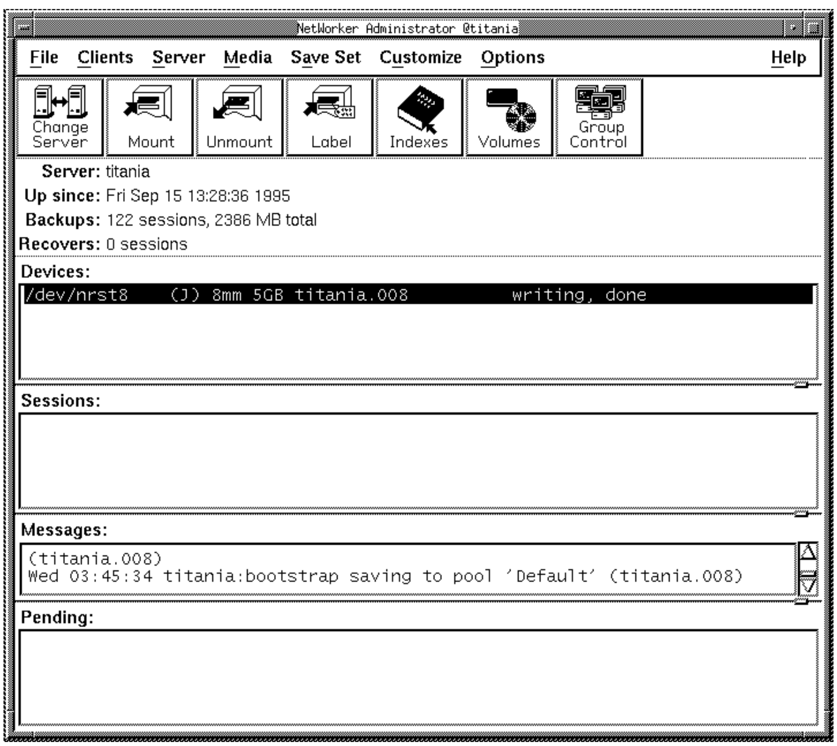

Use the Backup Administrator window to administer, monitor, and configure your backups. See Chapter 1, "Introduction" for a complete description of this window.

Using Preconfigured Settings

Backup is shipped with easy-to-use preconfigured settings. Preconfigured settings are the default choices Backup uses for backing up the clients on your network.

Backup provides preconfigured settings for the following:

- *pools* for sorting data to preselected backup volumes
- *schedules* for controlling when and how data is backed up
- *policies* for managing the file and media indexes
- *directives* for customizing the way Backup backs up data
- *label templates* for labeling backup volumes consistently and automatically
- *groups* for determining the groups to which clients belong
- *notifications* for informing you of Backup activity

You can use these preconfigured settings to back up your client data immediately after installing Backup. You have the option to change most of the preconfigured settings or create your own.

Once you enable the Default group in the Groups window, Backup backs up the server as a client of itself using the preconfigured settings.

▼ Labeling and Mounting the Backup Volumes

Backup volumes must be labeled before the Backup product can back up files to them. The Backup product uses these labels to track and identify the contents of the backup volumes. When you start a recover, Backup automatically determines the necessary volume(s) and requests them by the label name. This enables you to quickly locate backup volumes required for recovering a file.

You must mount the backup volumes after labeling them before you can proceed with the backup. Volumes are automatically mounted when a jukebox is used for backup.

See "Jukebox Labeling Dialog Box " in Chapter 8 for details on labeling jukebox volumes.

Follow these steps to label and mount backup volumes in a device other than a jukebox:

1. **Locate blank media that can be used for a backup.**
2. **Load the media into the Backup server backup device.**
3. **Highlight the device in the Devices display in the Backup Administrator window.**
4. **Click the Label speedbar button.**
   The Label dialog box appears, with a field for entering the name of the volume.

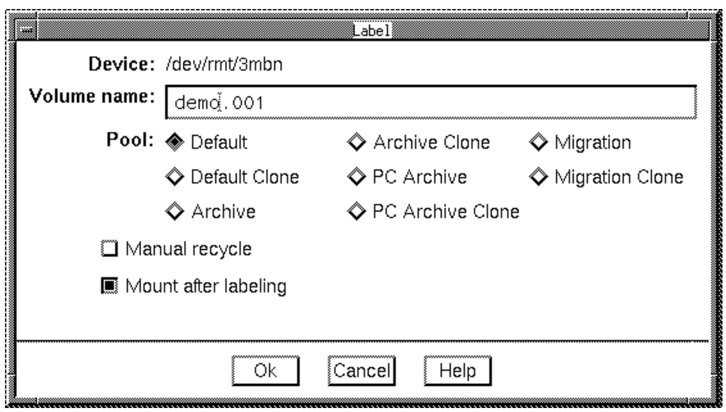

5. **Enter a name, such as "demo.001" into the Volume name field.**
   You do not need to select a pool because the Default pool is already selected.
6. **Click the Mount after labeling box.**
7. **Click the Ok button to label and mount the backup volume.**

A dialog box appears, asking you to confirm your request. Backup asks for confirmation because existing data on a volume is lost when you relabel a volume. Click OK to confirm the labeling request or Cancel to quit the request.

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

Now that you have completed labeling and mounting a backup volume, you are ready to back up files.

Note – See “Labeling and Mounting Backup Volumes ” in Chapter 5 of this manual for more information on labeling and mounting features provided by Backup.

Backup and Recover Features

Typically, the day-to-day operation of Backup occurs with scheduled, unattended backups. However, you can use the Backup product Backup program to initiate an immediate backup of selected files and the Backup Recover program to recover files.

Caution – Read the “Quick Tour” section in the *Solstice Backup 4.2 User Guide* now to become familiar with unscheduled backups and recovers.

After completing the “Quick Tour” in the *Solstice Backup 4.2 User Guide*, you will be familiar with Backup’s backup and recover features.

Using Backup’s online file and media indexes, users can browse their backed-up files and quickly identify files for recovery. Backup also automatically identifies which backup volumes contain the files needed for recovery. Adding a jukebox to your server enables you to automate your backups and recoveries.

Note – Backup offers many features to help you recover files in addition to those introduced in the “Quick Tour.” See the *Solstice Backup 4.2 User Guide* for complete details.

Viewing the Backup Schedule

Most backup activity occurs unattended with Backup scheduling features. The following section describes the Schedules window.

For this exercise, you will not perform any operations in this window; you will only view its contents.

By default, your backup server is considered a Backup product client. Backup backs up your server using the preconfigured Default schedule.

Open the Schedules window by selecting Schedules from the Customize pull-down menu. The Schedules window appears, as shown.

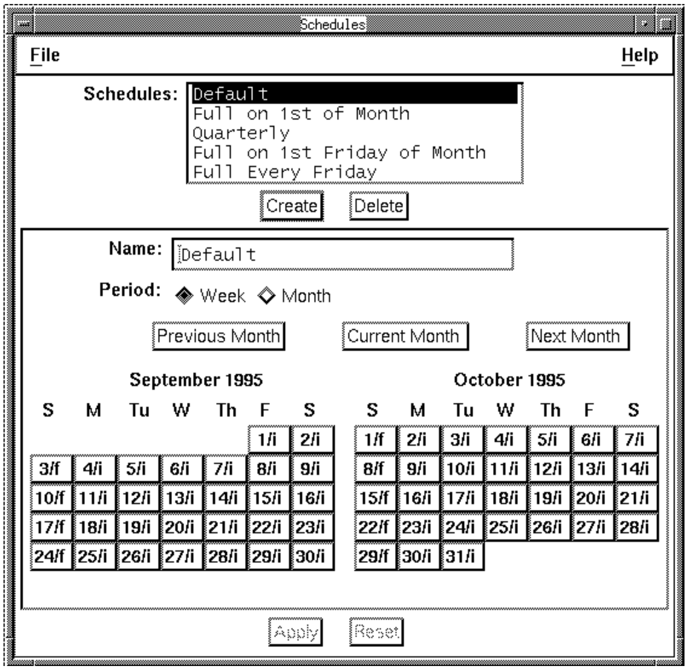

Notice each calendar day displays an "f" or an "i" next to the day. The "f" represents a full backup and the "i" represents an incremental backup.

The schedule named Default completes a full backup every Sunday, followed by incremental backups every other day of the week.

You can select one of the preconfigured schedules included with Backup from the Schedules window or create your own. We designed these schedules based on feedback obtained from Backup customers about the most popular ways to perform network-wide backups.

The Schedules window provides a calendar you can use to set up or create your own backup schedules. For more information about schedules, see "Setting up Backup Schedules " in Chapter 7 of this manual.

Caution – Regardless of the selections in the preconfigured schedules, Backup always performs a full backup the first time it backs up your files.

Viewing the Clients Window

To back up other client machines to your server, install Backup on the clients, then configure them using the Clients window. Backup backs up your clients using the selected configurations.

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

You do not need to make changes in this window for now. You may want to make changes and set up different backup schedules after you create more groups and clients.

After you enable the Default group, the server is backed up as a client of itself, using the preconfigured settings in the Clients window.

Note – The Clients window is discussed in detail in Chapter 3, "Configuring and Monitoring Clients."

The Clients window with preselected settings is shown below.

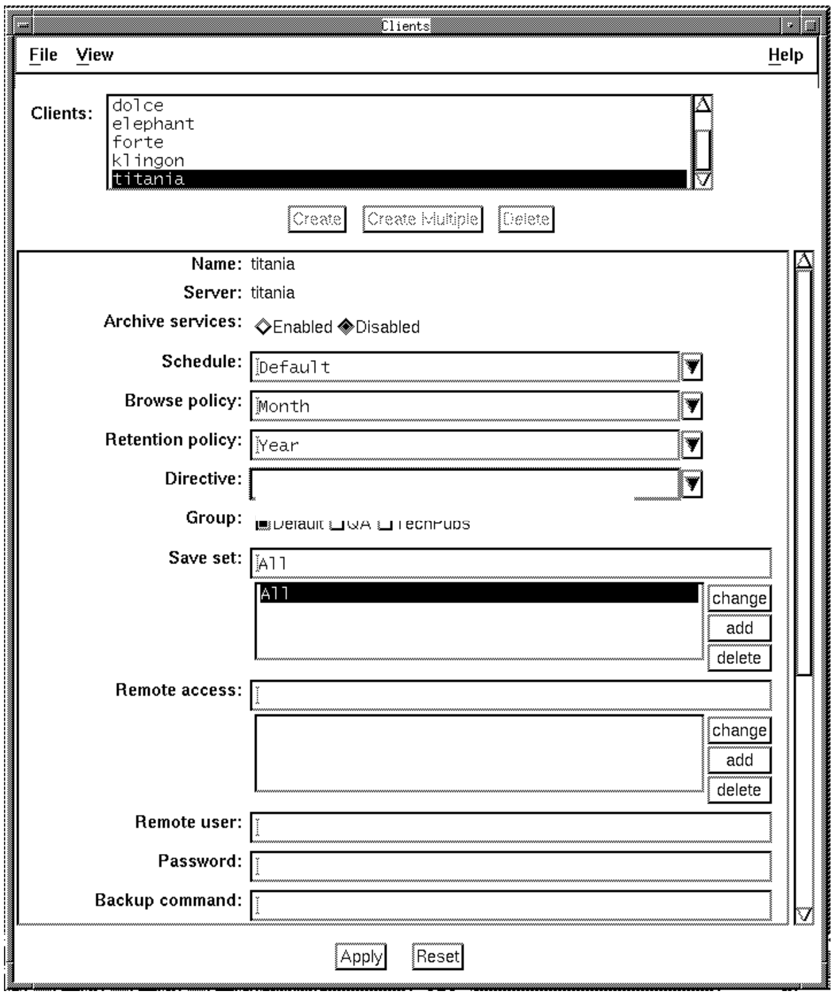

Enabling the Scheduled Backup

Now that you have viewed the Schedules and Clients windows, open the Groups window. In the Groups window you need to enable the Default group, to which your backup server belongs.

When you enable the Default group, the Backup product backs up all the clients in that group according to their individual selections. Backup backs up your server based on the preconfigured settings in the Default group and the Default backup schedule.

Open the Groups window by selecting Groups from the Customize pull-down menu. The Groups window appears, as shown.

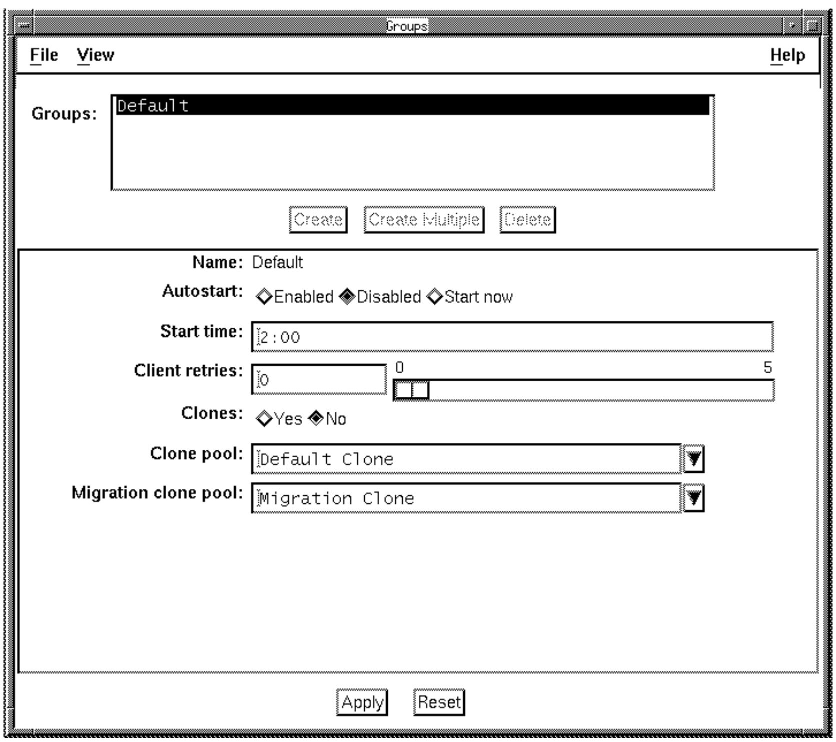

Notice that the Start time is 3:33 a.m. and the preselected choice for the Autostart choices is Disabled. If you want to change the time Backup backs up a group, simply enter a new start time.

To enable the Default group, follow these steps.

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

Your Backup server is automatically included in the group Default. By enabling the Default group, your server and any other clients you add to it in the future start backing up at 3:33 a.m. every night.

Note – The Groups window is described in detail in “Configuring Backup Groups ” in Chapter 7 of this manual.

▼ Starting a Backup Now

You viewed several different windows showing how to back up your Backup server with the preconfigured selections. Now, start a backup process to see how the Backup software backs up a client.

There are two ways to begin a backup immediately. You can click the Start now choice for Autostart options in the Groups window and click Apply before exiting the window, or you can click the Start button in the Group Control window.

In this example, you will start a backup in the Group Control window where you can also stop, restart, and monitor the progress of a backup.

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

1. **Click the Group Control speedbar button or select the Group Control command from the Server menu.**

The Group Control window appears, as shown.

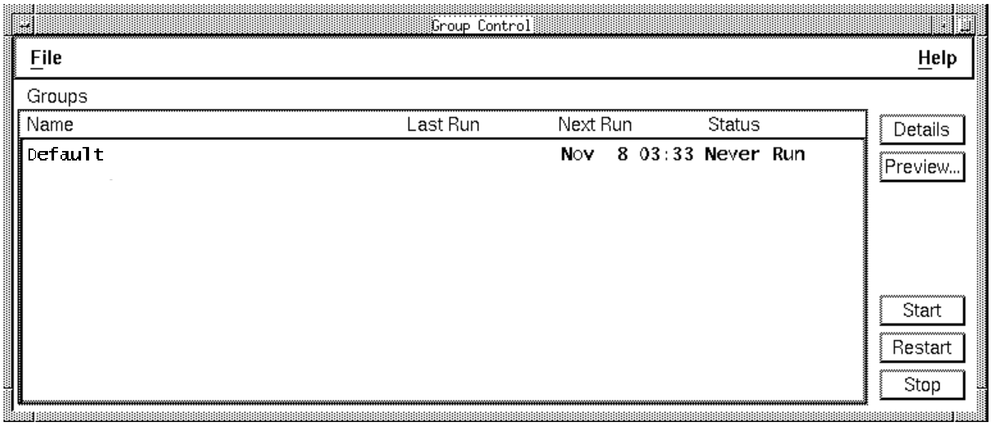

2. **Select the Default group in the Groups scrolling list.**

3. **Click the Start button.**

A confirmation dialog box appears, as shown. Click OK to start the backup or Cancel, if you change your mind.

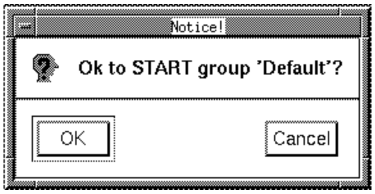

Backup begins a backup of the clients in the Default group and displays "running" in the Status column of the Group Control window. Normally the Default group would use the Default schedule to back up the data. However, the first time Backup performs a backup for a client, it performs a full backup for all of its data, regardless of the schedule assigned.

By starting a backup now, you override the predefined backup time for the Default group, which is 3:33 a.m. However, immediately starting a backup does not eliminate the regularly scheduled backup for the Default group. If the Default group is enabled, it will back up again at 3:33 a.m. as scheduled.

Verifying the Backup

After selecting the Start button, keep the Group Control window open to verify and monitor the progress of your backup. After Backup has successfully backed up your files, the Finished message appears in the Status field.

If you want more detailed information about the files being backed up, click the Details button. The Group Control Details dialog box appears, as shown.

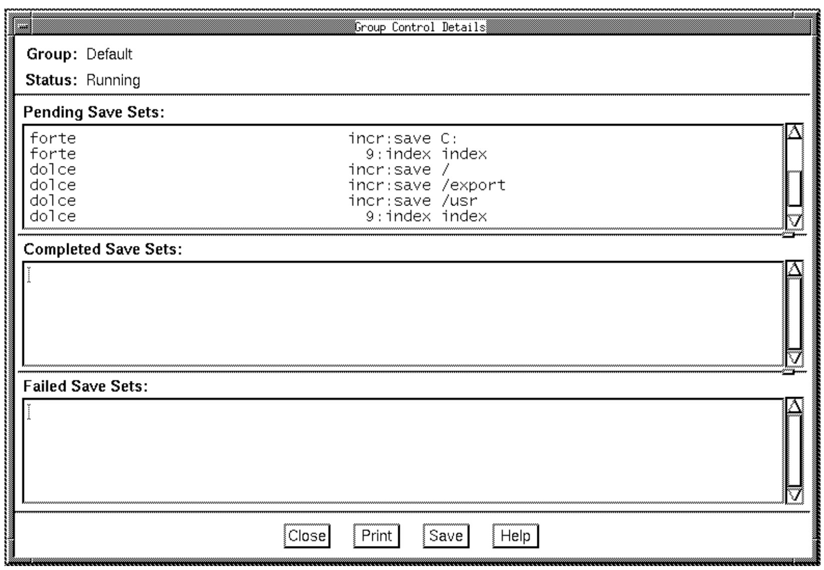

Note – For more information about the Group Control window and Group Control Details dialog box, see Chapter 4, “Configuring and Monitoring the Server.”

The Next Step

Listed below are suggestions for further configuration exercises.

- Add clients to the Default group and enable the group for an overnight backup. In the morning, check the results of the scheduled backup in your mail.

- Create your own backup schedules and apply them to the appropriate clients. Create several groups containing clients grouped logically by department. For example, create an Accounting group that backs up all of the machines in accounting at 12:00 a.m. and a Marketing group that backs up all of the marketing machines at 2:00 a.m.
- Run several recovery tests. After completing one or more backups, rename a file in your system. Use the Backup Recover window to recover the original version of the file.

Summary

You have completed a quick test drive of Solstice Backup and have seen that Backup is a powerful tool for protecting files across an entire network. Backup recovers your files in a quick and efficient manner.

You have also seen that Backup is easy to administer and configure from any machine on your network using the X Window System GUI.

Configuring and Monitoring Clients 3

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

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

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

Caution – You must have administrator privileges to add a client to the Backup product backups or change settings in the Clients window.

Configuring Backup Clients

Backup provides preconfigured settings for backing up clients. These settings also apply to the server, which automatically becomes a client of itself. You may use these preconfigured settings or create your own and apply them to clients in the Clients window.

Navigating the Clients Window

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

Open the Clients window by selecting Client Setup from the Clients pull-down menu.

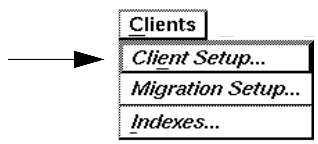

The Clients window appears, with the preconfigured settings as shown.

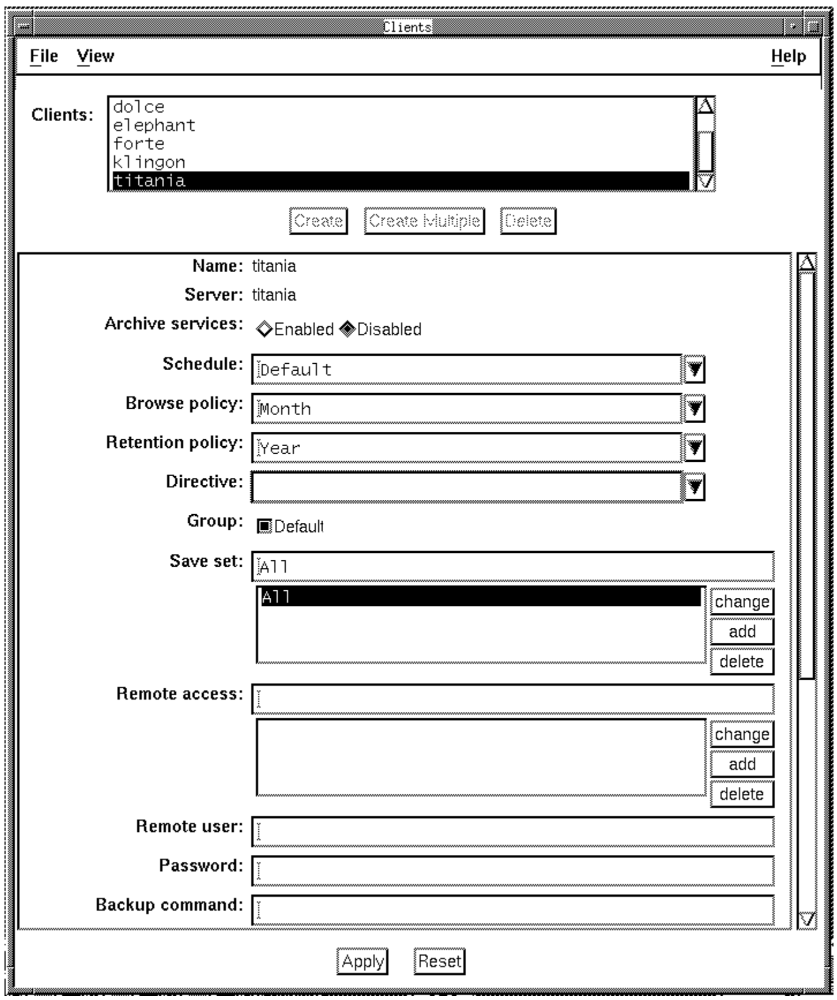

Caution – The server becomes a client of itself when you install Backup. Backup automatically adds the server hostname to the Clients scrolling list and provides the preconfigured settings shown above. You must enable the Autostart feature in the Groups window for Backup to back up the server automatically.

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

Note – For more information on the preconfigured settings available, refer to Appendix A in the *Installation and Maintenance Guide*.

- Clients scrolling list – displays an alphabetical list of *hostnames* for Backup clients of the selected server. Since the Backup server automatically becomes a client of itself, its *hostname* appears in the Clients list when you open the Clients window for that server.
- Name field – displays the name of the selected client. To add a client: click the Create button, enter the name of the client in the Name field, select the appropriate settings, and then click Apply.
- Server field – displays the name of the Backup server currently in use.
- Archive services choices – allows enabling or disabling of archive services for clients. This option is available only with the purchase of the optional Solstice Backup Archive Application. Refer to Chapter 9, "Archive Application," for more information.
  All Backup backup clients are automatically enabled for archive services when the Archive Application is installed.

- Schedule field – displays the assigned schedule for the selected client. Click the arrow button to access the Schedule scrolling list. The list displays schedules included with Backup and any new schedules you create in the Schedules window.
  One of the preconfigured Schedule settings included with Backup is Default, which performs a full backup every Sunday and an incremental backup every other day of the week.

- Browse policy field – displays the chosen browse policy for the selected client. Click the arrow button to access the Browse policy scrolling list containing the policies included with Backup and any new policies you create in the Policies window.
  The default Browse policy setting is Month. Entries for the client files backed up remain in the file index for one month. Backed-up files can be browsed and recovered in the Backup Recover window for one month.

- Retention policy field – displays the retention policy for the selected client. The retention policy determines how long the save sets on a volume will be recoverable before the volume is available for recycling. Click the arrow

button to access the Retention policy scrolling list containing the policies included with Backup and any new policies you create in the Policies window.
The default Retention policy setting is Year. The names of the backup volumes are retained for one year in the media index. Files can be recovered from the backup volumes for up to one year, unless the volumes are overwritten.

Note – Backup generates an error message if you attempt to set a Browse policy for the client that exceeds the Retention policy. For example, a Browse policy setting of one month with a Retention policy setting of one week would generate an error message. A volume is usually recycled once its retention has expired; therefore, the online file index entries are no longer valid.

- Directive field – displays the directives recognized by the Backup server. The Directive field is blank, by default. Click the arrow button to access the Directive scrolling list containing the directives included with Backup and any new directives you create in the Directives window.
- Group choices – displays all backup groups known to the current Backup server. To enable scheduled backups for the selected group, you must highlight the group name in the Groups window, then click the Enabled button for the Autostart feature.
  Backup provides a preconfigured group named Default, which starts a backup daily at 3:33 a.m. Additional backup group settings are displayed in the Clients window once they have been set up within the Groups window. You can assign a client to as many groups as you want.
- Save set field – displays individual client save sets (typically comprising one or more client filesystems) for the selected Backup client. The Save set scrolling list displays all of the filesystems entered for a Backup client. To backup up data from a single client filesystem, enter the client filesystem name.
  The Save set setting included with Backup is All, which backs up all of the local filesystems for the selected client.

- Remote access field – displays the remote access permission assigned to the selected client. A blank field indicates that only the client itself can recover its backed-up files. If you want specific users or other clients to have access to data for another client, enter their *user@hostname* or *netgroup* name in this field.

  Use the scrollbar to access the lower portion of the Clients window, as shown.

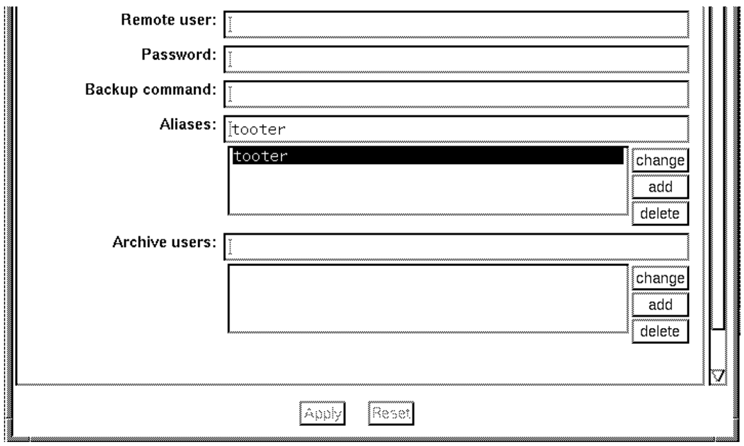

- Remote user field – use this field to restrict client permissions by entering the client user name. Backup uses the *user name* in this field to execute `save` and `savefs` on a client during a scheduled backup. When this field is blank, the *user name*, by default, is *root*.
- Password field – use this field only if the client is a PC running NetWare, to include the PC client in a scheduled backup. Enter the password used by the Backup server to log into the client system in the Password field and press [Return].

Note – If you enter a password in this field, you must also enter a user name in the Remote user field.

- Backup command field – enter the name of a program file you create that includes commands to execute before and after Backup backs up client data. Create the file in the directory where the `save` command is installed. See "Using the Pre- and Post- Processing Backup Command" in this chapter for details.

- Aliases field – use this field to enter an unlimited number of aliases for a selected Backup client. Client names may be in simple local name format (for example, *myclient*) or fully qualified DNS format (for example, *myclient.acme.com*). List all aliases for the client to help eliminate failed saves due to unrecognized or mismatched client names on the network.
- Archive users field – you can only access this field if the optional Archive Application is enabled on your Backup server. Refer to Chapter 9, "Archive Application," for more information.

Using the Pre- and Post- Processing Backup Command

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

The backup command must follow these conventions:

- the program name must begin with the prefix `save` or `nsr` and cannot exceed 64 characters
- the program file must reside in the same directory as the `save` command, typically `/usr/bin/nsr`
- the `save` command must be included in the program to ensure data is properly backed up
- all commands within the program file must successfully execute, otherwise Backup cannot complete the remaining instructions

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

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

Enter the program name in the Backup command field in the Clients window. When Backup begins a backup process, it checks for an entry in the Backup command field for the client. If a command is present, Backup runs that command instead of the `save` command normally used to back up client data.

Since Backup does not run the standard `save` command when a backup command is indicated, you must include the `save` command in your program. If you do not include the `save` command in the program, Backup will not back up the data.

Example – Backup Command

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

```
#!/bin/sh
case $0 in
/* )PATH=/bin:/usr/bin:`/bin/dirname $0`
c=`/bin/basename $0`
;;
* )PATH=/bin:/usr/bin:/usr/sbin
c=$0
;;
esac
export PATH
# print message at start of backup
echo “$c”: backup started at ‘date‘
# perform backup on client
save “$@” > /tmp/saveout$$ 2>&1
# print message at completion of backup
echo “$c”: backup complete at ‘date‘
# print save output last for savegrp
cat /tmp/saveout$$
rm -f /tmp/saveout$$
exit 0
```

Caution – You should immediately try backing up a client once you have created your program. This ensures that the backup command created does not prevent the client from backing up successfully during an unattended backup session. Click the Start button in the Group Control window to start an immediate backup.

Configuring Clients

Before a client can back up to a Backup server, the client software must be installed. The software can be installed locally on the client disk or NFS-mounted over the network. After the software is successfully installed on the client, add the client to your Backup software backup processes in the Clients window.

Note – Refer to the *Solstice Backup 4.2 Installation and Maintenance Guide* for instructions on installing the client software.

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

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

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

▼ Adding a New Client

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

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

7. **Select a set of directives for the client.**

8. **Add client aliases, if any, in the Aliases field.**

9. **Click Apply.**

Note – For details on how to add groups of clients who share common settings, see the instructions in “Create Multiple Button” on page 22 in Chapter 1 of this guide.

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

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

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

Follow these steps to back up a specific filesystem:

1. **Select All in the Save set scrolling list.**

2. **Click the delete button.**

3. **Enter the *pathname* into the Save set field for the filesystem that you want backed up. For example:**

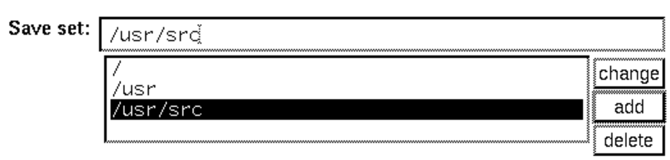

4. **Click the add button to add the client filesystem.**
   The *pathname* appears in the Save set scrolling list.

Note – This feature allows you to back up scheduled filesystems at different times. For example, if a client has a large amount of data, you may want to schedule its filesystems separately for backups. For more information, see “Scheduling Large Client Filesystems” on page 49.

By default, Backup clients can only browse or recover their own files. If your organization is concerned about security, we recommend you do not change the remote access for any of the clients.

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

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

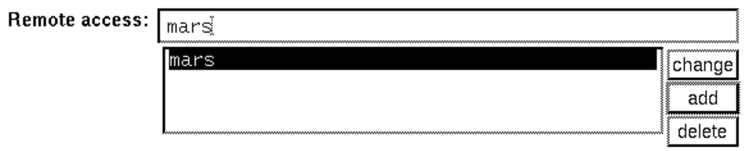

2. **Click the add button next to the Remote access scrolling list.**
   The *hostnames* or *netgroup* names added to the Remote access field appear in the Remote access scrolling list for the client.

3. **Decide whether to use the Remote user and Password fields. See** **"Navigating the Clients Window" on page 40 in this chapter for details.**

4. **Enter all aliases for the client in the Aliases field.**

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

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

▼ Scheduling Large Client Filesystems

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

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

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

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

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

1. Find the names and sizes of the locally-mounted client filesystems with the `df` command.

```
% df
```

Filesystem	kbytes8.0	used	avail	capacity	Mounted on
/dev/sd0a	15087	9898	3680	73%	/
/dev/sd0g	361474	313232	12094	96%	/usr
/dev/rs1c	559860	362532	141342	72%	/usr/src
/dev/sd1c	818627	695502	41262	94%	/export
/dev/sd0d	28181	10713	14649	42%	/var
/dev/rs0c	559860	435366	68508	86%	/home/atlas

The largest filesystem in this example is `/export`, with 818627 KB of space.

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

```
818627 kbytes / 300 kbytes/second = 2728.76 seconds

2728.76 seconds / 60 seconds/minute = 45.48 minutes
```

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

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

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

Use the Clients window to configure the client:

1. **Open the Clients window.**

2. **Select the client from the Clients scrolling list whose filesystems you plan to divide into different backup schedules. If the client has not already been added to the scrolling list, click Create to add the new client.**

3. **Configure the client by choosing a backup group, a directive, a policy, and the Monday Full schedule.**

4. **Delete All from the Save set list and add each filesystem you want to back up with the Monday Full schedule, one at a time.**
   Apply the Monday Full schedule to the following filesystems:

/dev/sd0a	15087	9898	3680	73%	/
/dev/sd0g	361474	313232	12094	96%	/usr
/dev/rs1c	559860	362532	141342	72%	/usr/src

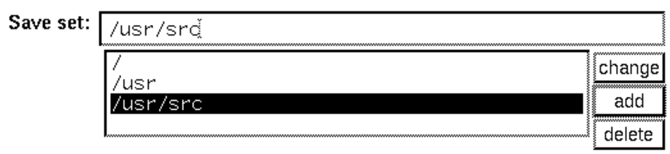

5. **Click the Apply button.**

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

1. **Click the Create button to add the duplicate client.**

2. **Configure the client by choosing a backup group, directive, policy, and the Wednesday Full schedule.**

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

```
/dev/sd1c   818627   695502   41262   94%  /export

/dev/sd0d    28181    10713   14649   42%  /var

/dev/rs0c   559860   435366   68508   86%  /home/atlas
```

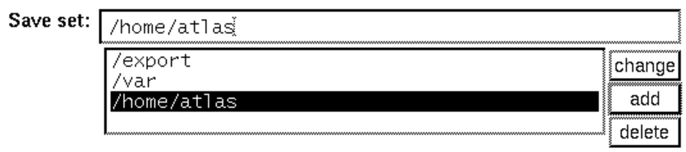

4. Click the Apply button.

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

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

▼ Removing a Client

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

1. Select the name of the client in the Clients scrolling list.

2. Click the Delete button. Backup asks for confirmation to delete the client.

Deleting a client means it cannot back up or recover its files from the Backup server. The backup history for the client still remains in the file and media indexes until you specifically remove the entries. See “Reclaiming Index Space” in Chapter 5 for instructions on how to remove the index entries for a Backup client.

Securing the Backup Environment

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

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

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

Restricting Client Access

Backup's preconfigured settings allow clients to browse and recover only their own files. To give other clients recover access to files on another client, the administrator must explicitly add the access to the Remote access list in the Clients window. See "Configuring Clients" on page 47 for more information.

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

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

The `nsrexecd` program runs on Backup client machines. This program provides a secure and restrictive way for Backup to start automatic backups on clients. The `nsrexecd` program allows you to restrict access to a select set of Backup servers. When you run the `pkgadd SUNWsbuc` command on a client, `nsrexecd` is started, and statements are added to the appropriate boot files to restart `nsrexecd` each time the client reboots.

Migration Setup

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

For complete instructions on the use of the Backup HSM Application, see Chapter 10, "Hierarchical Storage Management Application."

Manually Managing the Online Indexes

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

Note – See "Creating a New Policy" in Chapter 7 for more information on automatic index management.

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

This section provides you with the following concepts about the online indexes and their management:

- description of online indexes
- description of save sets
- description of the four online index actions

This section also includes instructions on manually managing the indexes and backup volumes:

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

Online Indexes

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

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

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

Save Sets

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

To conserve both index and backup volume space, Backup provides a variety of backup levels. For example, a *full* backup contains all files in a given filesystem, while an *incremental* backup contains only the files that have changed since the previous backup. Since directories often contain files that do not change along with files that do change, incremental backups conserve both index and volume space.

Note – For more information on backup levels, see the section “Understanding Backup Schedules” in Chapter 7 of this manual.

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

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

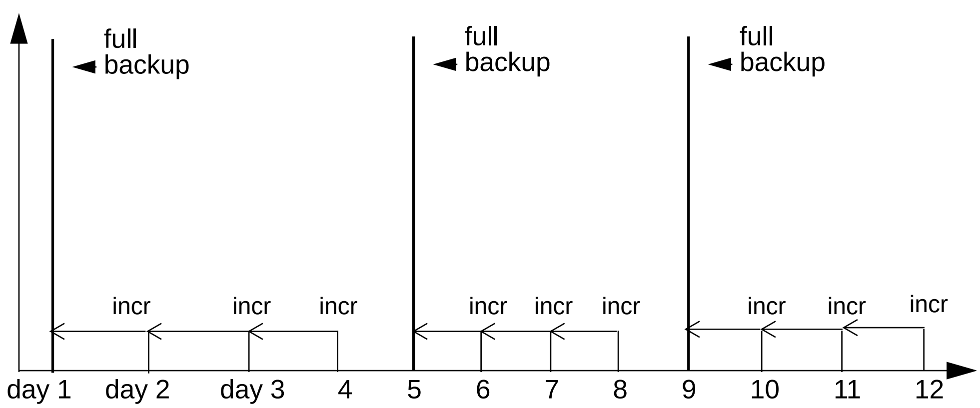

Caution – Incremental backups depend upon the previous full backup. Backup does not remove a save set until *all* of its dependent save sets have been removed.

Index Actions

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

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

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

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

Note – Remember to reclaim space after removing cycles.

▼ Generating File Index Reports

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

See the `nsrinfo` and `mminfo` man pages for more details on command line options used to customize the file index report.

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

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

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

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

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

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

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

```
scanning client `cutey' for savetime 811756903 (Fri Sep 22
01:01:43 1995) from the backup namespace

UNIX ASDF file `/var/adm/wtmp', fid = 5383.2066, file
size=105408, off=6105488, len=105672, app=backup(1)

1 file found
```

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

Navigating the Indexes Window

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

To open the Indexes window, click the Indexes speedbar button or select Indexes from the Clients pull-down menu.

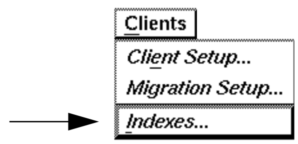

The Indexes window appears, as shown.

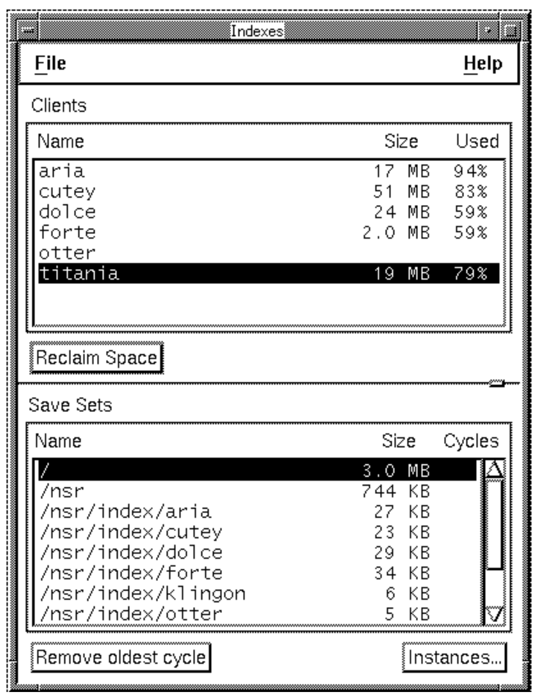

The Indexes window contains the following items.

- Clients scrolling list – contains the client names and the following information about each client file index:
  - Name – the Backup client name
  - Size – the allocated size of the client file index. The allocated disk space automatically grows as the index size increases.

- Used – the percentage of the index file in use. If 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.
- If the percentage is less than 100%, you can reduce the size of the index by first selecting the client and then clicking the Reclaim space button. This removes any holes in the file index created by removing index entries. Backup automatically removes index entries based on the browse and retention policies you selected for managing the online indexes.

- Reclaim space button – click to decrease the size of the client file index. Backup compresses the space in the index left by removed entries, thus freeing up disk space. A confirmation box appears when a reclaim space operation begins. Click OK to proceed or Cancel to discontinue.
- The Save Sets scrolling list – displays the save sets contained in the file index for the selected client. The save sets are grouped according to their unique save set names. For example, all backups of `/usr` in a client file index are grouped under the save set name `/usr`.
  - Name – shows client save set names
  - Size – displays an estimate of the amount of index space used by the save set group in the client file index
  - Cycles – displays the number of cycles for the save set group contained in the client file index. One cycle starts with a full backup and ends with the next full backup. A cycle includes the incremental and level 1-9 backups, if any exist. See Chapter 7, "Customizing Backups," for a description of Backup's backup levels.
- Remove oldest cycle button – used to remove the oldest cycle for the save set, if the index requires further reduction. Select the save set, then click the Remove oldest cycle button. After removing a save set, click Reclaim space to remove the holes left by removing the oldest cycle. A confirmation box appears once the removal of an oldest cycle has begun.
- Instances button – used to display entries contained in the file index for the selected save set. Instances are displayed in order, starting with the oldest entry. The Instances dialog box displays the save set ID, number of files, size, date backed up, and backup level for each occurrence of the selected save set. Click the Print button to print a copy of the information displayed. Click the Save button to save a copy of the current save set information to a separate disk file.

▼ Removing the Oldest Cycle

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

1. **Open the Indexes window by selecting Indexes from the Clients pull-down menu.**
2. **Select the client from the Clients scrolling list.**
3. **Select the save set group from the Save Sets scrolling list.**
4. **Click the Remove oldest cycle button to remove the oldest cycle for the save set group. A confirmation box appears, prompting you to confirm that you want to remove the oldest cycle of the selected save set.**
5. **Click Ok in the confirmation box to proceed or click Cancel to discontinue the operation.**

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

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

The following diagram illustrates what happens when you remove the oldest cycle from the index.

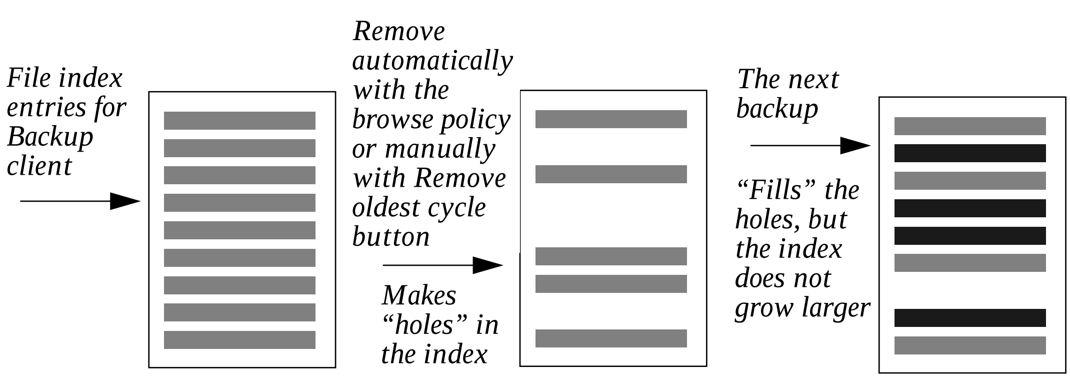

Note that removing index entries does not decrease the size of the file index, so it still takes up as much space as it did before. To decrease the amount of space the index uses, use the Reclaim space button.

Note – The last full cycle may not be removed using the Remove oldest cycle button because it might be needed for recovery.

▼ Reclaiming Index Space

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

1. **Open the Indexes window by selecting Indexes from the Clients pull-down menu.**
2. **Click the Reclaim Space button. You do not need to select individual clients or save set groups from the scrolling lists.**
   The Reclaim Space dialog box appears, as shown.

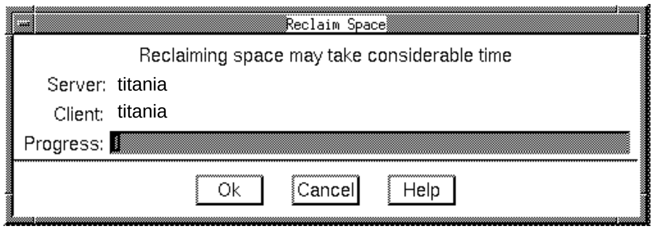

3. **Click the Ok button, to start the reclaim space operation. The Progress bar moves to show the progress of the index operation. When an index is busy, the Ok button is disabled.**
   The Reclaim Space dialog box changes, as shown.

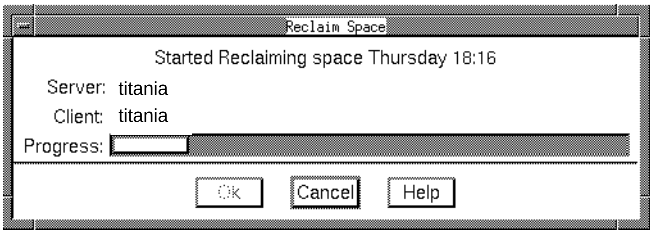

4. **Click the Cancel button to dismiss the dialog box. The reclaim space operation continues in the background. You can check the progress of the operation at any time by clicking the Reclaim space button to redisplay the confirmation box.**

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

The following diagram illustrates what happens when you use the Reclaim Space button.

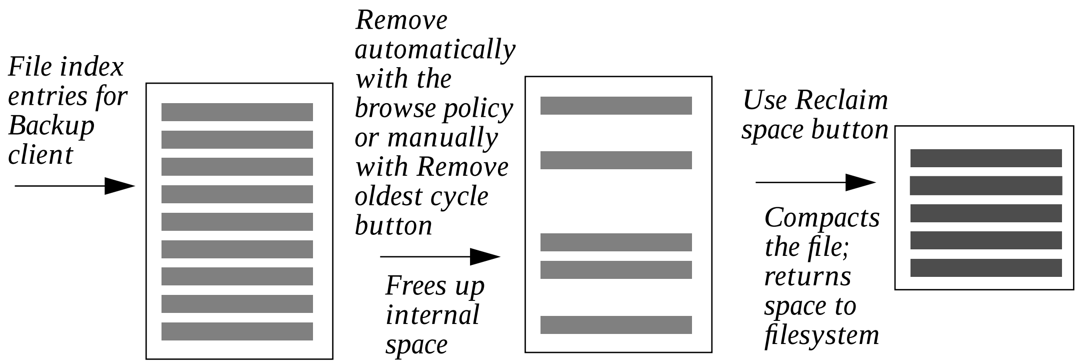

Caution – You may not simultaneously reclaim space and remove an oldest cycle for one client. The index is busy during both operations.

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

▼ Viewing Save Set Details

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

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

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

1. Select a client from the Clients scrolling list in the Indexes window.

2. Select a client save set from the Save Sets scrolling list.

3. Click the Instances button.

The Instances dialog box appears, as shown.

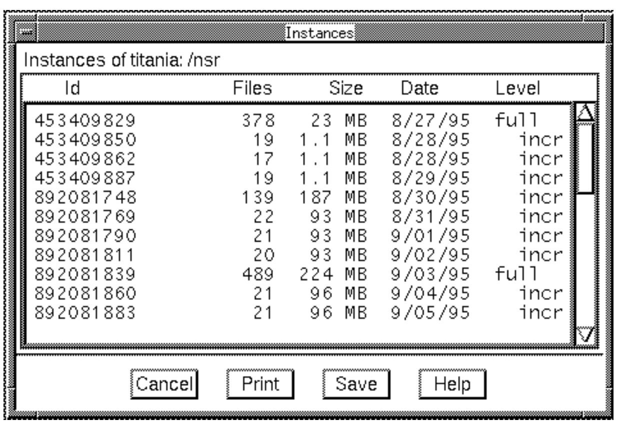

The title bar for the Instances dialog box displays the name of the currently selected save set. The Instances scrolling list contains the following information:

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

Configuring and Monitoring the Server 4

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

This chapter describes how to use the commands in the Server pull-down menu for configuring and controlling some of these Backup server operations.

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

- change server parallelism
- use devices concurrently
- add or change a Backup administrator
- change to another Backup server
- monitor backups in the Group Control window
- immediately start, stop, and restart a backup
- register Backup

Configuring the Backup Server

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

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

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

Navigating the Server Window

Use the Server window to accomplish the following tasks:

- change server parallelism
- use devices concurrently
- add or change Backup administrators
- provide important company and product information

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

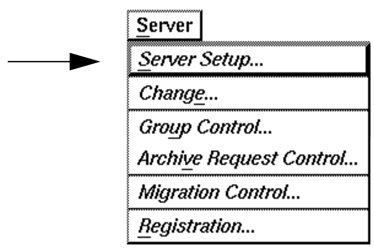

The Server window appears, as shown.

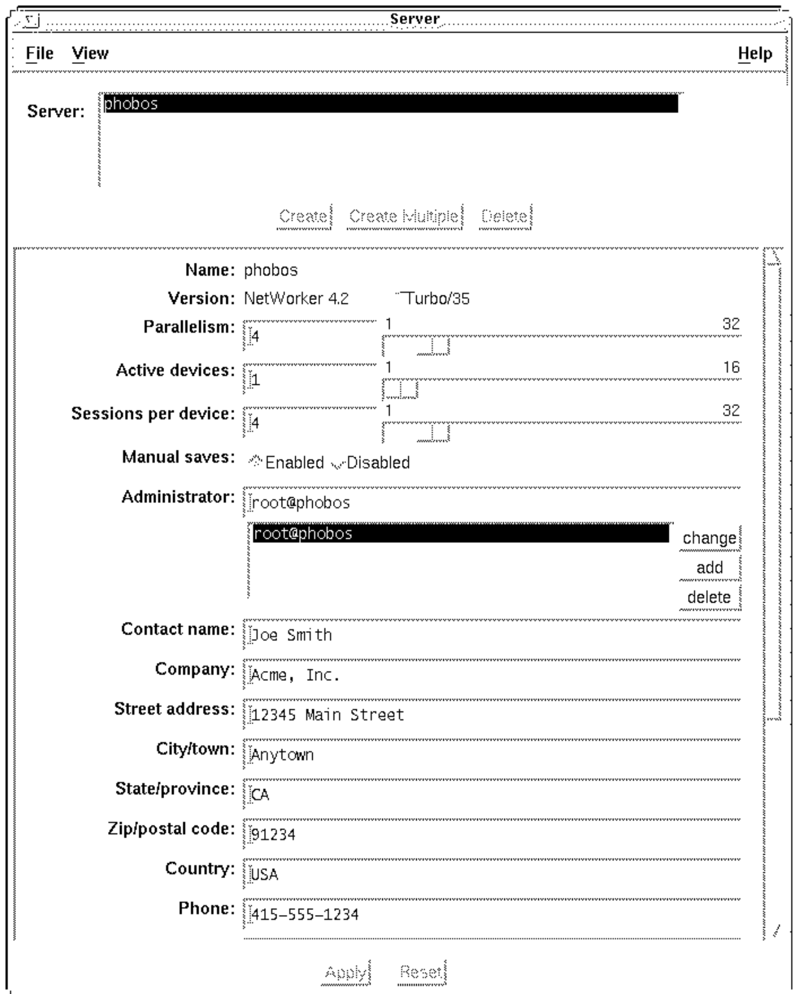

Use the scroll bar on the right side of the window, if necessary, to view the entire contents of the Server window.

The Server window contains the following items.

- Name field – displays the name of the current Backup server.

Note – If you have more than one Backup server on your network, you can use the Change Server speedbar button or command to change to a different server.

- Version field – displays the version of the Backup software currently installed on the Backup server.
- Parallelism field – displays the number of clients backing up in parallel. Use the mouse to slide the bar to the left or right to set the parallelism, or enter the number in the Parallelism field. The maximum value for parallelism varies, depending upon which Backup product you purchased.
- Active devices field – contains the maximum number of media devices Backup can use concurrently. An active device is one that the Backup server uses for backups and recovers. Use the slide bar to set the value, or enter the number in the Active devices field.

Note – The Active devices and Sessions per device fields are only available if you purchase Backup TurboPak. See the section "Using Multiple Devices Concurrently" on page 72 for more information.

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

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

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

The remainder of the window contains fields for entering customer information required for registering Backup. Refer to Chapter 5 in the *Solstice Backup 4.2 Installation and Maintenance Guide* for more information about registering Backup.

▼ Changing Parallelism

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

Note – The maximum value for parallelism varies, depending upon which Backup product you purchased.

You may want to adjust the number of clients backing up simultaneously if the server:

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

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

1. **Open the Server window by selecting Server Setup from the Server 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. **Click the Apply button.**

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

Note – To "turn off" the parallelism so that only one client can back up to the server at a time, set the Parallelism value to one.

Using Multiple Devices Concurrently

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

Note – You must purchase the Backup TurboPak to back up to multiple devices concurrently.

Enter the maximum number of active devices and the number of sessions per device in the Server window.

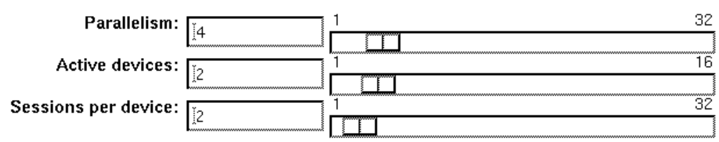

A save session is generated when a client starts a backup. A single client can simultaneously generate several save sessions.

The Active devices value sets the maximum number of backup devices Backup uses concurrently for backups. An active device is one the Backup server uses for backups or recovers.

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

Note – The Parallelism, Active devices, and Sessions per device values control the backup traffic from the Backup server to the media devices.

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

1. **Open the Server window by selecting Server Setup from the Server menu.**
2. **Set the Parallelism value by moving the slide bar to the right or left to display the value you want.**
3. **Set the Active devices by moving the slide bar to the right or left to display the value you want.**
4. **Set the Sessions per device by moving the slide bar it to the right or left to display the value you want.**
5. **Click the Apply button.**

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

▼ Adding or Changing Administrators

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

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

1. **Open the Server window by selecting Server Setup from the Server menu.**
2. **Highlight the current entry in the Administrator field.**

3. **Enter a new user group or *netgroup* name or make changes to the existing one. Netgroup names must be preceded by an ampersand (&) in the Administrator field.**
   You can restrict administrator access to a particular machine by listing it as shown in the following example:

```
root@server_name
```

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

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

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

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

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

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

▼ Changing the Backup Server

This section describes how to change to another Backup server on a multi-server network.

To open the Change Server dialog box, click the Change Server speedbar button or select Change from the Server pull-down menu.

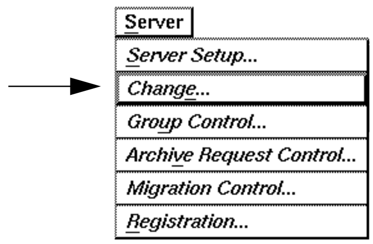

The Change Server dialog box appears, as shown.

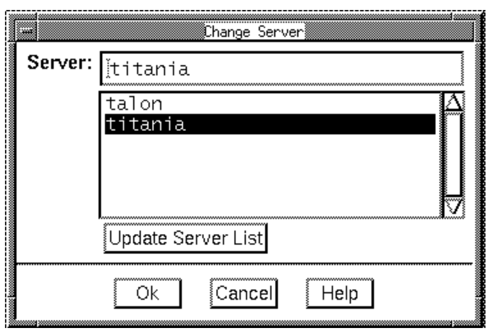

Follow these steps to change to another Backup server:

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

Caution – When you click the Update Server List button, Backup searches for every server available to your machine. If you have a large network, this search could take awhile.

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

Monitoring and Controlling Backups

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

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

- monitor a scheduled backup
- view the details of a scheduled backup
- preview a backup
- immediately start a scheduled backup
- stop a scheduled backup
- restart a scheduled backup

Navigating the Group Control Windows

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

Open the main Group Control window by clicking the Group Control speedbar button or selecting the Group Control command from the Server pull-down menu.

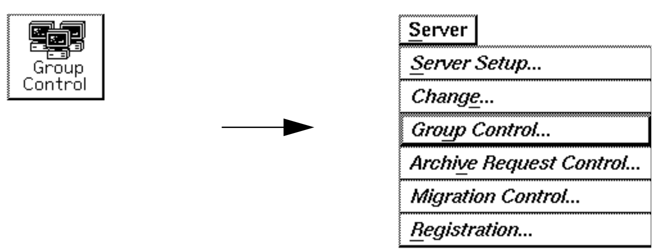

The Group Control window appears, as shown.

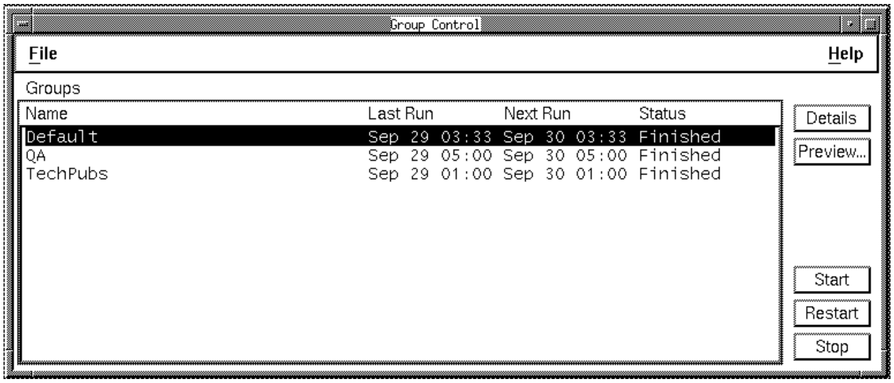

The Groups scrolling list displays information about all of the backup groups known to the Backup server.

- Name – displays the name of the backup group.
- Last Run – lists the date and time the group was last backed up.
- Next Run – shows the date and time for the next enabled scheduled backup.
- Status – displays one of the following messages about the progress of your backup:
  Running – the group is currently backing up

  Never Run – the group has never been backed up

  Finished – the backup has finished

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

  Preview Run– the preview of a backup is running

- Start button – click to start a back up immediately.
- Restart button – click to restart a backup that has been stopped.
- Stop button – click to stop a scheduled backup.
- Details button – click to view details of completed backups.
- Preview button – click to preview an upcoming backup.

Note – Use the Groups window to create new groups and to enable and set the backup time for a group.

The Group Control Details Dialog Box

To view more detailed information about a completed group backup, click the Details button in the Group Control dialog box.

Backup displays the Group Control Details dialog box, as shown.

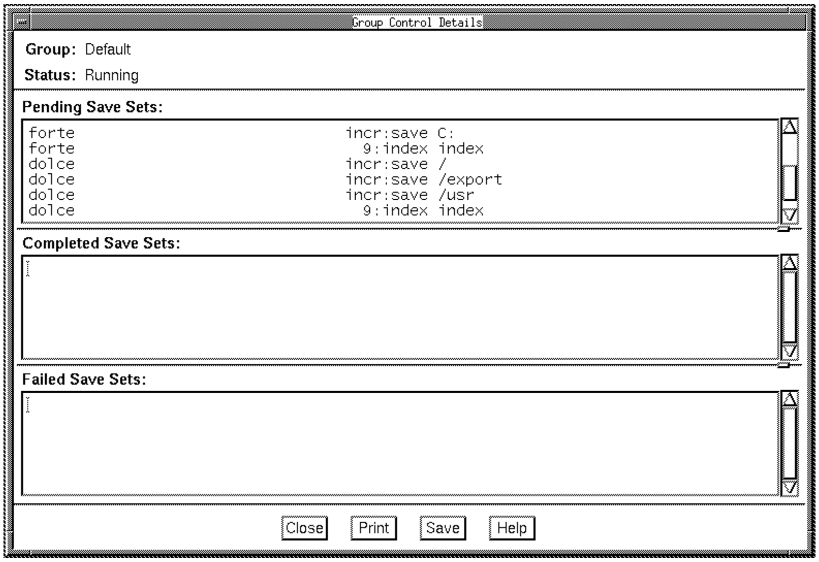

There are three message fields in the Group Control Details dialog box.

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

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

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

Typical reasons for failed save sets include the following:

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

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

Previewing a Group Backup

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

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

▼ Starting a Scheduled Backup Immediately

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

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

1. **Select the group you want to back up immediately from the Groups scrolling list.**

2. **Click the Start button.**
   A confirmation dialog box appears, as shown.

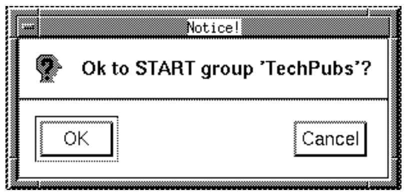

3. **Click OK to start the backup or Cancel to discontinue.**

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

Note – You can also start an immediate backup by selecting the Start now button in the Groups window available from the Customize pull-down menu.

▼ Stopping a Backup Group

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

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

1. **Select the group you want to stop from the Groups scrolling list.**

2. **Click the Stop button.**
   A confirmation dialog box appears, as shown.

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

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

▼ Restarting a Backup Group

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

Follow these steps to restart a backup:

1. **Select the name of the halted backup group from the Groups scrolling list in the Group Control window.**

2. **Click the Restart button.**
   A confirmation dialog box appears, as shown.

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

Backup resumes the scheduled backup for the group and displays "running" in the Status field.

Handling Open Files

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

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

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

Archive Request Control

The Solstice Backup Archive Application provides file archiving and retrieval services to a range of client machines. If you purchased this optional extension for your Backup server, the Archive Request Control command on the Server pull-down menu is enabled.

For complete instructions on using the Backup Archive Application, see Chapter 9, "Archive Application."

Migration Control

The Solstice Backup Hierarchical Storage Management (HSM) Application provides file migration and recall services to a range of client machines. If you purchased this optional extension for your Backup server, the Migration Control command on the Server pull-down menu is enabled.

For complete instructions on using the Backup HSM Application, see Chapter 10, "Hierarchical Storage Management Application."

Registering Backup

After installing and enabling the Backup server and client software and any new clients or options, you must register the software with SunSoft immediately. Use the Registration window to properly register your Backup software.

If you do not register the software, Backup automatically disables itself 45 days after being enabled. Once the software disables itself, you will not be able to back up any more files. You will, however, be able to recover previously backed-up files.

Note – Refer to the section "Registering Backup Products" in the *Installation and Maintenance Guide* for a complete description of the Registration window and the registration process.

Managing Media and Backup Devices 5

This chapter discusses the Media pull-down menu commands used for managing media and stand-alone backup devices. The Media pull-down menu commands used for accomplishing operations specific to a jukebox are described in Chapter 8, "Jukebox Software Module."

This chapter describes the following operations:

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

Note – If you back up and recover to multiple backup devices, you need to purchase Solstice Backup Server Edition or Solstice Backup plus TurboPak. If you back up to a jukebox, you also need to purchase the Jukebox Software Module.

Labeling and Mounting Backup Volumes

This section applies only if you label and mount volumes in a stand-alone backup device. If you need to label volumes for a jukebox see Chapter 8, "Jukebox Software Module."

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

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

Navigating the Label Dialog Box

Use the Label dialog box to label backup volumes in a stand-alone device.

To open the Label dialog box click the Label speedbar button or choose the Label command from the Media pull-down menu.

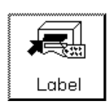

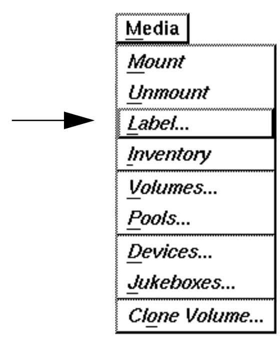

The Label dialog box appears, as shown.

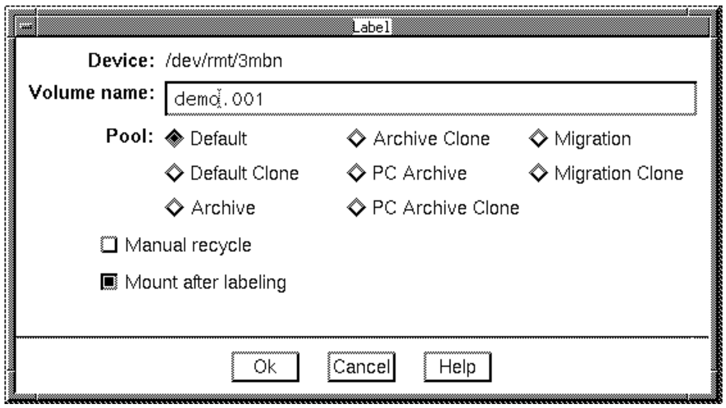

The following describes the contents of the Label dialog box.

- Device field – contains the name of the selected device. To select another device, click Cancel to close the Label dialog box, highlight the new device in the Backup Administrator window, and reselect the Label command.
- Volume name field – contains the label that will be assigned to the volume and displays the name of the volume mounted in the backup device. Backup automatically assigns the next sequential label from the label template associated with the chosen pool. If you do not want Backup to use the preselected label template, delete the contents in the Volume name field and enter a new name for the backup volume.
- Pool choices – choose a volume pool from the Pool choices. Only pools that have been enabled in the Pools window appear in the Label dialog box (most of Backup's preconfigured pools are automatically enabled). Backup automatically uses the label template associated with the selected pool to label the backup volume.

Note – For more information on volume pools, see the section "Using Volume Pools" in this chapter.

- Manual recycle toggle button – select if you want to manually recycle a backup volume. When a backup volume is marked "manual recycle," Backup disregards the assigned retention policy. Manual recycle mode is explained in the section "Volumes Scrolling List" in this chapter.
- Mount after labeling toggle button – selected by default. Deselect if you only want to label volumes and not mount them.

Caution – You may prelabel as many volumes as you like at one time without mounting them by deselecting the Mount after labeling toggle button.

- Ok button – click to label or mount a backup volume. If you make a mistake, click the Cancel button.

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

▼ Labeling and Mounting in One Operation

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

1. **Place a blank volume in the Backup server backup device.**
2. **Highlight the device in the Devices display in the Backup Administrator window.**
3. **Open the Label dialog box by clicking the Label speedbar button or selecting Label from the Media menu.**
4. **Choose a volume pool from the Pool choices.**
5. **Delete the name from the Volume name field if you want a different label than the one chosen by the label template, then enter the new name.**
6. **Make your selection for Manual recycle and click the Mount after labeling toggle button.**
7. **Click the Ok button when you are ready to label and mount the volume. If you make a mistake, click the Cancel button to cancel the operation.**
   A dialog box appears, asking you to confirm your request. Click OK to confirm the labeling request or Cancel to quit the request.

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

Caution – Relabeling a backup volume destroys all records of its contents under the old label. That is why Backup asks for confirmation with the question "About to relabel." For more information about relabeling backup volumes, see "Manually Managing the Online Indexes" in Chapter 3.

Understanding How to Label Backup Volumes

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

Note – Use a jukebox to automate volume labeling and mounting operations. See Chapter 8, "Jukebox Software Module," for more information.

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

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

Note – For more information on volume pools, read the section "Using Volume Pools" on page 109. For information about label templates, see the section "Using Label Templates" in Chapter 7.

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

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

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

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

atlas.001	this backup volume is from the Backup server named *atlas*
mars.010	this backup volume is from the Backup server named *mars*

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

Mounting and Unmounting Backup Volumes

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

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

Use the speedbar buttons to mount and unmount backup volumes or the Mount and Unmount commands in the Media pull-down menu.

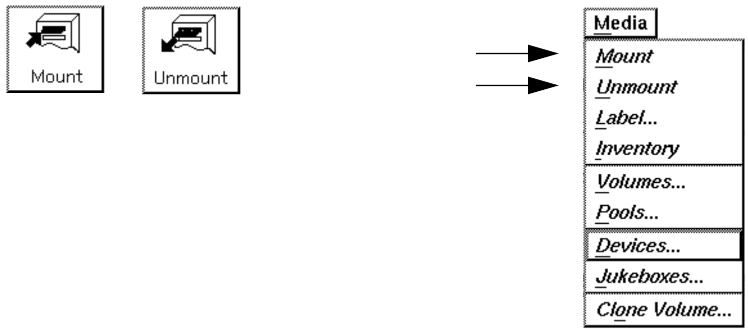

▼ Mounting Backup Volumes

When a backup volume is required for a stand-alone device, Backup uses the following hierarchy to select a volume from the appropriate backup volume pool:

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

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

1. **Select the device name in the Devices display of the Backup Administrator window.**

2. **Click the Mount speedbar button.**

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

Caution – To perform an unattended backup using a stand-alone device, you must premount backup volumes.

If you try to back up files when a backup volume is not mounted, Backup 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
```

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

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

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

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

▼ Unmounting Backup Volumes

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

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

Finding a Backup Volume Name

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

First method:

Click the Mount speedbar button to mount the volume in the backup device; the volume name appears next to the device name in the Devices display of the Backup Administrator window.

Second method:

Load the volume in a device and select the Label speedbar button. When the Label dialog box appears, the name of the label appears in the Volume name field. Once you determine the volume name, click Cancel to avoid accidentally renaming the volume.

Automatically Managing Backup Volumes

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

Note – For more information about browse and retention policies, see "Creating Index Policies" in Chapter 7.

Manually Managing Backup Volumes

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

Navigating the Volumes Window

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

To view information about your backup volumes, open the Volumes window by clicking the Volumes speedbar button or selecting Volumes from the Media pull-down menu.

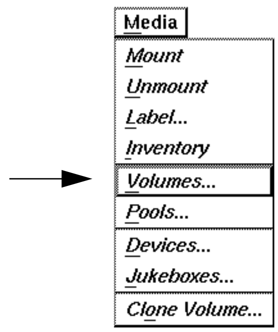

The Volumes window appears, as shown.

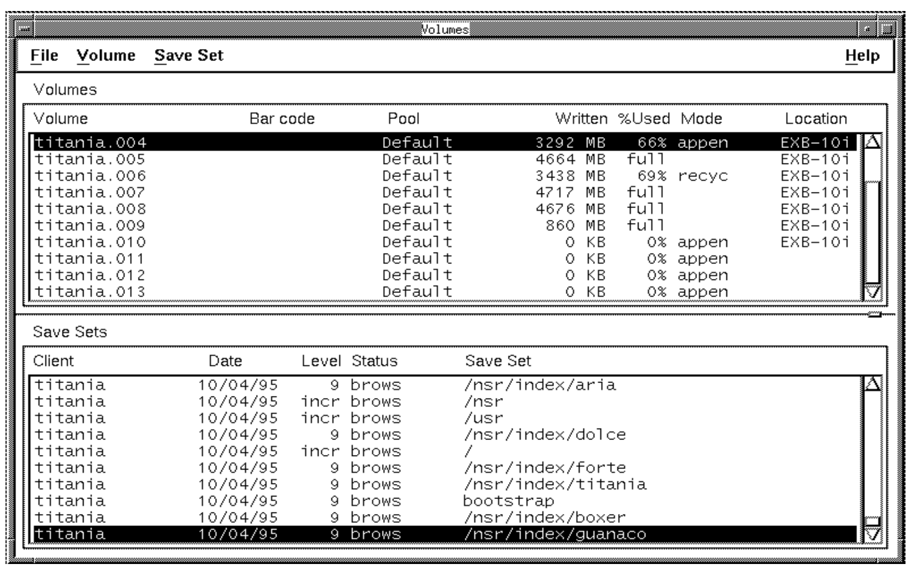

The Volumes window displays the following information.

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

Volumes Scrolling List

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

- Volume – backup volume name. One of the following designations may appear at the end of the volume name:
  - (A) – indicates an archive volume
  - (R) – indicates a read only volume
- Bar code – bar code label, if present
- Pool – name of the pool to which the backup volume belongs
- Written – amount of data written on the volume
- %Used – percentage of the estimated total capacity of the backup volume used. When 100% appears in the %Used column, the volume has met or exceeded its expected capacity. However, space may still be available because the end of tape mark has not been reached.
  The %Used column displays full when there is no more space on the backup volume for additional data and the end of tape mark has been reached. If the save sets on the volume have not yet passed the time period specified by their browse and retention policies, the Mode column is blank. If the save sets have passed their browse and retention policies, recyc appears in the Mode column, indicating the volume is available for recycling.

- Mode – status or kind of backup volume.
  - appen – (appendable) there is room for more data on the backup volume.
  - recyc – (recyclable) all save sets on the volume have passed the browse and retention policies, and the volume is ready to be relabeled. Once the volume has been relabeled, it can be overwritten with new data.
  - man – (manual recycle) both the volume and its save sets are marked for manual recycling. If the appen or recyc modes are already assigned to the volume, the man mode is displayed as an additional mode.
  - The following volume mode indicator may appear next to the volume name in the Volume column:
  - (R) – volume has a mode of "read only" assigned manually with the Change Mode command. When the mode is manually changed to read only, the previously assigned mode is removed and the Mode column is blank for the volume.
  - See "Changing the Mode of a Backup Volume" on page 98 for more information on manually changing the mode of a volume.
- Location – location of the backup volume, entered by the user as a reminder.

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

Save Sets Scrolling List

Use the Save Sets scrolling list, shown below, to view the save sets stored on a volume.

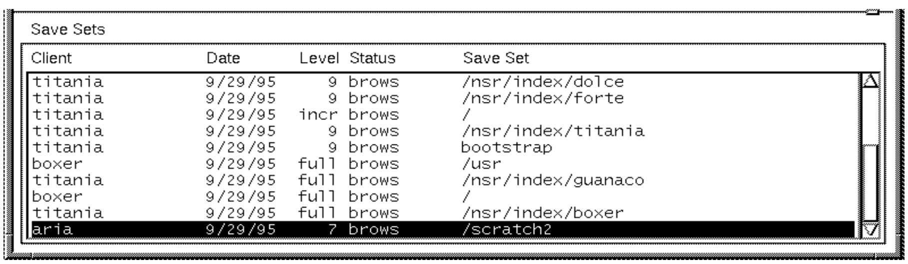
Save Sets

Client	Date	Level	Status	Save Set
titania	9/29/95	9	brows	/nsr/index/dolce
titania	9/29/95	9	brows	/nsr/index/forte
titania	9/29/95	incr	brows	/
titania	9/29/95	9	brows	/nsr/index/titania
titania	9/29/95	9	brows	bootstrap
boxer	9/29/95	full	brows	/usr
titania	9/29/95	full	brows	/nsr/index/guanaco
boxer	9/29/95	full	brows	/
titania	9/29/95	full	brows	/nsr/index/boxer
aria	9/29/95	7	brows	/scratch2

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

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

Save Set Recover window. If you want to recover both the data and the file index entries, use the scanner command. (See the scanner man page for information on using the scanner command.)
  - recyc – (recyclable) indicates that the save set has passed both the Browse and Retention policies time periods and is available for relabeling. The volume must be relabeled before you attempt to overwrite it with new backups.
  - scann – (scanned-in) indicates that the save set was scanned in using the scanner command. Backup's standard index management policies do not apply to scanned data. You must remove these save sets manually from the index. See "Navigating the Indexes Window" in Chapter 3 for instructions on manually removing save sets from the index.
  - inpro – (in-progress) indicates the save set is in the process of being backed up.
  - abort – (aborted) indicates the backup was aborted manually by the administrator or the system crashed.
  - susp – (suspect) indicates that a previous recovery attempt failed. In this case, the susp status is displayed in addition to one of the previously described status types.
- Save Set – path of the filesystem contained in the save set. The save set column also includes clone information. If the save set has a clone, it is marked has clones and the cloned save set is marked cloned *date time.*

Manually Managing Backup Volumes

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

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

Changing the Status of a Save Set

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

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

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

1. **Select the save set from the Save Sets scrolling list in the Volumes window.**

2. **Use the Change Status command in the Save Set pull-down menu in the Volumes window to manually change the status of the save set.**

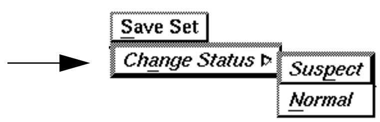

If you want to undo changes made manually to the status of a save set, you can use the Change Status command to return the save set status to its former setting.

Setting the Location for Backup Volumes

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

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

1. **Highlight a backup volume in the Volumes scrolling list.**

2. **Select Set Location from the Volume pull-down menu. The Set Location dialog box appears, as shown.**

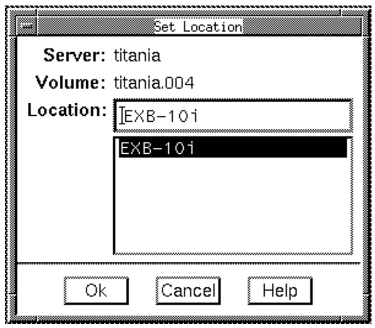

3. **Enter the location of the backup volume. You may want to specify the location of backup volumes stored on a shelf or volumes stored offsite.**

4. **Click the Ok button to designate the location for the selected backup volume. The Volumes window now reflects the location for that backup volume.**

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

Note – If you move backup volumes from one jukebox to another, inventory them so Backup can recognize them and update their location with the name of the new jukebox. If you move volumes offsite, manually assign them the location of the offsite facility before you remove them from the jukebox.

Note – When Backup requests that you mount a backup volume during a recovery, you can find the volume by referring to its location in the Volumes window.

Changing the Mode of a Backup Volume

The mode of a backup volume determines whether Backup can write data to it during a backup. See "Volumes Scrolling List" earlier in this chapter for a description of volume modes.

Backup automatically assigns the appen and recyc modes. This section provides a detailed description of how to manually change volume mode.

The Volume pull-down menu in the Volumes window contains the Change Mode command. This command allows you to manually change the mode of a currently-selected volume to appendable, read only, or recyclable.

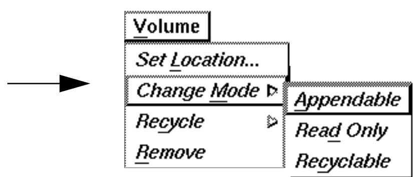

The following table provides a guide to existing volume modes that you can manually change.

Current Volume mode	Change to Appendable	Change to Read Only	Change to Recyclable
Recyclable	yes*	yes	no
Appendable	no	yes	yes

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

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

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

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

1. **Click the Volumes speedbar button or select the Volumes command from the Media pull-down menu.**

2. **Highlight the volume in the Volumes scrolling list.**

3. **Select Change Mode from the Volume pull-down menu in the Volumes window.**

4. **Drag and release on a mode from the Change Mode choices: Appendable, Read Only, or Recyclable.**
   A confirmation box appears asking if you want to change the volume to the mode you have chosen, as shown.

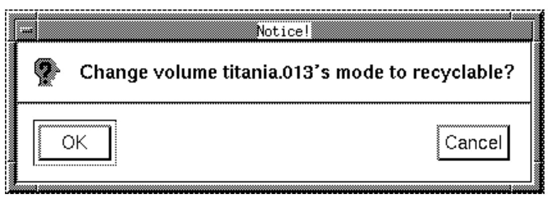

5. **Click the Ok button if you want to change the mode of the selected volume.**
   Notice the mode of the volume changes in the Volumes scrolling list. Click Cancel to leave the mode unchanged.

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

Caution – Do not manually change the mode of a volume to recyclable if you still need any of the save sets on the volume and do not have a cloned volume.

If a cloned volume does not exist when you manually change the mode of a volume to recyclable, all of the save sets on the volume change to a status of recyclable, regardless of the save set browse or retention policies. If a cloned volume exists, the Status of the save sets remains unchanged.

Original backup volumes and cloned volumes share the same set of save set entries. As long as one of the volumes exist, Backup does not remove save set entries, since they are still needed for the remaining volume.

Recycling Volumes

You can override the volume retention policy by changing it to manual recycle. Select Recycle from the Volume pull-down menu. The Manual and Auto commands appear, as shown.

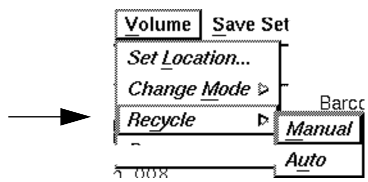

You may want to change to manual recycle if you need to keep the save sets on a volume longer than its retention policy specifies. You can change a volume marked manual recycle back to automatic recycle, so the volume will once again use its original retention policy.

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

1. **Select a volume from the Volumes scrolling list.**

2. **Select the Recycle command from the Volume pull-down menu in the Volumes window; the Manual and Auto commands appear.**

3. **Select Manual from the Recycle menu. Backup identifies the volume by adding "man" to the volume mode in the Volumes scrolling list.**

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

Removing a Backup Volume

To remove a volume from the media index, select Remove from the Volume pull-down menu in the Volumes window. Typically, you should only remove a backup volume from the media index if the volume has become physically damaged or unusable.

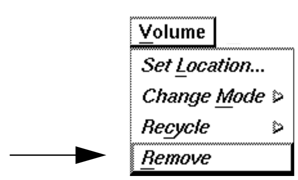

When you select a volume for removal, Backup checks to see if the volume has a clone. If the volume has a clone, Backup does not remove the entries from the online file index. If the volume does not have a clone, Backup removes the entries from the online file index.

Note – If a backup volume is lost or destroyed, you should remove its entries from both the file and media indexes. Otherwise, Backup may request the nonexistent volume for data recovery.

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

1. **Highlight the volume in the Volumes scrolling list.**
   The volume save sets appear in the Save Sets scrolling list.

2. **Select Remove from the Volume menu.**
   The notice "Remove volume *name* from the media index?" appears, as shown.

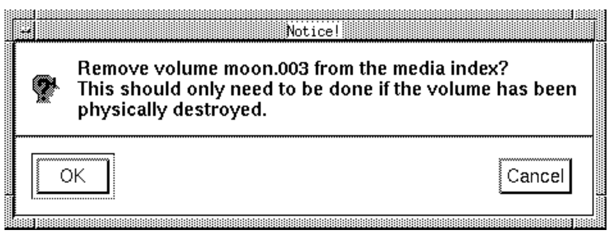

3. **Click the Ok button to remove the volume from the media index.**
   If a removed volume does not have a clone, the save sets are also removed.

4. **Click the Cancel button if you do not want to remove the volume from the media index.**

Caution – If you accidentally delete a backup volume from the media index and have not relabeled it yet, you can still recover the data by using the scanner program. See the man pages for more information on the scanner program.

5

Cloning Backup Volumes

Backup protects the data on your network by backing up files on a daily basis. Cloning provides additional protection by allowing you to make identical copies of the data on your backup volumes on an as-needed basis.

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

You must have all of the following to clone save sets on your backup volumes with Backup:

- Backup, Server Edition or Backup, Network Edition plus TurboPak
- multiple backup devices
- Active devices value in the Server window set to 2 or more
- source and destination volumes

Note – Backup also has the ability to clone individual save sets manually and automatically. See "Cloning Save Sets" in Chapter 6 for more information.

▼ Navigating the Clone Volume Windows

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

To open the Clone Volume window, select the Clone Volume command from the Media pull-down menu.

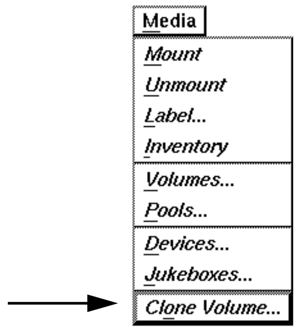

The Clone Volume window appears, as shown.

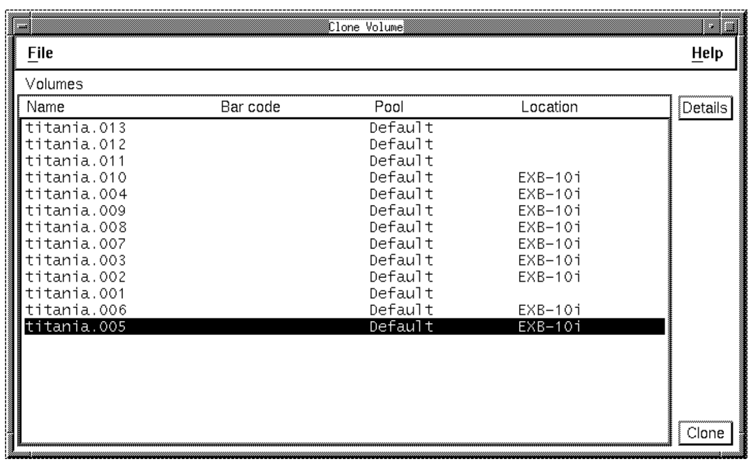

The following describes the contents of the Clone Volume window.

- Volumes scrolling list – contains a list of all backup volumes known to the server. The list includes the volume and pool name, bar code label, and the volume location. You can clone save sets from an unlimited number of backup volumes at the same time. Clone backup volumes after a scheduled backup has completed to avoid overloading your Backup server and backup devices.

- Details button – click to view backup volume details.
  The Clone Volume Details dialog box appears, as shown.

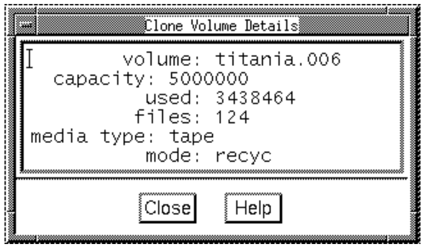

  The Clone Volume Details dialog box contains the following items:

  - volume – name of volume.
  - capacity – size of the media.
  - used – space used on the backup volume.
  - files – number of files on the backup volume.
  - media type – kind or type of backup volume.
  - mode – mode of a backup volume is either appen (appendable), full (full), or recyc (recyclable). See “Viewing the Volumes Scrolling List” in this chapter for an explanation of the different volume modes.

- Clone button – click to open the Clone Volume Status dialog box to begin cloning the selected volumes.
  The Clone Volume Status dialog box appears, as shown.

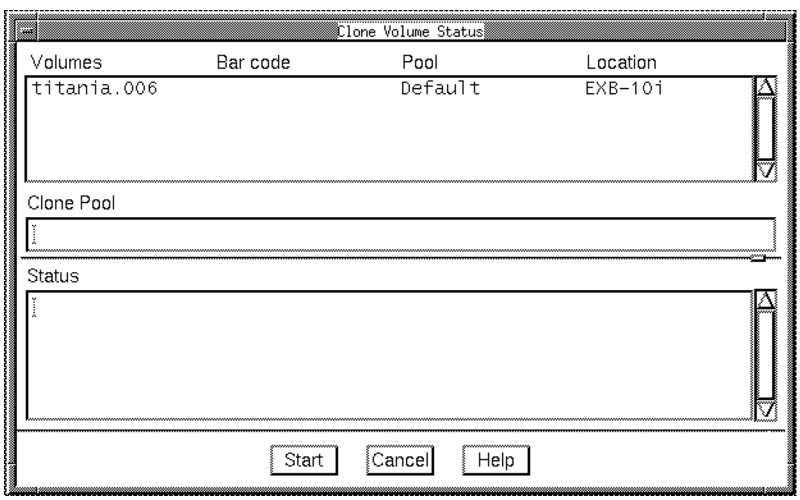

The Clone Volume Status window contains the following items:

- Volumes scrolling list – lists the volume(s) selected for cloning, the pool to which it belongs, and the volume location.
- Clone Pool field – enter the name of the clone pool. Cloned backup volumes must belong to a clone pool, either one you have created or one of the preconfigured clone pools. If you do not enter a pool name, Backup automatically selects the Default Clone pool.
- Start button – click to begin cloning.
- Status scrolling list – displays the progress of the cloning operation. The same messages appear in the Backup Administrator window.

Example – Cloning Backup Volumes

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

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

1. **Load, label, and mount a blank destination volume in the media drive.**

Note – If you use a stand-alone device for the Backup software backups, a message appears in the Messages display of the Backup Administrator window requesting that you mount a volume. If you use a jukebox for Backup backups and do not load, label, and mount a volume, Backup selects an “appendable” volume from the clone pool entered in the Clone Status window.

---

2. **Open the Clone Volume window by choosing the Clone Volume command from the Media menu.**

3. **Select the volumes containing full backups for the engineering department from the Volumes scrolling list.**

4. **Click the Clone button in the Clone Volume window to open the Clone Volume Status window.**

5. **Enter the name of the clone pool in the Clone Pool field.**
6. **Click the Start button in the Clone Volume Status window to begin the cloning process.**

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

Understanding the Volume Cloning Feature

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

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

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

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

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

Note – Cloned media does not contain the file or media indexes used for recovering data. Since an offsite facility using Backup will not have these indexes, they must use the scanner command to recover the cloned data. For complete information about the scanner command, see the man pages.

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

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

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

Entering Index Entries for Cloned Data

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

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

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

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

Note – For more information about cloning data see "Understanding the Cloning Save Sets Feature" in Chapter 6.

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

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

Recovering from Cloned Volumes

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

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

Backup considers volumes for recovery in the following order:

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

Using Volume Pools

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

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

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

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

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

Backup provides the following preconfigured pools you can begin using right away: Archive, Archive Clone, Default, Default Clone, Full, NonFull, Migration, Migration Clone, PC Archive, PC Archive Clone, and Offsite. These pools and any pools you create fall into the following “pool type” categories: Backup, Backup Clone, Archive, Archive Clone, Migration, and Migration Clone.

Note – For information about the preconfigured pools provided by Backup, refer to Appendix A in the *Solstice Backup 4.2 Installation and Maintenance Guide*.

All Backup pools require a label template, described in this chapter in the section “Using Label Templates with Volume Pools” on page 121. Before creating a new pool, create a label template with the same name. You should 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.

Caution – For each pool, you must select at least one of the following choices: a group, client, save set, or level. If you do not make a selection for at least one of these items, you will not be able to create a pool.

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

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

You can further manage backed-up data:

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

Navigating the Pools Window

To open the Pools window, select the Pools command from the Media pull-down menu.

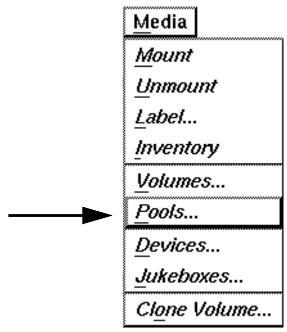

The Pools window appears, as shown.

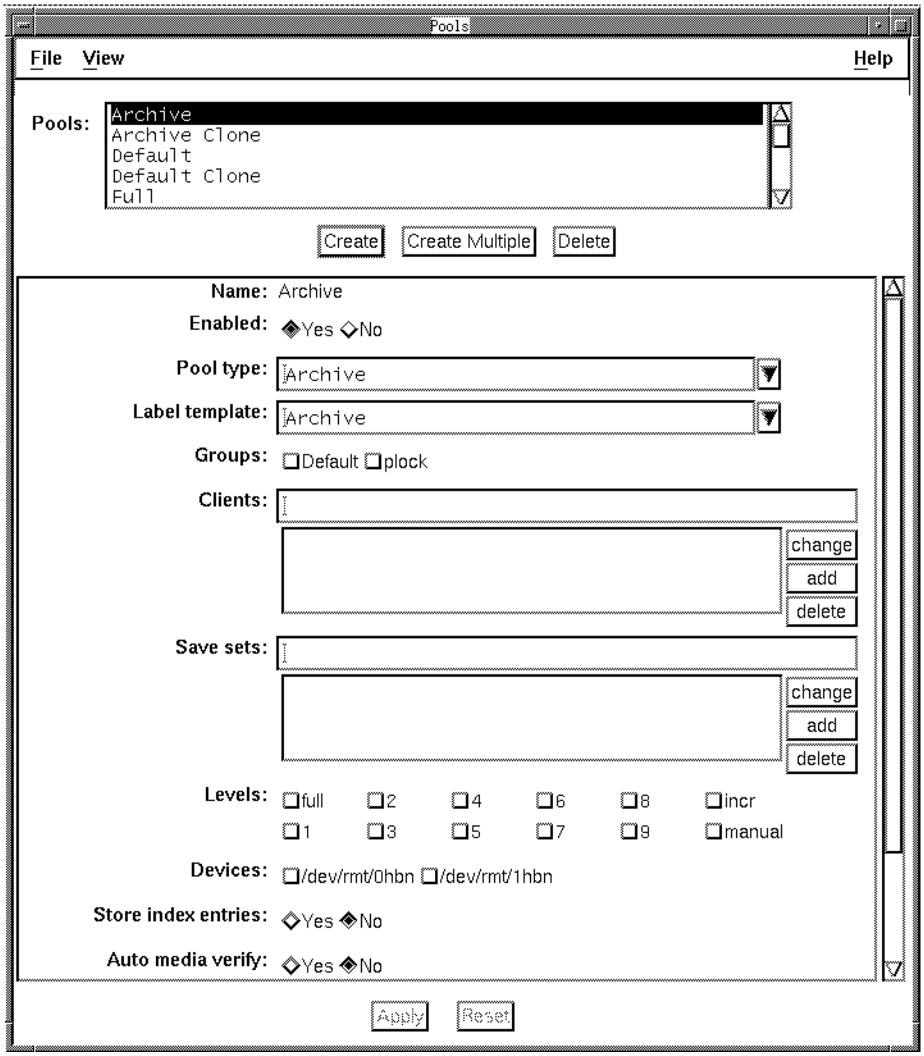

Use the scroll bar or resize the window to view the entire Pools window. The Pools window contains choices required for sorting client data to selected backup volumes.

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

Caution – You cannot create pools that share identical settings for Pool type, Groups, Clients, or Save Sets. If any of these settings for a new pool match an existing pool, you receive a warning message similar to the following:

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

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

Note – You cannot delete or modify any of Backup's preconfigured pools.

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

Note – With the exception of the Full, NonFull, and Offsite pools, Backup's preconfigured pools are already enabled.

- Pool type choices – select the appropriate type of pool for your backup volumes. The preconfigured pools include: Archive, Archive Clone, Backup, Backup Clone, Migration, and Migration Clone.
- Label template choices – select a label template, either a preconfigured label template or one created in the Label Templates window. Use the same name for both the label template and the pool.
- Groups choices – choose the groups you want included in the pool. The groups displayed represent all of the groups known to the Backup server. There is no limitation on how many groups you can include in a pool. You can also select individual clients for a volume pool without selecting a group. If you do not make a selection from the Groups choices, Backup does not sort data based on a group.
- Clients field – entering a client name in this field is optional. You may want to enter the name of a Backup client:
  - to add one or more clients to a pool without selecting a group.
  - to pick a specific client belonging to a group. If you pick a specific Backup client from a group, you exclude all the other clients from that group. If you select a group, the client you select must belong to that group.
- Save sets field – entering a save set in this field is optional. You may want to enter the name of a client save set:
  - to send a specific filesystem, from a specific client, to a volume pool without selecting a group (see “Example 3 – Secure Device” on page 118).
  - to send a specific filesystem to a pool (see “Example 5 – Specific Filesystem” on page 119” in this chapter).
- Levels choices – select the backup level(s) to associate with the media pool. For example, if you click full, the pool will only include full backups for the groups assigned to the pool. See “Selecting the Backup Levels“ in Chapter 7 for more information on backup levels.
- Devices choices – displays the devices recognized by the Backup server. This is where you indicate which device to use for the selected pool. For example, you might want one pool for optical disks and one for tapes, if you have both backup devices on the same server. If you do not choose a device, Backup uses any available device.
- Store file index entries choices – allows you to choose, with Yes and No choices, whether or not you want to include the index entries from the backup in the online index. If you choose No, the index entries will not be included in the file index. However, an entry for the save set is still made in the media database. For example, if you are using the preconfigured pool

Offsite, you may not want the index entries online. You must select No for the Archive, Archive Clone, Default Clone and PC Clone pools. You must select Yes for the Migration and Migration Clone pools.

Caution – If you select No for Store file index entries, you will not have online file index entries available for browsing in the Backup Recover window. To recover data without file index entries, you must use either the Save Set Recover window or the scanner command. All parts of the save set must be intact in order for scanner to recover the data. See the man pages for more information on the scanner command.

- Auto media verify choices – click Yes to enable automatic verification of your backup media. Backup compares a record near the end of the tape with the data stored on disk. If Backup detects a discrepancy, the backup terminates and the status of the offending volume is marked as full. The savegroup completion message indicates whether or not media errors occurred during the backup.
- Apply button – click to apply all of the selections you have chosen for a particular pool.
- Reset button – click to restore your original selections and override any new selections you may have chosen.

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

Examples

This section contains step by step instructions for creating a new pool and several suggestions for why you would want to create your own volume pools. If you are already familiar with the Pools feature, skip to the shortcut directions on page 119.

Example 1 – Creating a New Pool

Suppose you want to create a pool named “Test” that selects all of the full backups from a group called “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. You do not expect that you will need to recover this data, but plan to keep it for a year.

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

1. **Open the Label Templates window by selecting Label Templates from the Customize menu. See the section "Using Label Templates" in Chapter 7 for complete instructions for creating a label template.**
   The Label Templates window appears similar to the following after you create the Test label template.

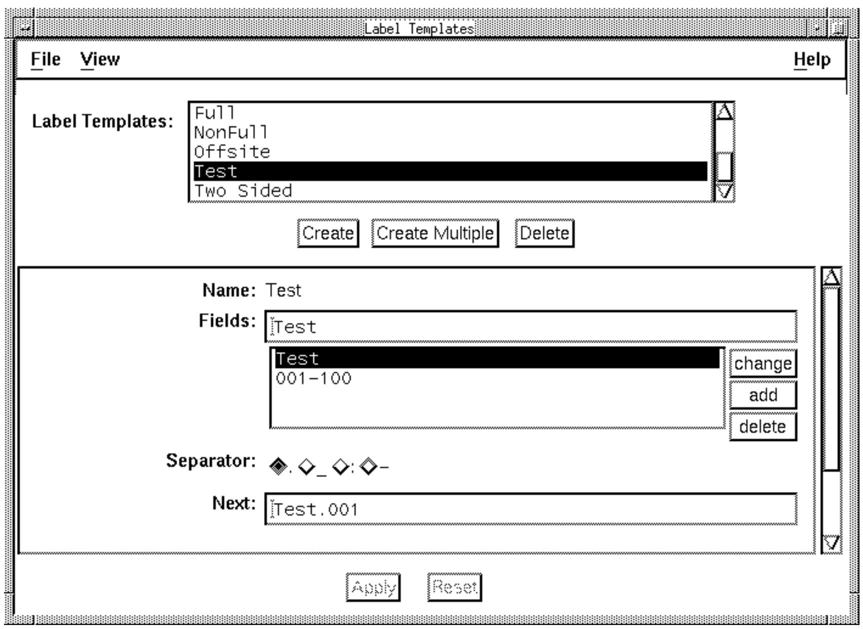

   After creating the label template, open the Pools window and create a pool named "Test."

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

Note – You cannot create a new pool if Backup is busy with a backup process – this ensures that Backup will send the backup data to the correct pool.

---

3. **Select Yes from the Enabled choices.**
4. **Select Backup from the Pool type choices.**
5. **Choose the label template Test, which you previously created.**
6. **Choose the group QA, because it contains the clients you want to include in the Test pool.**
7. **Do not enter anything in the Clients or Save sets fields, since you do not need to selectively include any clients or save sets.**
8. **Select full from the Levels choices.**
9. **Select a backup device from the Devices choices.**
10. **Select Yes from the Store index entries choices.**
11. **Select No from the Auto media verify choices.**
12. **Click the Apply button to apply the selections.**

The Pools window looks like the following after you create the pool Test.

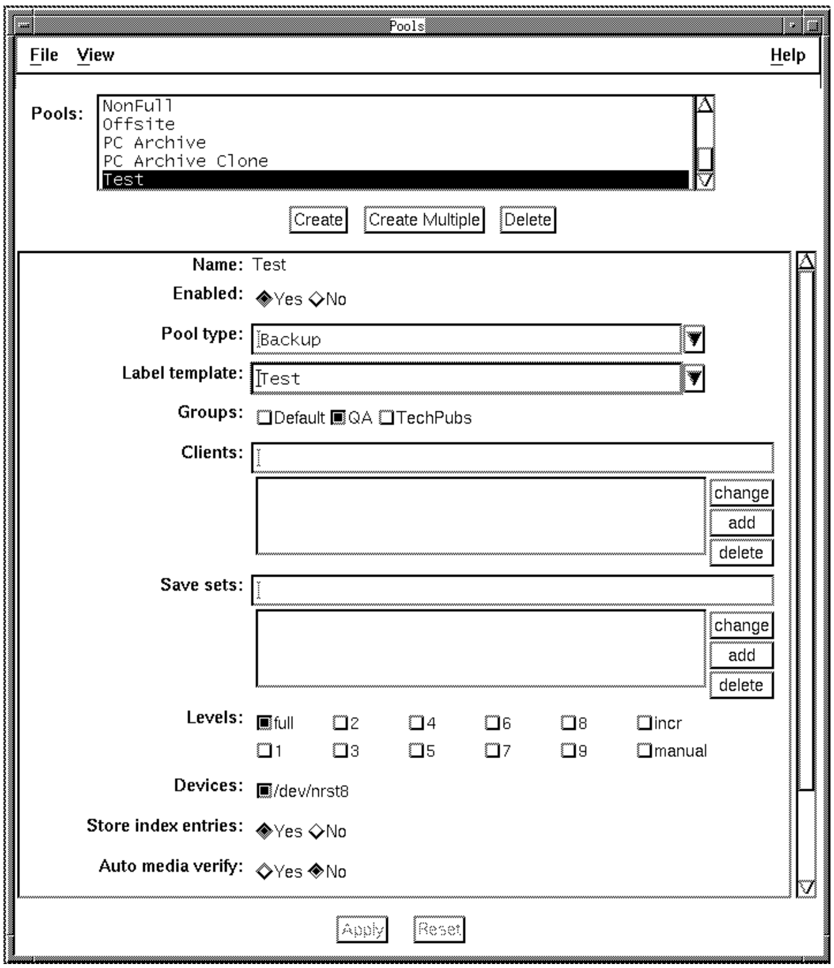

Once you create a new pool, you can easily modify it by choosing the pool from the Pools scrolling list and changing the current selections. When you are done making the modifications, click the Apply button. If you change your mind and do not want to implement the new selections, click the Reset button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

Example 2– Separate Departments

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

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

Example 3 – Secure Device

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

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

Example 4 – Targeted Backups

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

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

Example 5 – Specific Filesystem

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

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 will back up all electronic mail for all clients. You enter the following in the Save sets field:

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

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

Shortcut

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

To create a volume pool, follow these steps:

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

9. **Click the Apply button to apply the selections.**

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

▼ Deleting a Volume Pool

To delete a pool, follow these steps:

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

Caution – Give careful consideration to the pools you create. Once you create a pool, you cannot rename it nor delete it, as long as it contains backup volumes.

Understanding Volume Pools

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

By contrast, volume pools establish a logical and systematic method for tracking, organizing, and sorting your data and backup volumes. For example, you may want to create a separate pool for each department in your company or a pool 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 offsite. Volume pools allow you to select and sort specific files or data to preselected backup volumes.

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

Using Label Templates with Volume Pools

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

Adding and Deleting Backup Devices

With some versions of Backup, the installation script requires that you enter at least one device *pathname* for a Backup backup device. Since Backup supports multiple devices (either in a series, concurrently, or in a jukebox), you can add more devices. If your version of the server software supports more than one backup device in a series, Backup backs up to the volume in one device until it is full, then backs up to the volume in another device until it is full, and so on.

Use concurrent devices to spread your backups evenly across your backup devices. For unattended backups on a jukebox, use the optional Backup Jukebox Software Module.

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

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

DLT

half-inch magnetic tape drives

quarter-inch cartridge tape drives

4mm (DAT) drives

8mm tape drives

8mm 5 gigabyte tape drives

3480 tape drives

VHS tape drives

optical disk drives

Navigating the Devices Window

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

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

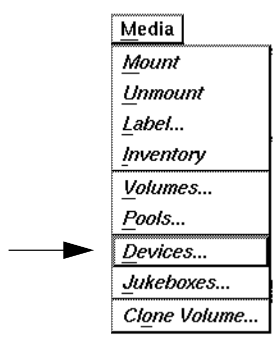

The Devices window appears, as shown.

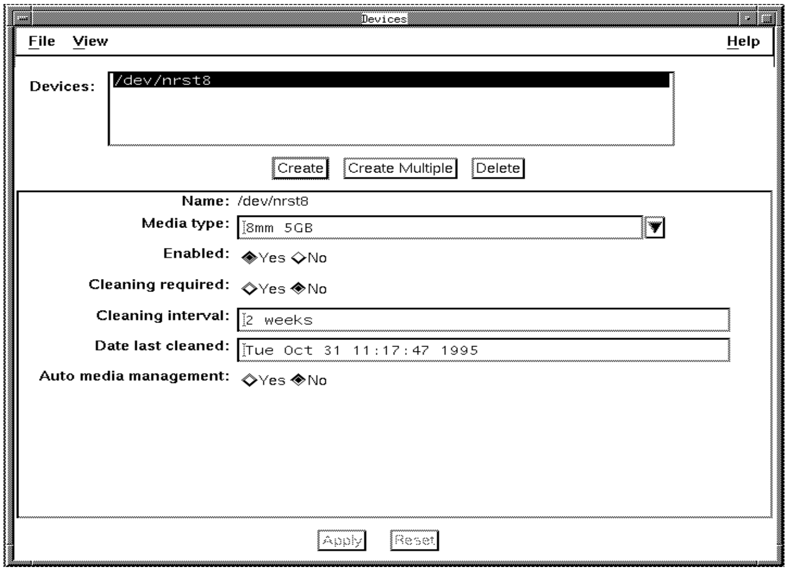

The Devices window contains the following items:

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

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

The choices represent media for half-inch magnetic tape drives (himt), quarter-inch cartridge tape drives (qic), four millimeter tape drives (4mm), eight millimeter tape drives (8mm), eight millimeter five gigabyte tape drives (8mm 5GB), 3480 tape drives (3480), digital linear tape drives (dlt), VHS tape drives (vhs), and optical disk drives (optical).

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

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

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

- Cleaning required choices – automatically set to Yes when Backup determines that a jukebox device needs to be cleaned and set to No after the device has been cleaned.
  A user schedules future cleaning of a device by setting this attribute to Yes. When the value of Cleaning required changes from Yes to No, the value of Date last cleaned is automatically set to the current date and time. If required, you can manually enter a value for Date last cleaned.

- Cleaning interval field – specified number of days, weeks, or months between the Date last cleaned and the next scheduled cleaning of the device. This field accepts: 2days, 2 days, 2weeks, 2 weeks (any combination of numerals followed by day, days, month, or months, with or without a space between).

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

- Date last cleaned field – records the date a device was last cleaned and is updated whenever a device is cleaned.

- Auto media management field – gives Backup exclusive control over media loaded in the backup device. If you select Yes, Backup automatically labels, mounts, and overwrites a volume it considers unlabeled. If you select No, Backup ignores the media and does not consider it for backup.

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

▼ Adding a New Backup Device

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

- the *pathname* of the device
  The *pathname* for optical devices varies depending upon the manufacturer. For example, for AIX servers the *pathname* for a no-rewind tape device is similar to */dev/rmt0.1*.
- the media type of the device

Caution – A no-rewind backup device is required (for tape drives) because Backup writes a file mark on the backup volume at the end of each backup and then appends data onto the volume based on the position of the file mark. If the device rewinds the media, the file mark position is lost and previously written data is overwritten by the next backup.

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

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

Any additions you make appear in the Devices window.

▼ Deleting a Backup Device

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

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

Using the Auto Media Management Feature

Auto media management further automates the backup procedure by freeing the administrator from the task of premounting and labeling a volume before backup. When Auto media management is enabled in the Devices dialog box, Backup assumes that the backup volume loaded in the device is handled exclusively as Backup media. Backup volumes that appear unlabeled are considered blank and are automatically labeled, mounted, and overwritten with new data.

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

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

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

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

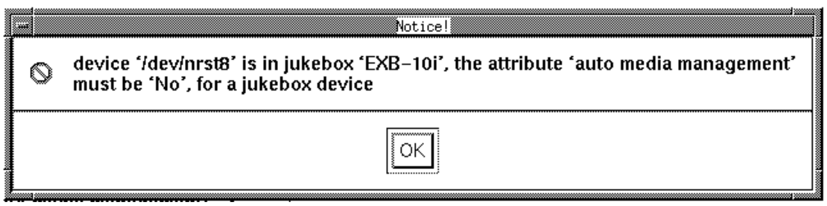

Note – For more information on managing media automatically with jukeboxes see Chapter 8, "Jukebox Software Module."

Recovering and Cloning Save Sets 6

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

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

Note – You can clone entire Backup backup volumes for offsite storage. See "Cloning Backup Volumes" in Chapter 5 for details.

Recovering Save Sets

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

To recover data from the backup media when the browse period has expired and the index entry has been removed, use the Save Set Recover window.

Note – When you recover by save set, you only recover the data; you do not recover file index entries for the data.

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

Note – Recover data whose entries have been removed from the online file index with the save set recover feature in Backup Administrator (`nwadmin` command). Recover data whose index entries are still in the online file index with Backup Recover (`nwrecover` command).

Navigating the Save Set Recover Window

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

To open the Save Set Recover window, select the Recover command from the Save Set pull-down menu.

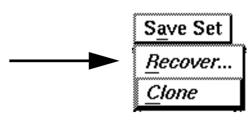

The Save Set Recover window appears, as shown.

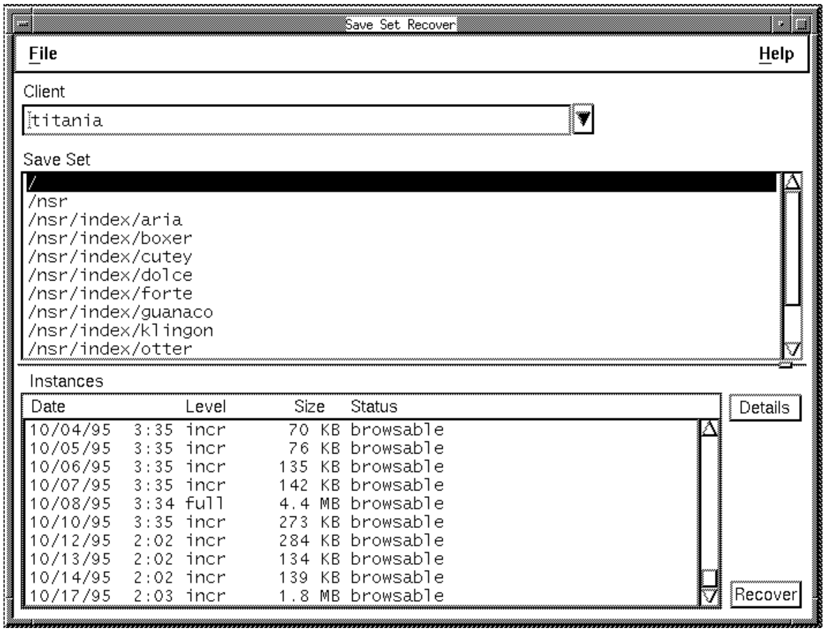

The Save Set Recover window contains the following items.

- Client field – enter the name of the client system for which you want to recover a save set, or use the arrow button to select the client from the Client scrolling list. The list contains the clients known to the current Backup server. If you do not enter a client name, Backup searches for all client save sets.
- Save Set scrolling list – contains the save sets for the selected client. Click one save set at a time to see its versions in the Instances scrolling list.
- Instances scrolling list – contains versions of the selected save set. The list includes the save set Date, Level, Size, and Status. The save set status also includes clone information. If the save set has a clone, it is marked "has clones" and the actual cloned copy is marked "cloned *date time.*"
  To select a save set, click once. To deselect a save set, click again. The following describes the save set status choices:
  - browsable – the save set still has an entry in the online file index.

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

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

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

Save Set Recover Details Dialog Box

The Save Set Recover Details dialog box provides detailed information, including the size of the data you are recovering and the name and location of the backup volume. To open the Save Set Recover Details dialog box, click the Details button in the Save Set Recover window.

The Save Set Recover Details dialog box appears, as shown.

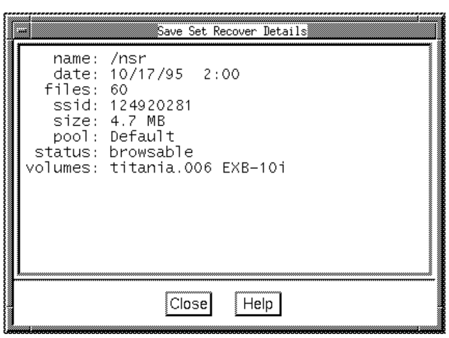

The Save Set Recover Details window contains the following items:

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

Save Set Recover Status Dialog Box

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

The Save Set Recover Status dialog box appears, as shown.

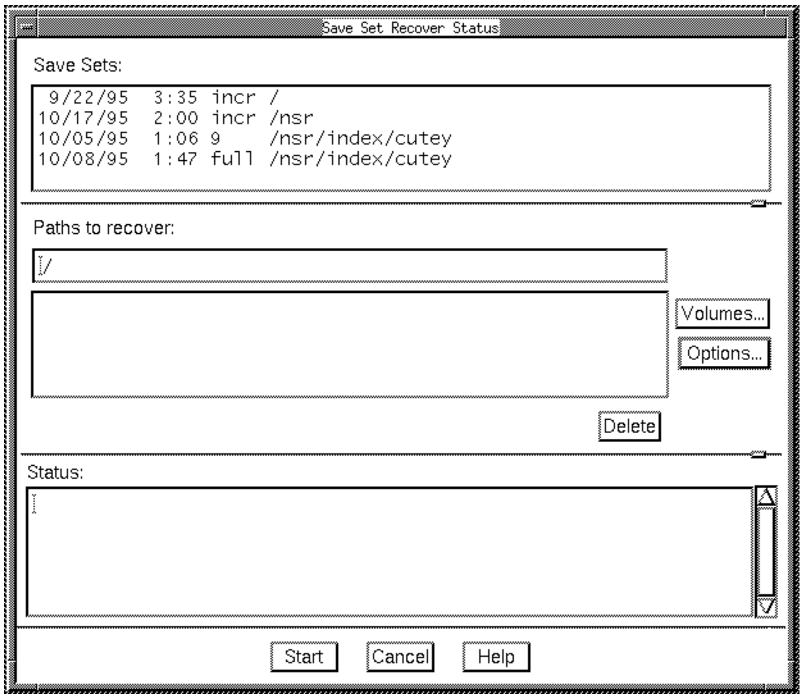

The Save Set Recover Status dialog box contains the following items:

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

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

For example:

The save set *pathname* in the Save Sets scrolling list is `/nsr/index/cutey`. You only need to recover a single file named *help* from the save set. To recover the file *help*, enter the following in the Paths to recover scrolling list:

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

Press [Return] to add *help* to the list. Backup recovers only the file *help*.

For example:

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

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

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

- Delete button – click to delete a file or directory selected from the Paths to recover scrolling list.
- Status scrolling list – displays messages that report on the status of the recover. The Messages display of the Backup Administrator window displays the same information.
- Start button – click to start a recover.
- Volumes button – click to open the Save Set Recover Volumes Required dialog box to determine which backup volumes are needed to recover the data. Backup displays all the volumes required for the save sets listed in the

Save Sets scrolling list. This information may be useful especially if you discover that the backup volume is located offsite and need to make special arrangements to retrieve it.

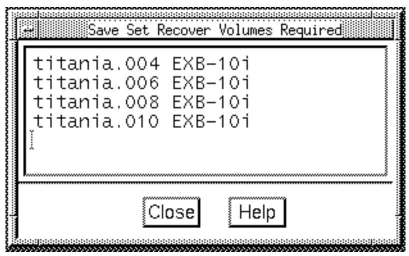

- Options button – click to open the Save Set Recover Options dialog box. Use this dialog box to determine where and how you will recover data.

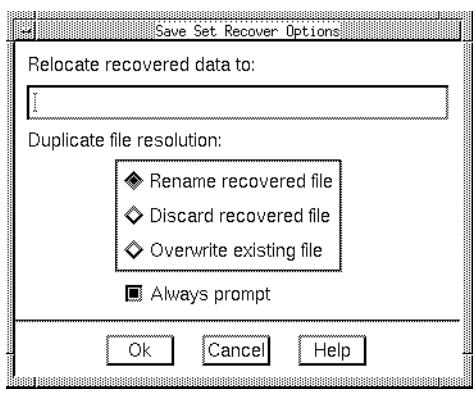

The Save Set Recover Options dialog box contains the following items:

- Relocate recovered data to field – enter the full *pathname* of the directory where you want to locate the recovered data. If you do not enter a *pathname*, Backup automatically recovers the data to its original location.
- Duplicate file resolution choices – select one of these items to determine what Backup should do if a file with the same name already exists. If you do not select one of the choices, Backup automatically selects the Rename recovered file choice.

- Rename recovered file – renames the recovered file with a *.R* extension when a filename conflict occurs. If the *.R* file exists, Backup adds another *.R* extension. For example, a file named *document.R* will be renamed to *document.R.R*.
- Discard recovered file – discards the recovered file. The existing file remains current.
- Overwrite existing file – replaces the existing file with the recovered file. The existing file is lost and the recovered file becomes current.

- Always prompt toggle button – select for prompts when there is a filename conflict, and turn off if you do not want prompts when there is a filename conflict. If the toggle button is not selected, the Duplicate file resolution choice is applied without confirmation.

Examples

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

Example 1 – Recovering Individual Files

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

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

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

1. **Open the Save Set Recover window by selecting Recover from the Save Set pull-down menu.**

2. **Enter the name of the client system in the Client field, or click on the arrow button to select a client from the Client scrolling list.**

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

3. **Select the save set in the Save Sets scrolling list that contains the appropriate filesystem.**
   Once selected, different versions of the save set appear in the Instances scrolling list.

4. **Select the save sets with the correct dates from the Instances scrolling list.**

5. **Click the Details button to see how many files are in the save set.**

6. **Click Ok when finished with the Save Set Recover Details dialog box.**

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

1. **Click the Recover button.**
   The Save Set Recover Status dialog box appears.

2. **Enter the entire pathname of each file to be recovered in the Paths to recover field.**
   The user does not need to recover the entire save set, only a number of files.

3. **Press the [Return] key to add the file to the Paths to recover scrolling list.**

4. **Click the Volumes button to open the Save Set Recover Volumes Needed dialog box to determine the backup volumes needed for recovery.**

5. **Click the Options button to open the Save Set Recover Options dialog box.**

6. **Enter the full pathname of the home directory for the user in the Relocate recovered data to field.**

7. **Select Rename recovered file or Discard recovered file from the Duplicate file resolution choices to prevent any existing files from being overwritten with recovered files.**

8. **Click the Always prompt button to be notified every time Backup finds an existing file with the same name.**

9. **Click Ok to apply the choices and to close the dialog box.**

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

---

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

Example 2 – Recovering a Disk Partition

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

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

The MIS personnel need to do the following to recover the data back to the original location on the disk partition:

1. **Properly prepare the Backup client disk partition where the data was originally located.**
2. **Open the Save Set Recover dialog box by selecting Recover from the Save Set pull-down menu.**
3. **Enter the name of the Backup client where the data was originally located in the Client field.**
4. **Select the save set containing the original project software from the Save Set scrolling list.**
   Different versions of the save set appear in the Instances scrolling list.

Note – Typically a save set is comprised of data from a single disk partition because Backup does not automatically cross over mount points when backing up data.

---

5. **Choose the most recent full backup and any incremental backups performed since the last full backup from the Instances scrolling list.**
   The incrementals are necessary because the engineering team needs any changes to the software since the last full backup.

6. **Click the Details button for one save set at a time to determine its size and number of files it contains.**
   The Save Set Recover Details dialog box appears.

7. **Click Ok to close the dialog box.**

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

1. **Click the Recover button.**
   The Save Set Recover Status dialog box appears. The Save Sets scrolling list contains the save sets selected for recovery.

Note – Do not enter a specific file or directory name in the Paths to recover field because the entire save set is needed for recovery.

---

2. **Click the Volumes button to determine the volume(s) needed for recovery.**
   The Volumes Required for Recovery dialog box appears.

3. **Click Ok to close the dialog box.**

4. **Click the Options button.**
   The Save Set Recover Options dialog box appears.

5. **Enter the full** *pathname* **of the disk partition in the Relocate recovered data to field where the files will be recovered.**
   Leave the field blank if you want to recover the files to their original location.

6. **Select Overwrite existing file from the Duplicate file resolution choices to overwrite any files that still exist on the original disk partition.**

7. **Deselect the Always prompt button.**
   There is no need to be notified every time Backup finds an existing file with the same name.

8. **Click Ok to apply the choices and to close the Save Set Recover Options dialog box.**

9. **Begin the recovery process by clicking the Start button in the Save Set Recovery Status dialog box.**
   As the recovery process begins, messages appear in the Status display that provide information on the progress of the recovery.

Shortcut

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

Follow these shortcut steps to recover a save set:

1. **Choose the Recover command from the Save Set pull-down menu.**

2. **Enter or select the client system whose save sets you want to recover.**

3. **Select the versions of the save sets you want to recover from the Instances scrolling list.**
   You will probably want the last full and any subsequent incrementals to recover the latest version of the files.

4. **Click the Recover button.**
   The Save Set Recover Status dialog box appears.

   a. **If you need to relocate the files, click the Options button, and enter a** *pathname* **in the Relocate recovered data to field, then click Ok.**

   b. **If you want to recover a specific directory or file, enter the complete** *pathname* **in the File to recover field.**
      If you do not enter a *pathname*, the entire save set is recovered.

5. **Click the Start button to begin recovery.**
   Unless you choose to relocate the recovered save set, the data is restored to its original location. You will be notified if any file conflicts occur.

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

Understanding the Save Set Recover Feature

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

Preserving Access Control Lists

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

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

Recovering Nonbrowsable Data

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

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

Note – For more information on browse and retention policies, see “Creating Index Policies” in Chapter 7.

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

Note – The scanner command also recovers data from media whose entries have been removed from the file index, as long as all parts of the save set are intact and the storage media has not been relabeled. With the scanner command you have the choice of scanning in data only, data and index entries, or just index entries. For more information on the scanner command, refer to the man page.

Recovering Large Amounts of Data

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

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

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

Note – For more information on the Backup Recover program, refer to Chapter 4 in the *Solstice Backup 4.2 User Guide*.

6

Cloning Save Sets

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

Note – Besides cloning individual save sets, Backup also clones entire backup volumes. See the section “Cloning Backup Volumes” in Chapter 5.

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

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

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

You must have two backup devices to clone data. Use one backup device for the source volume and use the other for the destination volume. If you have a jukebox with two devices, Backup mounts the volumes required for cloning automatically. If you have stand-alone devices, you need to mount the volumes manually.

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

Cloning save sets requires the following:

- Backup, Server Edition or Backup, Network Edition plus TurboPak
- multiple backup devices
- Active devices value in the Server window set to 2 or more
- source and destination volumes

For information on cloning save sets automatically, see “Automatic Cloning” in this chapter.

Navigating the Save Set Clone Window

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

Use the Save Set Clone window to clone save sets manually. Use the Groups window to enable automatic cloning. (See “Configuring Backup Groups” in Chapter 7 for information.)

Save Set Clone Window

Use the Save Set Clone window for browsing and selecting individual save sets for cloning.

To open the Save Set Clone window select the Clone command from the Save Set pull-down menu.

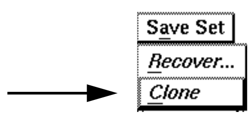

The Save Set Clone window appears, as shown.

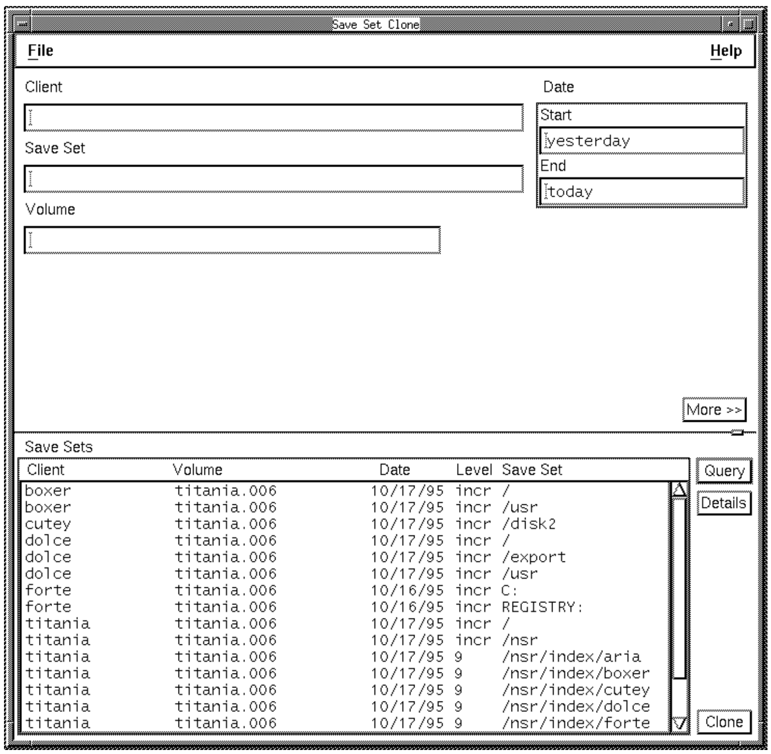

The upper portion of the window contains fields for selecting the save sets you want to browse and select for cloning. You can fine-tune the save sets you select by entering information in as many fields as necessary. The more fields you use for selecting save sets for browsing, the faster Backup locates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

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

The following describes the portions of the Save Set Clone window used for selecting save sets to browse for cloning.

- Client field – enter the client name whose save sets you want to browse. If you do not enter a client name, Backup locates and displays all of the client save sets.

- Save Set field – enter the name of the save set you want to browse.
- Volume field – enter the name of the volume whose save sets you want to browse.
- Date field– enter the range of dates when the save set was created. By default, Backup has yesterday in the Start field and today in the End field. Enter the date when you want to begin the search in the Start field and the date when you want to end the search in the End field. You may enter the dates in a variety of ways:
  - completely written out – October 1, 1995
  - numerically – 10/1/95
  - with a phrase, for example – two weeks ago

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

- More button – click for additional browsing choices in the Save Set Clone window. The window changes by adding more choices for selecting criteria for browsing. See "Additional Browsing Choices" in this chapter for more information.

The following describes the portions of the Save Set Clone window used for selecting and cloning save sets.

- Query button – click to display the save sets chosen for cloning in the Save Sets scrolling list.

Caution – The save sets selected for browsing do not appear in the Save Sets scrolling list until you click the Query button.

- Save Sets scrolling list – displays the save sets currently selected for browsing, including the Client name, Volume name, Date, Level, and Save Set name.
  To select a save set for cloning, click it once. To deselect a save set version, click it again.

- Details button – click to open the Save Set Cloning Details dialog box to view more details about a specific save set before cloning.
- Clone button – click to clone the save sets you selected in the Save Set scrolling list. The Save Set Clone Status dialog box appears. You may clone an unlimited number of save sets in one session.

Save Set Clone Details Dialog Box

The Save Set Clone Details dialog box contains more detailed information about the save set selected in the Save Sets scrolling list. Select only one save set at a time to view its details. The dialog box displays information for the last save set selected.

The Save Set Clone Details dialog box is shown below.

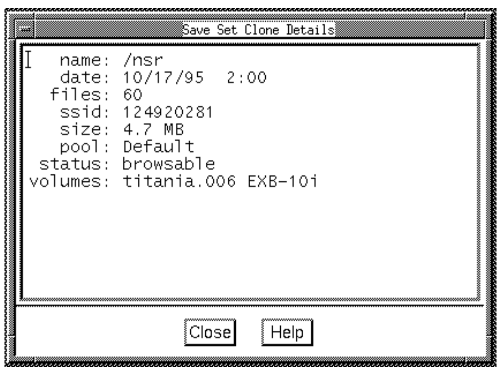

The Save Set Clone Details dialog box contains the following items.

- name – save set name.
- date – day and time the save set was backed up.
- files – number of files in the save set.

Note – The number of files could be zero if the save set was an incremental backup and the files have not changed since the last backup.

- ssid (save set ID) – internal number Backup assigns to the save set for purposes of tracking the data.
- size – size of the save set.
- pool – name of the pool to which the backup volume belongs.
- status – status of the save set, which also indicates if a save set is a clone or has clones.
- volume – backup volume name and location.

The information in the Save Set Clone Details dialog box helps you better prepare for cloning selected save sets.

Save Set Clone Status Dialog Box

The Save Set Clone Status dialog box displays the save set selected for cloning. Use this dialog box to enter the name of the clone pool to which the data will be cloned and to view the progress of the cloning operation.

Caution – Volumes used for cloned data must belong to a clone pool. Use either the preconfigured clone pools, or a clone pool you have created yourself. For more information see “Using Clone Pools” in this chapter.

The Save Set Clone Status dialog box is shown below.

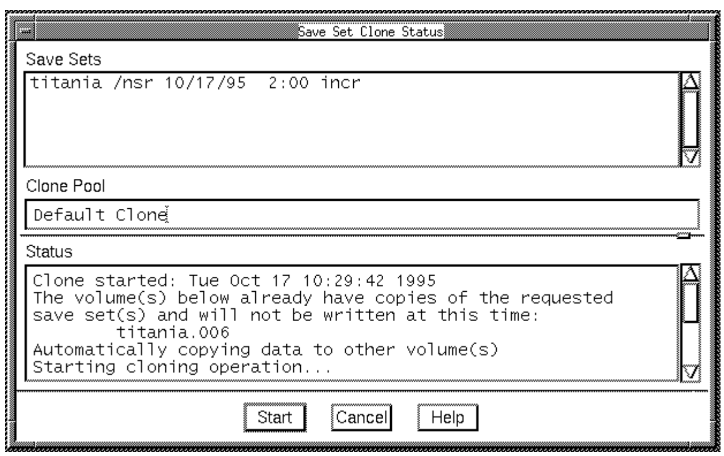

The Save Set Clone Status dialog box contains the following items.

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

Additional Browsing Choices

To view additional choices for browsing save sets to select for cloning, click the More button in the Save Set Clone window.

The Save Set Clone window changes to display additional fields, as shown. To return the window to the previous view, click the Less button.

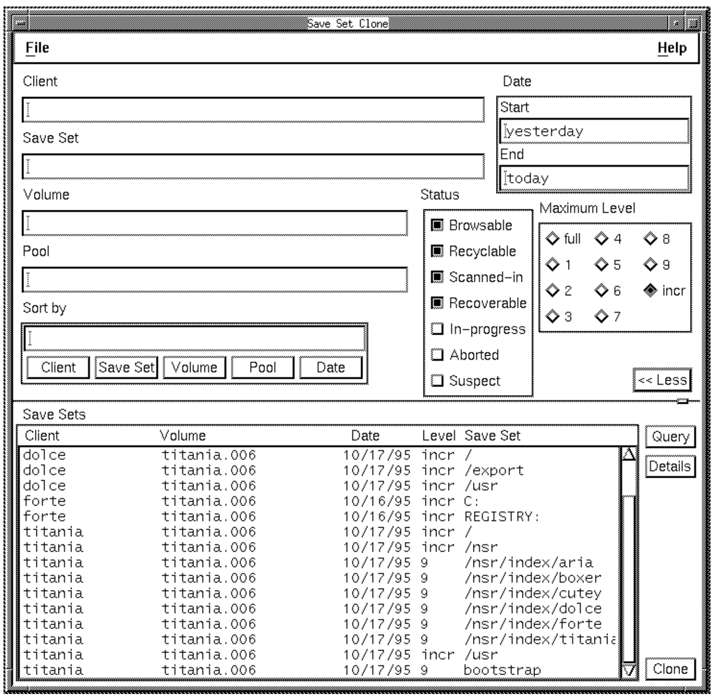

Additional choices for browsing save sets to select for cloning include the following.

- Pool field – enter the name of the pool whose save sets you want to browse.
- Status choices – choose the status of the save sets you want to browse. By default Browsable, Recyclable, Scanned-in, and Recoverable are already chosen. The following describes the save set status choices:
  - Browsable – the save set still has an entry in the online file index.
  - Recyclable – all the save sets have passed both the Browse and Retention policy time periods, and the volume is available for relabeling and overwriting. The data is still recoverable as long as the volume has not been relabeled.

- Scanned-in – the save set was scanned in using the scanner command.
- Recoverable – the entry for the save set has been removed from the client file index, but the save set is still recoverable from the media.
- In-progress – the save set is in the process of being backed up.
- Aborted – the save set was aborted manually by the administrator during backup or by a system crash.
- Suspect – a previous recovery attempt might have failed.

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

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

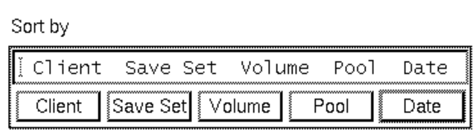

Try entering information in the different fields, clicking buttons, and selecting save sets, to become familiar with the Save Set Clone window.

Examples

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

Example 1 – Manually Cloning Save Sets

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

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

Note – An offsite location will not have the file and media indexes for the cloned data. Therefore, the company location receiving the cloned data must use the scanner command to retrieve the data. See the man page for more information about command line options for the scanner command.

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

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

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

1. **Open the Save Set Clone dialog box by selecting Clone from the Save Set menu.**

2. **Enter the following: client name in the Client field, filesystem name in the Save Set field, and starting and ending dates to search for in the Date field.**

3. **Accept the default Status choices: Browsable, Recoverable, Recyclable, and Scanned-in.**

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

1. **Click the Query button; the Save Sets scrolling list displays all save sets that meet the criteria entered in the upper portion of the dialog box.**

2. **Sort the save sets by date by clicking the Date button in the Sort by field to make sure that the correct dates have been selected.**

3. **Select the save set versions for cloning.**

4. **Click the Details button to determine the size of the data and which backup volume contains the save sets.**

5. **Click the Clone button. The Save Set Status Clone dialog box appears.**

6. **Enter the name of the clone pool to which the backup volumes will belong. Click the Clone button in the Save Set Status Clone dialog box. The Status field displays the progress of the cloning process.**

Using the Save Set Clone window is useful for occasionally cloning individual save sets or for cloning a small number of save sets on an on-going basis.

Example 2 – Automatically Cloning One Client

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

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

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

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

1. **Open the Groups window by choosing Groups from the Customize pull-down menu.**

2. **Create the group *Documents*. See "Configuring Backup Groups" in Chapter 7 for more information.**

3. **Select the appropriate choices in the Groups window. The following are critical choices:**
   - **Autostart choices – Enabled**
   - **Clones choices – Yes**
   - **Clone pool field – Default Clone**

4. **Open the Clients window by selecting Client Setup from the Clients menu.**

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

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

Example 3 – Automatically Cloning Several Clients

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

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

Following is an overview of the tasks the system administrator needs to accomplish.

- Create two schedules: one to perform incremental backups for the clients and one to perform full backups and clones. Create one schedule named *Acct Incr* and another schedule named *Acct Full Clone.*

Note – By creating two schedules and selecting them for the appropriate set of clients, the system administrator avoids cloning all of the incremental backups.

Acct Incr Schedule

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

Acct Full Clone Schedule

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

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

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

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

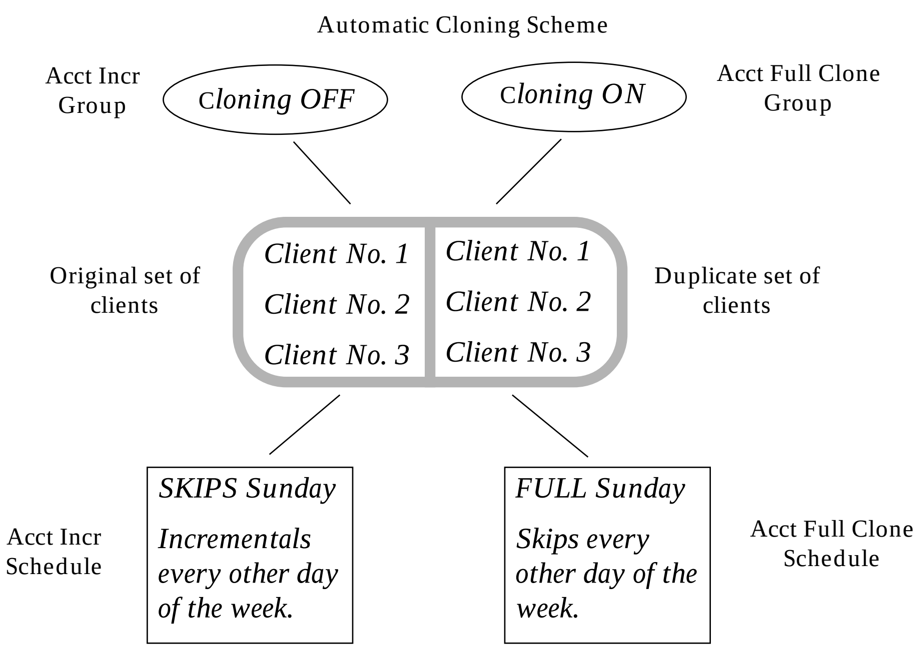

Automatically clone the data by following these steps:

1. **Open the Schedules window by selecting Schedules from the Customize pull-down menu. (For complete instructions on creating schedules, see "Setting up Backup Schedules" in Chapter 7.)**
   The *Acct Incr* schedule performs incremental backups every day of the week, except Sunday when it skips the backup.

Acct Incr Schedule

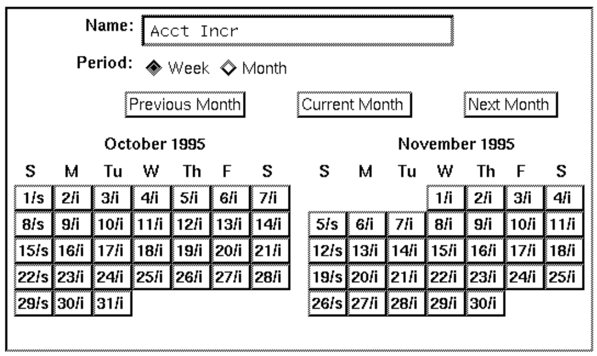

The *Acct Full Clone* schedule skips a backup six days of the week, but performs a full backup and clones the data automatically on Sunday.

Acct Full Clone Schedule

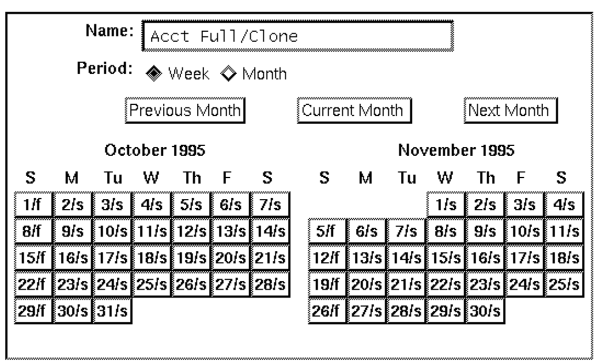

2. **Choose Groups from the Customize menu, after creating the two schedules, to open the Groups window.**

3. **Create the first group, Acct Incr. (For detailed instructions on creating a group see "Creating a New Group" in Chapter 7.)**

4. **Click the buttons for the following critical choices:**
   - Autostart choices – Enabled
   - Clones choices – No

   The window with the *Acct Incr* group looks similar to the following example:

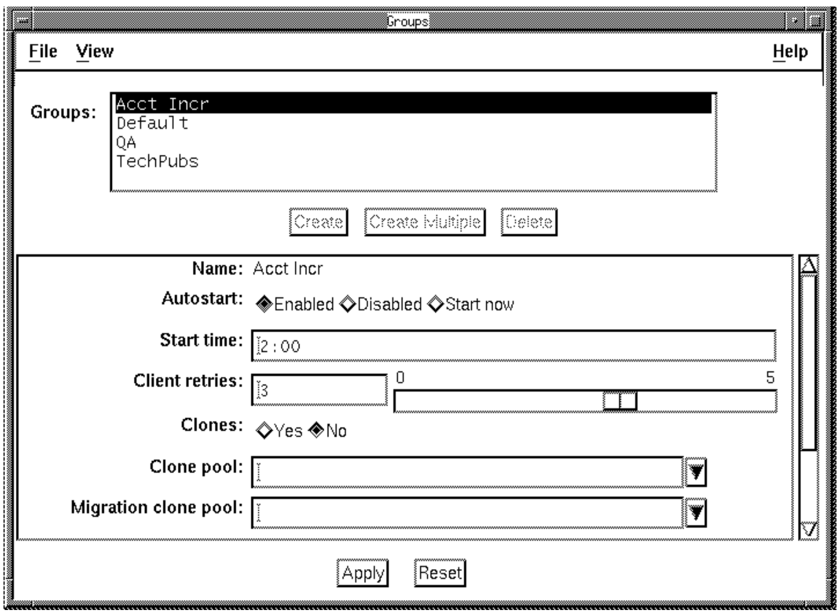

5. **Create the second group,** *Acct Full Clone***.**

6. **Click the buttons for the following critical choices:**
   - Autostart choices – Enabled
   - Clones choices – Yes
   - Clone pool field – Default Clone

Note – Using two different entries for Start time for the backups helps reduce network traffic.

The *Acct Full Clone* group looks similar to the following example:

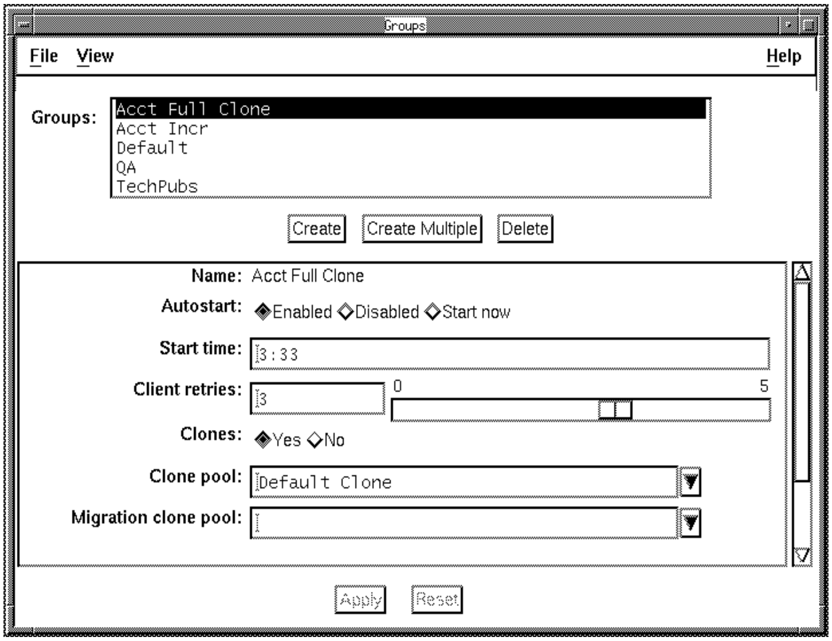

Next, the system administrator creates two sets of identical clients and selects their schedules and groups.

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

7. **Make the other appropriate selections in the Clients window.**
8. **Click the buttons for the following critical choices:**
   - Schedule scrolling list – Acct Full Clone
   - Group choices – Acct Full Clone
9. **Click the Apply button after making the choices for each client.**

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

Shortcut

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

To clone a save set, follow these steps:

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

Understanding the Cloning Save Sets Feature

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

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

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

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

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

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

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

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

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

Automatic Cloning

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

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

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

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

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

Using Clone Pools

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

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

Note – For complete instructions on creating pools, see “Using Volume Pools” in Chapter 5.

Customizing Backups 7

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

Caution – Backup provides preconfigured settings, allowing you to immediately back up your Backup clients. Customizing your backups is not required, but is provided as an option so you can create your own unique configurations. Refer to Appendix A in the *Solstice Backup 4.2 Installation and Maintenance Guide* for information on the preconfigured settings.

The following are tasks for customizing Backup:

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

Setting up Backup Schedules

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

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

Navigating the Schedules Window

Create new backup schedules for your Backup clients in the Schedules window.

To open the Schedules window, select Schedules from the Customize pull-down menu.

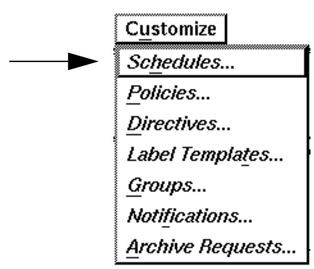

The Schedules window appears, as shown.

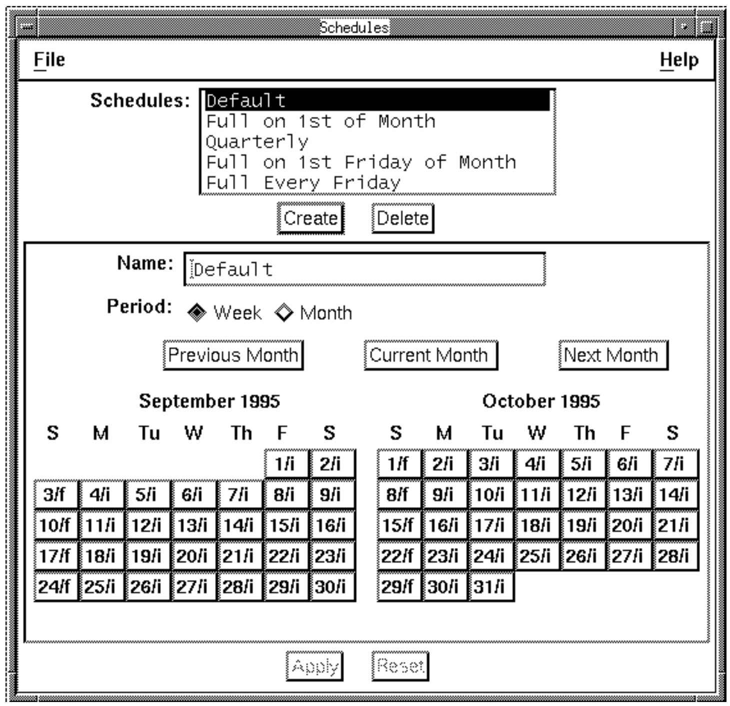

The Schedules window contains the following items:

- Schedules scrolling list – displays the names of the backup schedules known to the Backup server, including the preconfigured schedules provided with Backup and any new schedules that you create.
- Name field – displays the name of the currently-selected schedule. Use the Name field to enter the name of a new schedule after you click the Create button.
- Period choices – let you set the schedule to apply to a weekly or monthly period. Week is the default selection.
  When you select Week and set up a schedule, the backup level you select is applied to that day of the week for all weeks in the calendar (for example, level fulls every Sunday).

  When you select Month and set up a schedule, the backup level you select is applied to that day of the month for all months in the calendar (for example, level fulls on the fifteenth of each month).

- Previous Month button – click to view the backup schedule for the previous month in the calendar.
- Current Month button – click to view the backup schedule for this month and next month in the calendar.
- Next Month button – click to view the backup schedule for the upcoming month in the calendar.

Note – Each time you click the Previous Month button, Backup displays the previous two months in the calendar year; each time you click the Next Month button, Backup displays the next two months in the calendar year.

- calendars – display the day of the month and the backup level scheduled for each day. The date on each calendar day is followed by a forward slash (/) and one of the following numbers or characters indicating the level and type of scheduled backup for that day:
  - f – full backup
  - 1-9 – level one through nine backups
  - i – incremental backup
  - s – skip scheduled backup

  An asterisk (*) next to a backup level means you have forced a different level (an override) for that day. For example, you scheduled a full backup on every Monday of the month, but the second Monday is a holiday. You can force a "skip backup" for the second Monday and force a full backup on Tuesday instead.

  Before setting up a backup schedule, decide which schedule best fits the needs of your network and learn about the backup levels Backup supports.

Examples

This section illustrates setting up backup schedules with two examples. We suggest that you read through the examples before you try to set up any new schedules on your own.

To create a new schedule, name the schedule and select your choices in the Schedules window. Then apply the schedule to a client in the Clients window.

Example 1 – A Weekly Schedule

This schedule is suitable for Backup clients with files that change often. You want to do a full backup every Monday because a weekly full backup simplifies the disaster recovery process. You also want to do a full backup on Mondays because someone will be available to change the media and monitor the server. The other days of the week, incremental backups protect the files that have changed since the last backup. You are going to name the backup schedule “Monday Full” to remind yourself that this schedule does a full backup every week, on Mondays.

To set up the “Monday Full” schedule, follow these steps:

1. **Open the Schedules window by selecting Schedules from the Customize pull-down menu.**
2. **Click the Create button.**
3. **Enter “Monday Full” in the Name field.**
4. **Click the Week button, if it is not already selected.**

5. **Select any Sunday in the calendar for the current month and pull down the menu of backup levels, as shown.**

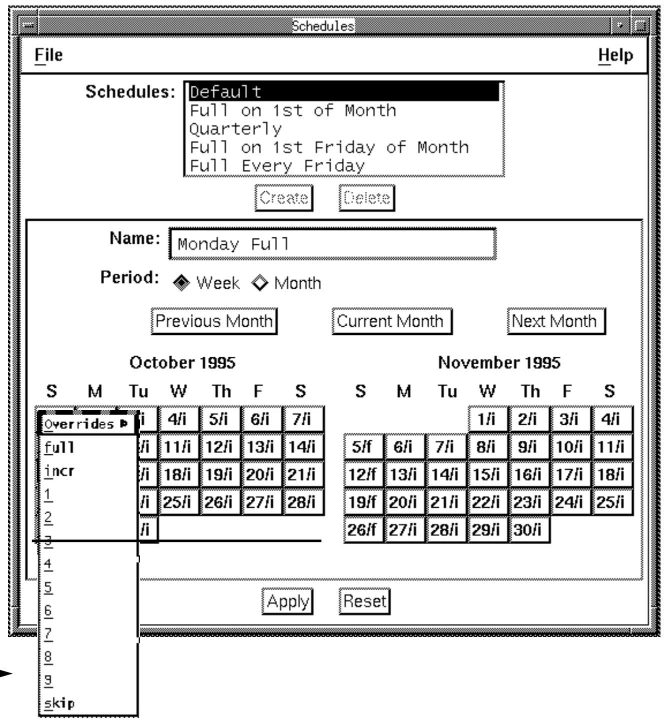

The Backup level menu

6. **Select incr from the backup level menu.**

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

Next, set the calendar to do a full backup every Monday:

1. **Select any Monday in the calendar for the current month and pull down the backup level menu.**
2. **Select full from the backup level menu.**
3. **Click Apply to save the schedule.**

Note – If you see the error message “user *user_name* needs to be on administrator’s list,” it means you do not have permission to make changes to Backup configurations.

See “Adding or Changing Administrators“ in Chapter 4 for more information.

All the Mondays in both calendars display “f” next to the date, indicating a full backup takes place instead of the incremental backup. When you scroll through all the months, you see how the schedule set up for the current month is maintained throughout the entire calendar.

Example 2 – A Monthly Schedule

Use a monthly schedule for clients with files that do not change often. These clients only need a full backup performed once a month.

In this schedule, a full backup takes place on the first day of the month. A level 5 backup takes place in the middle of the month, which backs up all the files that have changed since the full backup. All other days of the month, an incremental backup protects daily changes.

Name this schedule “Monthly Full” to remind you that this schedule does a full backup only once a month. Skip the level 5 backup with an override on November 16 because it is a company holiday; instead force the backup on November 17.

To set up the “Monthly Full” schedule, follow these steps:

1. **Select Schedules from the Customize pull-down menu to open the Schedules window.**
2. **Click the Create button.**
3. **Enter “Monthly Full” in the Name field.**
4. **Click the Month button, if it is not already selected.**

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

Note – When you set up one month in the Schedules window, that setup applies to all calendar months, and carries over from year to year. In the same way, if you set up a weekly schedule, it applies for all calendar weeks. This allows you to plan all the weeks or months in one session, without having to repeat the process every week or month.

Set the month for a level 5 backup on the sixteenth day:

1. **Select the sixteenth day (16/i) of the current month.**
2. **Select "5" from the backup level menu.**

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

Next, override the scheduled level 5 backup on November 16 with a level "skip," to skip the backup for that day, and schedule the level 5 backup for November 17:

1. **Use the Previous Month or Next Month buttons to display the November calendar.**
2. **Select the sixteenth day in the November calendar.**
3. **Pull down the backup level menu, then select Overrides.**

4. Slide the cursor to the right to display the Overrides menu.

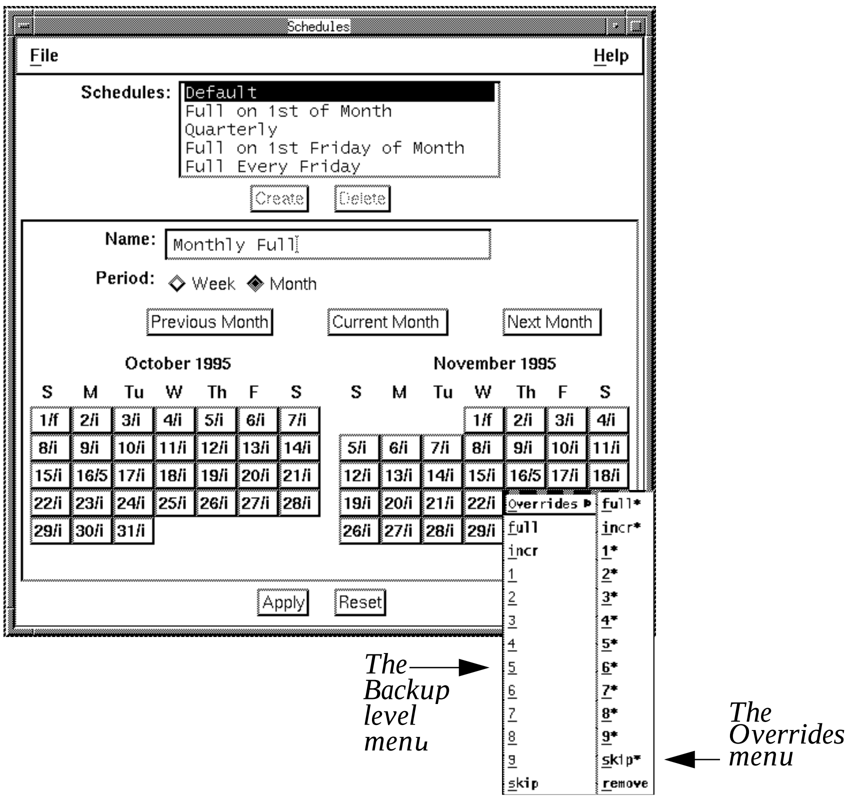

5. Select skip* from the Overrides menu.

6. Select the seventeenth day in the November calendar.

7. Select 5* from the Overrides menu.

8. Click Apply to save the schedule.

The date on November 16 displays an "s*," indicating you have forced a "skip backup" instruction for that day. The date on November 17 displays a "5*" to indicate a level 5 backup will take place instead of the incremental. Overrides must be scheduled explicitly each time. An override does not automatically repeat within a monthly or weekly period.

Shortcut

Use this shortcut section if you are an experienced Backup user, or you have reviewed the examples preceding this section.

To create a backup schedule, follow these steps:

1. **Select Schedules from the Customize pull-down menu to open the Schedules window.**

2. **Click the Create button in the Schedules window.**

3. **Enter a name for the new schedule in the Name field.**

4. **Choose a weekly or monthly period for the schedule.**

5. **Place the cursor on a calendar day, and select it to pull down the backup level menu. Set the backup level for each day (or week, or month) by selecting one of the following choices:**
   - full for a full backup.
   - incr for an incremental backup.
   - skip to skip a backup.
   - any of the numbers [1] through [9] for a level backup.
   - Overrides from the backup level menu to bring up the Overrides menu. Choose an override backup level to apply to the selected calendar date.
   - remove from the Overrides pull-down menu to remove an existing override.

6. **Click the Apply button to create the new schedule when you are through setting backup levels.**

You may set up as many of the Backup software backup schedules as you like, as long as you give each one a unique name. To assign a backup schedule to a client, you must use the Clients window. See Chapter 3, “Configuring Clients,” for more information.

Understanding Backup Schedules

Backup makes it easy to set up your backup schedules. However, deciding which backup schedule best fits your requirements needs planning.

When creating backup schedules, consider the following items.

- How long do you want to keep the backed-up data?

- How many versions of the data do you want to maintain?
- How much data do you have to back up?
- How many backup volumes do you want to use?
- How much time do you have to complete a network-wide backup?
- Do you want to use just a few backup volumes to recover from an entire disk crash?

For example:

Suppose the capacity of an 8mm tape cartridge is about 5 gigabytes and the maximum transfer rate is around 500 kilobytes per second. Based on these figures it would take more than 5 hours to back up 10 gigabytes of data. To back up a network with a lot of data, you may use more than one schedule to stagger the full backups over several days.

Think about how many backup volumes you want to keep – this number will depend on how often the data changes and how long you want to keep the files available for recovery.

If you run only incremental backups every night, you will need more backup volumes to fully recover from a disk crash. If a site has 10 gigabytes of data and 5% of all the data is modified each day, that means you need to back up 500 megabytes of data every day. At 400 kilobytes per second, 500 megabytes takes about 25 minutes to back up and fills about one tenth of an 8mm tape cartridge. If you hold on to backups for three months, you will need to maintain 12 8mm tapes.

You must also determine a policy for recovering files. For example, if the users expect to recover *any version* of a lost file for at least three months, maintain *all* the backup volumes for a three-month period. On the other hand, if the users only expect to recover the *latest* version of a lost file, use level [1-9] backups to decrease the quantity of backup volumes you need to maintain. For more information on policies see “Creating Index Policies” on page 177.

Selecting the Backup Levels

Backup supports four kinds of backup levels.

- Full – backs up all files, regardless of whether or not they have changed.
- Level [1-9] – backs up files that have changed since the last lower backup level. Each backup level is represented by numbers 1 through 9, where 1 represents the fullest backup and 9 represents the most minimal backup.

- Incremental – backs up files that have changed since the last backup, regardless of the level.
- Skip – skips the scheduled backup. For example, you may want to skip a backup on a holiday if you know that no one will be available to change or add more backup volumes.

If you do not need to maintain every version of a backed-up file online, you may use a backup scheme that includes occasional full backups followed by level [1-9] and incremental backups. Different backup levels allow you to trade off the number of backup volumes and amount of time required to complete a backup with the number of backup volumes and amount of time required to recover from a disk crash.

The following diagram illustrates how backup levels work.

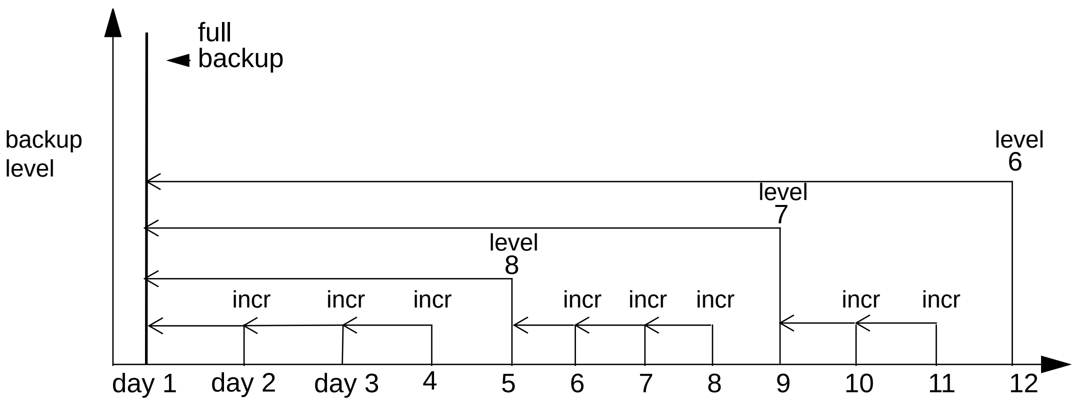

Assume you use a new backup volume for each daily backup.
On day 1, a full backup is run. On day 2, the incremental backs up everything that has changed since the full backup. On day 3, the incremental backs up everything that has changed since day 2. On day 4, the incremental backs up everything that has changed since day 3. At this point, you have 4 backup volumes.
To recover from a disk crash, you need all four of them – the one with the full backup (day1), and all the volumes with incremental backups.

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

On day 9, the level 7 backs up everything that has changed since the full backup. You still 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. Carefully planning your backup strategy allows you to recover everything on a disk with a minimum number of backup volumes.

Note – You also control the size and time it takes to back up your data by using directives, which compress and eliminate unnecessary data from your backups. See the section "Using Directives" on page 184.

Creating Index Policies

This section describes how to create policies that automatically manage the online indexes. You can also manage the indexes manually with the Indexes and Volumes windows. See "Manually Managing the Online Indexes" in Chapter 3 for a description of manual index management and index policy concepts.

When Backup starts a backup, it creates entries for the saved files in the online indexes. Backup maintains two types of indexes: a file index and a media index. The file index contains entries for the individual files which are backed up in groups as save sets. The media index stores entries for the backup volumes that contain the backed up save sets.

The media index does not recognize individual files, but instead recognizes the save sets to which the files belong. Backup maintains one file index per client and one media index per Backup server. Backup uses the indexes as databases to locate files marked for recovery.

As these indexes grow, they take up more disk space. Backup uses browse and retention policies to manage and reduce the size of the online indexes. Use the Policies window to create browse policies for managing the file index entries and retention policies for managing the media index entries. A policy (preconfigured or one you create yourself) can be used as either a browse or a retention policy.

Apply the policies to individual clients in the Clients window by selecting them from the Browse and Retention policies scrolling lists.

Note – Backup provides the following preconfigured policies: Decade, Month, Quarter, Week, and Year. Refer to Appendix A in the *Solstice Backup 4.2 Installation and Maintenance Guide* for more information on preconfigured policies.

Navigating the Policies Window

Use the Policies window to create new policies for your browse and retention policies.

Select Policies from the Customize pull-down menu to open the Policies window.

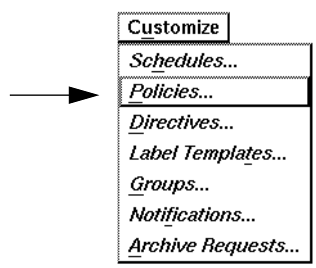

The Policies window appears, as shown.

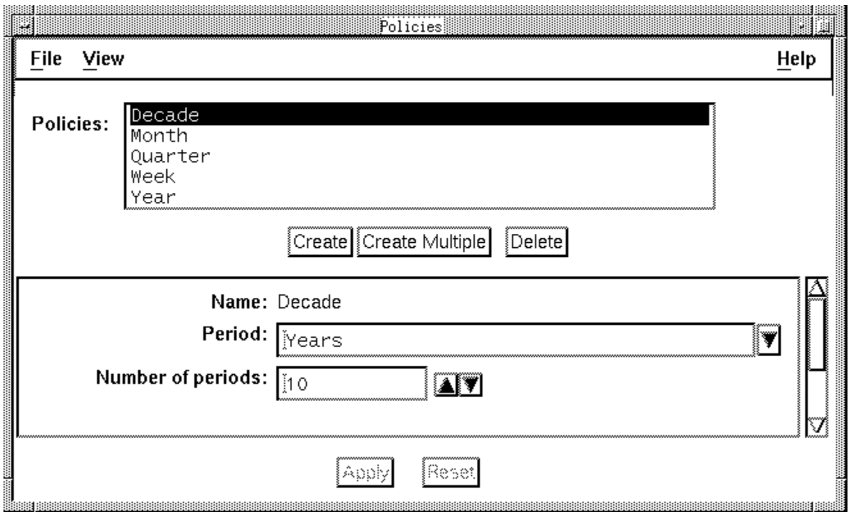

The Policies window contains the following items:

- Policies scrolling list – displays the names of all the policies known to the Backup server. When a policy is highlighted in the scrolling list, its name, period, and number of periods appear in the lower section of the window.
- Name field – displays the name of the currently selected policy. Enter the name of a new policy in the Name field after you click the Create button.
- Period field – displays the period of time you want the policy to apply. Use the arrow button to display the Periods scrolling list choices.
- Number of periods field – displays the number of periods applied to the policy. Use the arrow buttons to select the number of periods you want. For example: A choice of Months for the Period and 3 for the Number of Periods means Backup automatically removes index entries from the online file index that are older than three months.

▼ Creating a New Policy

To create a new policy, follow these steps:

1. Click the Create button.

2. Enter a name for the policy in the Name field.

3. **Select a Period for the policy.**
4. **Select the Number of periods for the policy by clicking the arrow button to select from the scrolling list or by entering a number in the field.**
5. **Click the Apply button to apply your selections. After you click Apply, the newly created policy appears in the Policies window scrolling list and in the Clients window as both a Browse policy and a Retention policy.**
   If you make a mistake, click the Reset button to reset the window to the previous selections.

You can enter any name for the policy that you want. Using the time period for policy names helps you remember the length of the policy. For example, when you specify Quarter as a browse policy for a client, it is easy to tell that the entries are kept in the online file index for three months.

Deleting a Policy

Follow these steps to delete a policy:

1. **Select the policy you want from the scrolling list.**
2. **Click the Delete button. A warning message appears, as shown.**

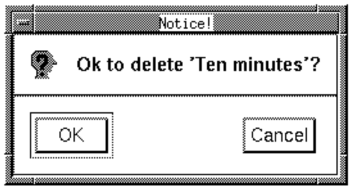

3. **Click OK to delete the policy, or click Cancel if you change your mind.**

Note – You cannot delete the preconfigured Backup policies or policies currently applied to a client.

Changing a Policy

Follow these steps to change a policy:

1. **Select the policy in the scrolling list.**
2. **Select the Period and Number of periods for the policy and make the changes. However, you cannot change the name of the policy.**
3. **Click the Apply button to change the policy.**

Caution – Backup does not allow a client browse policy to exceed its retention policy. This means that an index entry for a save set must be removed from the file index before it can be marked recyclable or removed from the media index.

Understanding Index Policies

Every backup adds entries to the indexes, so the indexes require management in order to control their size and contents. Backup automatically manages the contents of the indexes by using the policies you chose for each client. You can conserve disk space by using automatic index policies to remove old index entries.

Each client requires two separate policies: a browse policy and a retention policy. Backup compares the backup date of the entries in the indexes with the time period specified by the policies. The browse policy determines how long entries for your files remain in the online file index and thus browsable in the Backup Recover window. File index entries older than the browse policy plus one cycle are automatically removed from the online file index, reducing disk space requirements. The retention policy determines how long entries for your volumes are retained in the media index and are available for recovery. Media index entries older than the retention policy plus one cycle are marked as recyclable in the media index, making the backup volume available for relabeling and overwriting with new data. The files on the volumes are deleted when you overwrite the volume with new data.

If you create new policies for automatic index management using the Policies window, the policies appear as selections for both the browse and retention policies in the Clients window. You determine the amount of time files remain browsable and recoverable for a client.

The following diagram illustrates the interaction between an entry in the indexes and the index policies.

Automatic Index Management

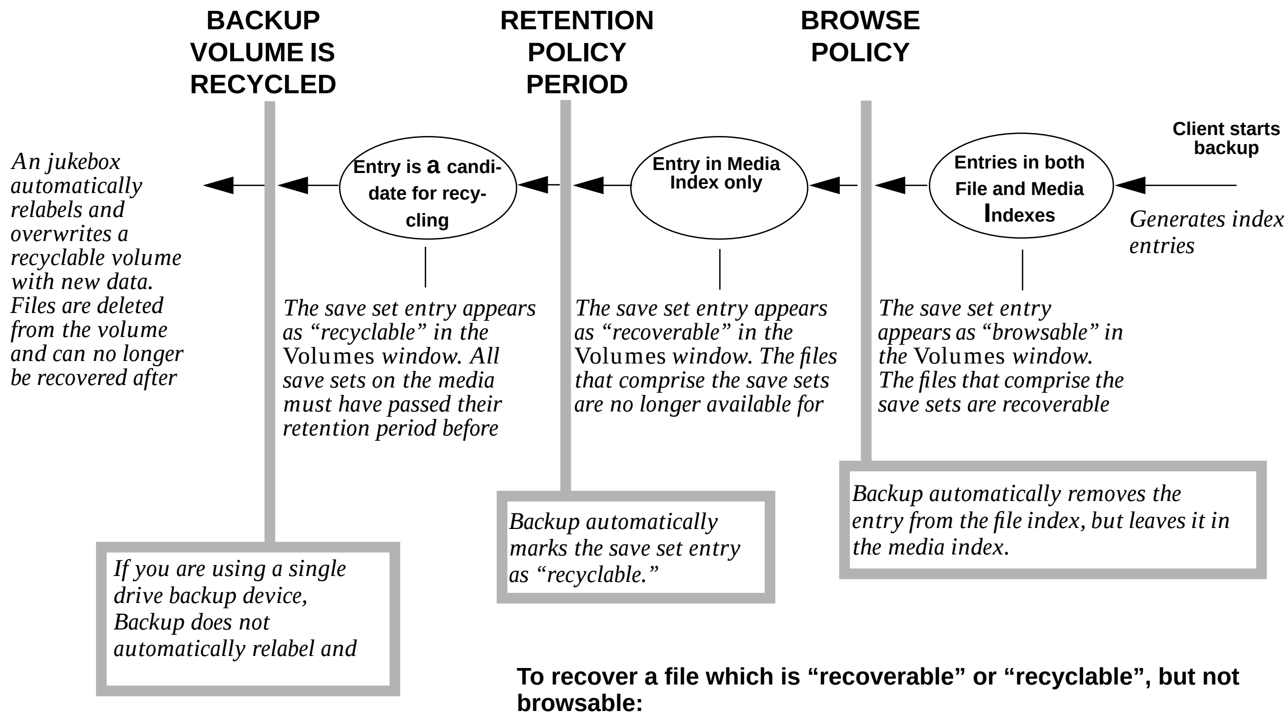

To recover a file which is "recoverable" or "recyclable", but not browsable:

1. *You must know which save set contains the file you want to recover. Use the* Save Set Recover *window to recover the save set containing the file you need.* ***Note:*** *using the* Save Set Recover *window to recover files only recovers the data, not the index entries.*

Using Directives

Directives can reduce the amount of data you back up, possibly even eliminating the need to change backup volumes on the days you perform a full backup.

A directive contains instructions to assist the backup process. Sometimes directives require execution of special programs called Application Specific Modules (ASMs), which are Backup modules that perform specific operations on files. Refer to the uasm man page for more information on ASMs.

Backup is shipped with the following preconfigured directives: DOS standard, Index, NetWare standard, NT standard, NT with compression, Unix standard, and Unix with compression. Each set of directives covers the most important and most useful backup instructions.

Note – For more information about preconfigured directives refer to Appendix A in the *Solstice Backup 4.2 Installation and Maintenance Guide*.

Navigating the Directives Window

Use the Directives window to create new directives or change the preconfigured directives.

To open the Directives window, select Directives from the Customize pull-down menu.

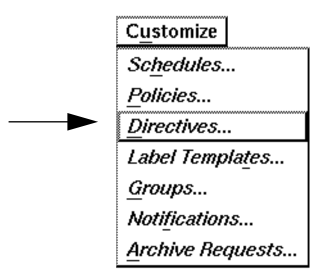

The Directives window appears, as shown.

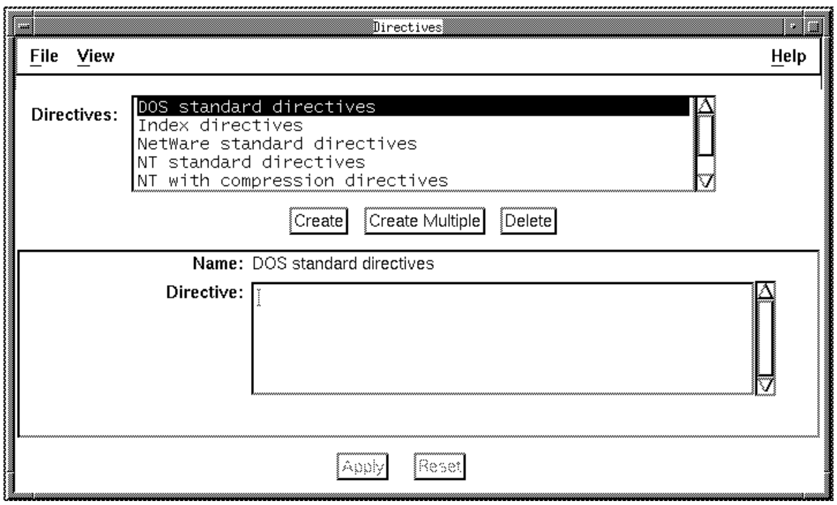

The Directives window contains the following items.

- Directives scrolling list – displays the names of the directives known to the Backup server.
- Name field – displays the name of the currently selected directive. After you click the Create button, enter the name of a new directive in the Name field.
- Directive field – shows the syntax for the currently-selected set of directives. Each directive is preceded by the directory to which it applies (for example, << /tmp >>).

▼ Adding or Changing Directives

Backup backs up your data effectively with the existing preconfigured directives. However, you can create new directives or change existing ones.

To create a set of directives or modify existing instructions, follow these steps:

1. **Click the Create button to add a new directive or select an existing directive from the Directives scrolling list.**

2. **Type your instructions in the Directive field.**

3. **Click the Apply button, after you add or change a directive, to apply the change.**

Note – For more information about directives, including algorithms and the correct syntax used to apply them, refer to the **nsr** man page.

Understanding Directives

Suppose your Backup client is a fast machine, but your network is slow or busy. You may want to *compress* client files in order to use less backup volumes and network bandwidth backing up its files. Using the *compressasm* directive can reduce the space consumed on a backup volume by as much as 50% (actual savings may vary). If you use compressasm on all the files that back up, a full backup of 8 gigabytes may fit onto a single backup volume.

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

Use the *skip* directive to omit certain files from the backup entirely. You can use pattern matching to apply the instructions to specific files in the directory tree. For example, you can have Backup skip backing up *core* files and ".o" files.

Using Label Templates

Backup generates labels for backup volumes according to the rules of a label template. Label templates provide a method for consistently naming and labeling your backup volumes. You can use the preconfigured label templates or create custom label templates using the Label Templates window.

Note – Backup also provides the ability to label your jukebox volumes with bar code labels. See "Understanding Labeling with Bar Codes" in Chapter 8 for more information.

This section describes how to use label templates, introduces preconfigured templates, describes how to create a custom label template, and provides examples for your use.

Backup uses label templates to label and keep track of your backup volumes. All of the backup volumes belong to volume pools that require corresponding label templates. For more information on pools, see the section "Using Volume Pools," in Chapter 5.

It does not matter to Backup 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 Backup needs a backup volume for recovery, it requests it by name in the Pending display of the Backup Administrator window.

Backup provides the following preconfigured label templates: Archive, Archive Clone, Default, Default Clone, Full, Migration, Migration Clone, NonFull, Offsite, PC Archive, PC Archive Clone, and Two Sided. These preconfigured label templates allow you to easily begin labeling your backup volumes without having to create any templates yourself.

Backup also provides preconfigured volume pools with names corresponding to the preconfigured label templates. The preconfigured volume pools automatically use the preconfigured label templates of the same name. The number range for all of the preconfigured label templates starts at "001" and ends with "999" to allow for expansion of the volume pools.

Note – For more information about preconfigured label templates refer to Appendix A in the *Solstice Backup 4.2 Installation and Maintenance Guide.*

Choosing a Label Template

Before you choose a label template, you must first choose the device that contains media ready for labeling from the Devices display in the Backup Administrator window. Next, choose a label template in the Pools window, then start the labeling process by selecting Label from the Media pull-down menu.

For a label template to appear as a choice in the Label window, the associated pool must be enabled in the Pools window.

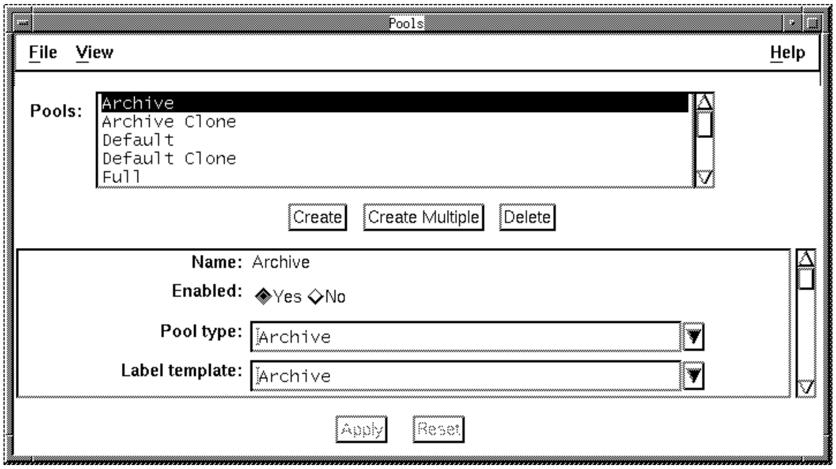

Caution – Whether you use a preconfigured or custom label template, remember to attach an external label to each volume with the correct label template name.

Navigating the Label Templates Window

This section explains how to create your own custom label templates using the Label Templates window.

Open the Label Templates window by selecting Label Templates from the Customize pull-down menu.

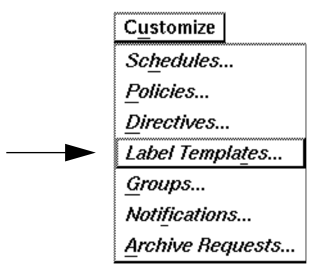

The Label Templates window appears, as shown.

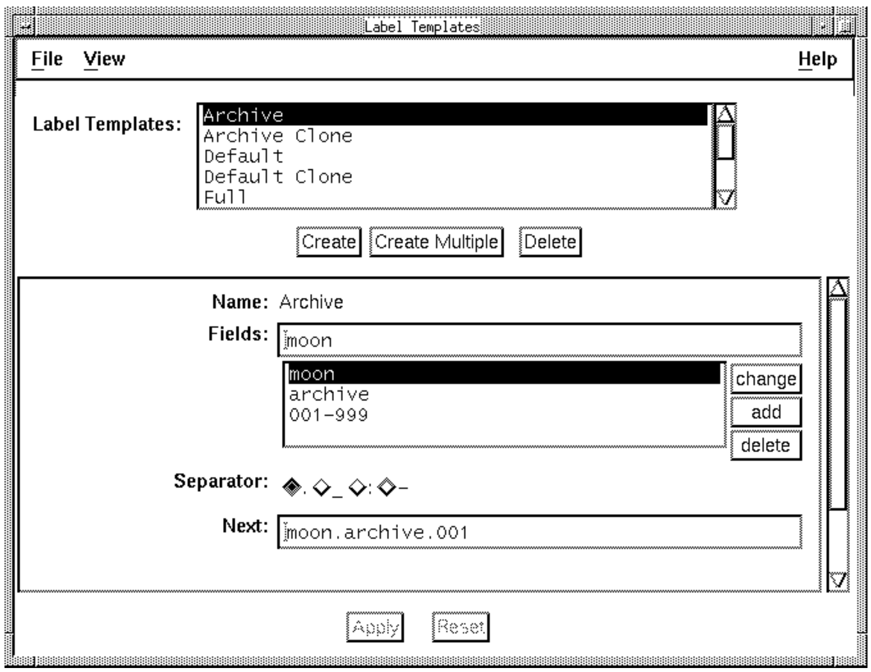

The Label Templates window contains the following items:

- Label Templates scrolling list – displays the label templates known to the Backup server. The preconfigured label templates and any templates you create appear in this list.
- Name field – displays the name of the selected label template. Use the Name field to enter the name of a new label template after you click the Create button.
- Fields field – use for entering the fields that comprise a label template. The scrolling list displays the components of the label template.
  Each component appears as one line in the display. There are four types of components:
  - range of numbers – for example, *001-999*
  - range of lowercase letters – *aa-zz*
  - range of uppercase letters – *AA-ZZ*

• character string – for example, *Accounting, 001*
If you enter a range of numbers or characters, the label number increments with each new label. For example:

first label	Engineering.001
second label	Engineering.002
third label	Engineering.003

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

- Separator choices – use the period, dash, colon, or underscore to separate each component of the label template. If you do not select a separator choice, the label components will not have separators (for example, AA00aa).
- Next field – displays the next label Backup will use to label the backup volume according to the template. If you want to force a label or start the labeling scheme at a certain point, enter the starting label in the Next field. Backup continues to generate labels from that point on, according to the rules of the template. If you want Backup to generate the first name for you, leave the Next field blank.

Examples

Use the examples in this section to help you create a custom label template that identifies your media in a logical and meaningful order.

If your label template matches your backup media storage system, it helps you organize and locate the media in storage. For example, if you store backup media in bins or racks, include a place for the bin or rack number in the label template.

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

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

Example 1 – A Storage Oriented Template

In this example, assume you have only one Backup server and one jukebox. Your backup media is magnetic tape, arranged in numerical order, and organized in three racks. You want to create a scheme that is simple so anyone can easily store and find the backup volumes on the rack shelves.

This example shows three racks, named “1,” “2,” and “3.” Each rack has five shelves; the top shelf is “1” and the bottom shelf is “5.” Each shelf holds 100 tapes, left-to-right.

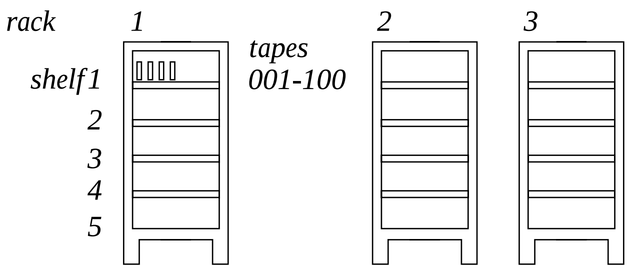

Using this labeling scheme, a tape labeled “2-3-54” will be stored in the second rack, on the third shelf, positioned between the fifty-third and fifty-fifth tapes on the shelf. Backup labels the tapes in the jukebox consecutively according to the label template. Backup starts with the tapes in the first rack, on the first shelf, and numbers the tapes 1-1-001 to 1-1-100. Then, the tapes on the second shelf of the first rack are labeled, numbered from 1-2-001 to 1-2-100, and so forth, until the last tape in the third rack is labeled “3-5-100.” If you want to add a fourth rack to your tape storage system, you can easily do so by changing the label template.

Note – Your label template should allow for expansion of your backup media storage system. For example, it is better to create a template for 100 tapes and not use all of them, than to create a template for only 10 tapes and run out of labels.

Follow these steps to create the storage oriented label template:

1. **Open the Label Templates window by selecting Label Templates from the Customize menu.**
2. **Click the Create button.**
3. **Enter the name of the labeling scheme in the Name field, for example: “Engineering backups.”**
4. **Enter the component of the first field into Fields. This field indicates the rack number, so enter “1-3” (without the quotes).**
5. **Click the add button.**

Note – You must enter the same number of characters for all numbers in the label template.

For example, use:

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

Add the second and third components of the label template:

1. **Enter the second component of the label template into Fields. This component represents the shelf number, so enter “1-5” in the field.**
2. **Click the add button.**
3. **Enter the third component of the label template, the tape position numbers. Enter “001-100” into Fields.**

4. **Click the add button.**
5. **Click the dash (-) Separator button to choose the separator for each component of the label template.**
6. **Click the Apply button to apply the new template to the Backup server.**

The Label Templates window now looks similar to the following.

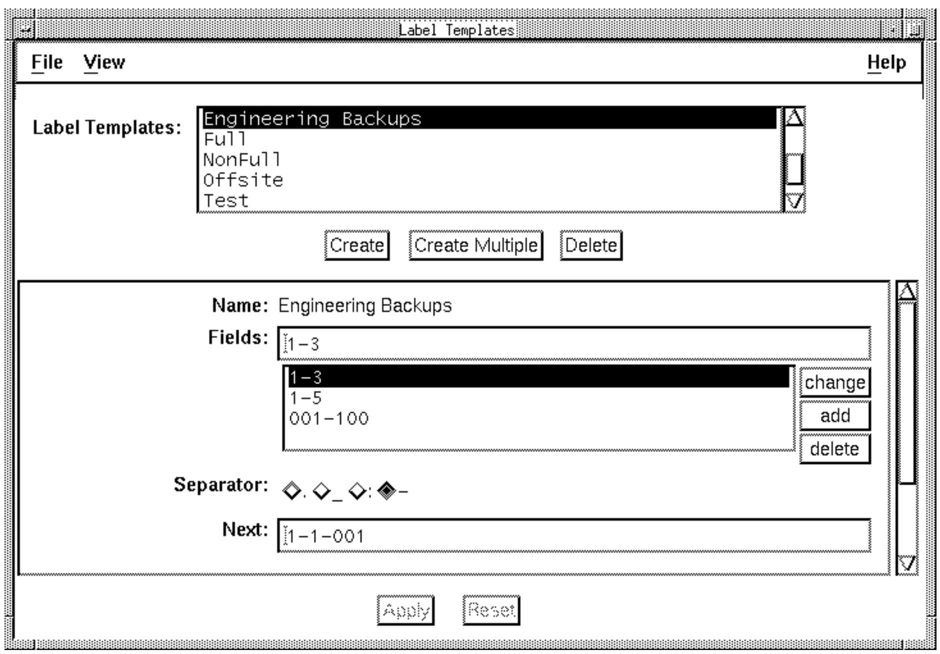

Your new template appears in the Label Templates scrolling list, the components of the label template appear in the Fields scrolling list, and the Next field displays the label of the next tape.

Example 2 – A Sequential Label Template

This example uses a sequential labeling scheme. In this scheme, there is no storage system pattern to follow – the labels are simply a way for Backup to identify the backup volumes. This label template will generate lots of labels – you will probably never “run out” of labels for your backup volumes.

The label template in this example has two fields: “aa-zz” and “00-99.” It will generate 67,600 different labels (26^2 times 10^2). To give you an idea of how Backup generates 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 a sequential label template:

1. **Open the Label Templates window by selecting Label Templates from the Customize menu.**
2. **Click the Create button.**
3. **Enter the name of the labeling scheme into the Name field, for example, “Personnel backups.”**
4. **Enter the component of the first field into Fields. In this example, this field contains pairs of character strings, so enter “aa-zz” into the field.**
5. **Click the add button.**

Add the second component of the label template:

1. **Enter the second component of the label template in Fields. This field contains double digits, so enter “00-99” in the field.**
2. **Click the add button.**
3. **Click the period (.) Separator button to choose the separator for each component of the label template.**
4. **Click the Apply button to apply the new template to the Backup server.**

This template is useful if you do not want to attach any meaning to the labels, but want to be able to generate lots of them over the years and not worry about having to come up with a new template.

Example 3 – A Two-sided Media Label Template

If you have an optical backup device, you can use the Two Sided preconfigured template or create a custom label template that accommodates two-sided media. The label template for two-sided media has *a* and *b* fields to designate the two sides of the optical disk.

Follow these steps to create a label template for two-sided media.

1. **Open the Label Templates window by selecting Label Templates from the Customize menu.**
2. **Click the Create button.**
3. **Enter a name for the label template in the Name field (for example, "TwoSidedLabel").**
4. **Add the following components to the Fields scrolling list: "server_name," "001-999," "a-b."**

Caution – Use "a-b" to distinguish between the two sides of the media.

---

5. **Click the period (.) Separator button to choose the separator for each component of the label template.**
   The Label Templates window now looks similar to the following.

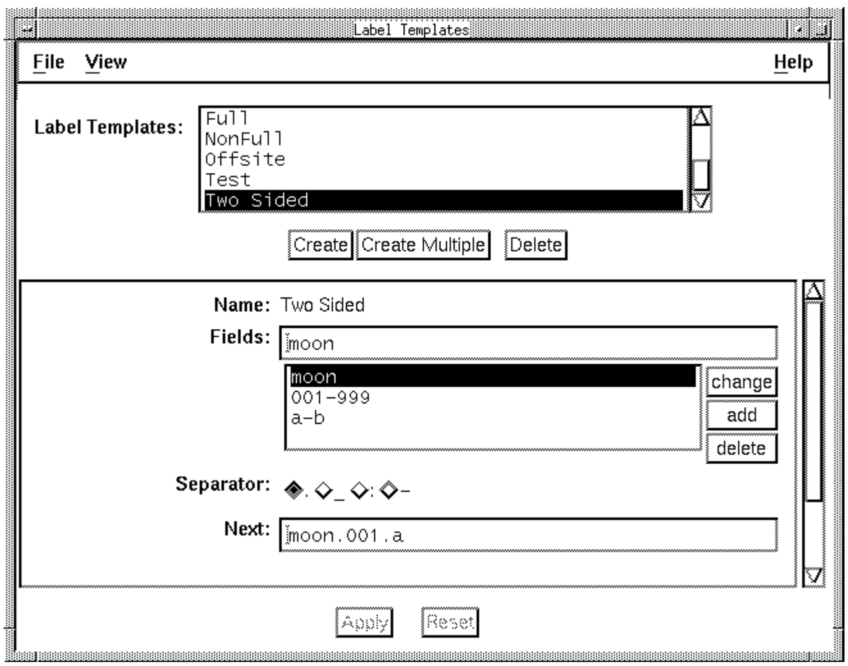

6. **Click Apply to save and apply the label template.**

Using the example shown, Backup labels the first side of the optical disk as *moon.001.a* and the second side as *moon.001.b.*

Caution – When you create a label template, the order in which you enter each component of the template is important. Backup applies each component in a left-to-right order, starting with the first one you enter. Backup displays the components from top to bottom in the Fields scrolling list, with the first component at the top.

Configuring Backup Groups

Backup allows you to bundle clients into backup *groups* so you can easily start selected groups of clients backing up at different times. When you create the different groups, be sure to stagger their backup times to avoid overloading the server.

Backup is shipped with the Default group to which Backup automatically adds all of the clients. You must enable the Default group before it will automatically back up at 3:33 a.m. You are permitted to make changes to the Default group, but you may not delete it.

If you have a large number of clients, consider creating several groups with different backup times to help reduce network traffic. For example, start the backup of the engineering department machines at four o'clock in the morning, and all the other clients on the network at midnight. You can have any number of backup groups on your Backup server. The clients in each backup group start their automatic backups according to the start time of the group.

Note – Create groups in the Groups window. Select groups for individual clients in the Clients window.

Navigating the Groups Window

This section explains how to create new groups using the Groups window.

Open the Groups window by selecting Groups from the Customize pull-down menu.

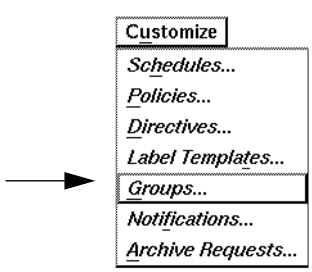

The Groups window appears, as shown.

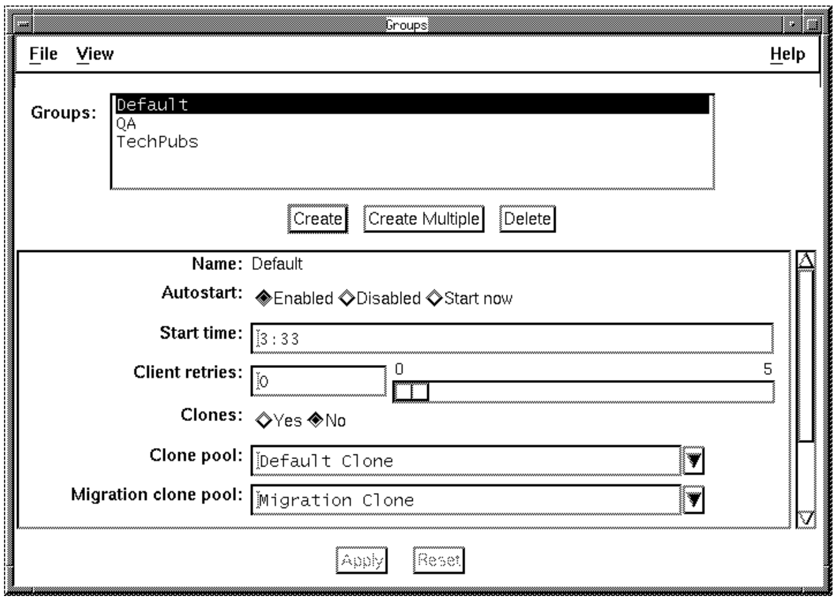

- Groups scrolling list – displays all backup groups known to the Backup server. When you first open the Groups window, you see the preconfigured Default group in the Groups scrolling list. When you select a name in the scrolling list, it also appears in the Name field along with its backup time in the Start time field.
- Name field – displays the name of the group selected in the Groups scrolling list. Use this field to enter the name of a new group after you click the Create button.
- Autostart choices – allow you to disable, enable, or immediately start a network-wide backup. The Default group comes disabled. To allow Backup to back up your clients in the Default group, you must select Enabled from the Autostart choices. The following lists the available Autostart choices:
  - Enabled – allows network-wide backup of the highlighted group to start at the scheduled time
  - Disabled – prevents network-wide backup of a group from starting
  - Start now – immediately starts a network-wide backup of a group, instead of using the scheduled start time
- Start time field – displays the time Backup starts the backup of the clients in the group selected in the Groups scrolling list. The Default group is preconfigured to back up at 3:33 am.

- Client retries field – the number of times Backup will attempt to back up a failed client during a scheduled backup. Use the sliding bar or enter a number in the field. For example, if you enter the number four and a client fails to back up, Backup will attempt to back up that client as many as four times.
- Clones choices – select Yes to automatically clone the data being backed up for a group. Select No if you do not want to clone the data being backed up in a group. In order to clone data you need the Backup TurboPak software to provide support for concurrent devices.
- Clone pool field – enter the name of a pool or use the arrow button to select the clone pool from the scrolling list to which cloned data is backed up.
- Migration clone pool field – enter the name of a pool or use the arrow button to select the migration clone pool from the scrolling list to which cloned migration data is backed up.

▼ Creating a New Group

You may choose to add all of your clients to the Default group. However, that may not be an efficient way to accomplish backing up client data across the network.

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

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

Caution – When you select a start time for your different groups, make sure you schedule them far enough apart so that one group has completed backing up before the next group starts.

5. **Use the Client retries sliding bar to set a value for the number of times Backup will attempt to back up a failed client.**

6. **Select one of the Clones choices. If you click Yes, then select or add the name of the clone pool in the Clone pool field.**

7. **Click the Apply button to apply the changes.**
   The name of the group appears in the Groups scrolling list. If you make a mistake, click the Reset button to clear your changes and start over.

Note – If the error message "user *user_name* needs to be on administrator's list" appears, it means you do not have permission to make configuration changes. See "Adding or Changing Administrators" in Chapter 4 for more information.

After you create a backup group and set the start time, the group appears as a check box in the Clients window. To assign clients to the backup group, click the check box displayed for the group in the Clients window. See the section "Configuring Clients" in Chapter 3 for more information.

▼ Retrying a Failed Client Backup

Backup normally attempts to back up a client in a group once during a scheduled backup. If Backup cannot back up a client in a group (because the client machine is down, for example), it will not try again unless you set a value for Client retries.

Use the Client retries sliding bar in the Clients window to set the number of times you want Backup to keep trying to back up failed clients.

1. **Drag the sliding bar to the right or left to set the number of client retries.**
   The number of times you set for Backup to attempt backing up a failed client appears in the Client retries field.

2. **Click the Apply button.**

▼ Changing a Backup Time

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

1. **Open the Groups window.**

2. Select the group for which you want to change the backup time from the Groups scrolling list.

3. Delete the time in the Start time field and enter a new time, using 24-hour clock convention. (For example, enter 11 p.m. as 23:00.)

4. Click the Apply button to apply the changes.
If you make a mistake, click the Reset button to clear your changes and start over.

▼ Deleting a Backup Group

Backup does not let you delete a backup group while clients still belong to it. Before you try to delete a group, check the Clients window to make sure the group does not have clients. The Default group cannot be deleted.

To delete a group, follow these steps:

1. Open the Groups window.

2. Select the name of the group you want to delete from the Groups scrolling list.

3. Click the Delete button. A dialog box appears to confirm the request to delete the group.

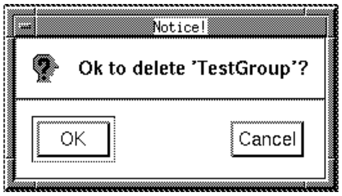

4. **Click Ok to delete the group.**
   If the group does not have clients, the group is deleted and its name disappears from the Groups scrolling list. If the group you are trying to delete still has clients, the following warning appears:

Setting up Event Notification

System administrators need to know about events taking place in the Backup environment. To keep you informed of backup activity, Backup displays messages on the system console (*/dev/console*) or mails them to you electronically.

These notices alert you to important Backup events. If a group of clients did not complete a nightly backup for example, you are informed of this by the savegroup completion notices in your electronic mail.

Backup provides preconfigured notifications for the following events:

- *Bootstrap – prints information needed for recovering from a disk crash, including the server index, media index, and nsr.res file.*
- Cleaning cartridge – reports cleaning cartridge status.
- Device cleaning – reports device cleaning status.
- *Index size* – suggests it is time to reduce the online index.
- *Migration* – reports file migration status. See Chapter 10, "Hierarchical Storage Management Application," for information on using Backup HSM.
- *Registration* – requests that you properly register your Backup products.
- *Savegroup completion* – reports on the success of a scheduled backup.
- *Tape mount request 1*, *2*, and *3* – reports that Backup is requesting mounting of a backup volume.

Note – For more information on the preconfigured notifications refer to Appendix A in the *Solstice Backup 4.2 Installation and Maintenance Guide*.

Navigating the Notifications Window

This section describes the contents of the Notifications window and how to use it for creating new notifications.

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

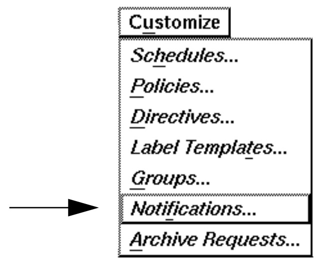

The Notifications window appears, as shown.

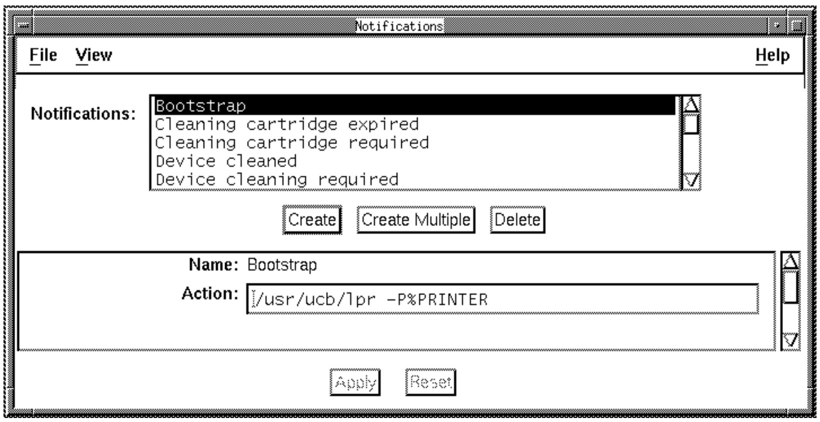

The Notifications window contains the following items.

- Notifications scrolling list – displays the names of all the notifications known to the Backup server. Information about a selected notification is displayed in the lower half of the window.

- Name field – displays the name of the currently selected event notification. Use this field to enter the name of a new notification after clicking the Create button.
- Action field – displays what type of action Backup takes with the notice. For example, Backup can use electronic mail to notify individuals that a volume needs to be mounted in the server. The action can be any UNIX command.
  You can change the Action field to customize how Backup posts the notifications.

▼ Creating a New Notification

To create a new notification, follow these steps:

1. **Open the notifications window by selecting Notifications from the Customize pull-down menu.**
2. **Click the Create button, and enter the name of the new notification in the Name field.**
3. **Enter the code or commands in the Action field.**
4. **Click the Apply button.**

Note – Select Details from the View menu if you want to use the expert mode options for creating a more detailed notification. All event choices are enabled, by default. Click to disable event choices you do not want.

Summary

Backup provides powerful and versatile backup configurations. With the ability to customize Backup's backup configurations, you can create a highly individual and specific method for backing up the data on your network.

Jukebox Software Module 8

This chapter provides information for configuring and operating Backup with your jukebox module. See Chapter 5, “Managing Media and Backup Devices,” for additional information pertinent to all backup devices.

The software displays the term “jukebox” so this chapter also uses the term “jukebox” to refer to an autochanger. Both terms refer to a variety of backup devices, including: carousel, library, near-line storage, datawheel, and autoloader.

Configuring Jukeboxes

Jukeboxes automate the task of loading, mounting, and labeling backup media. Before Backup can back up to a jukebox, you must do the following:

- install and enable Backup or Backup plus TurboPak
- enable the Backup Jukebox Software Module
- load and label your backup volumes

Several jukebox configuration changes occur when you install and enable the jukebox. After installation, you can make changes to the available slots, devices, bar code choices, cleaning devices, and auto media management settings.

Caution – You cannot add or create jukebox resources using the Backup Administrator program. You can only modify jukeboxes previously installed and configured using the **jb_config** program.

If you want to install additional jukeboxes on your Backup server, refer to the*Solstice Backup 4.2 Installation and Maintenance Guide.*

Navigating the Jukeboxes Window

Use the Jukeboxes window to see the names of the jukeboxes Backup recognizes and to remove jukeboxes no longer in use.

The Jukeboxes window displays the configurations selected when you installed the Jukebox Software Module. After installing, enabling, and configuring the jukebox, you can make changes to the following in the Jukeboxes window: Available slots, Devices, Bar code reader, Match bar code labels, and Auto media management.

Open the Jukeboxes window by selecting the Jukeboxes command from the Media pull-down menu, as shown.

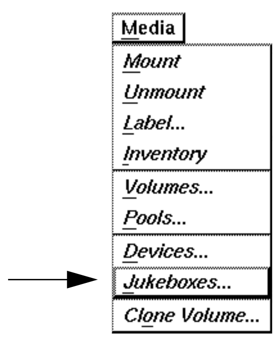

The Jukeboxes window appears, as shown.

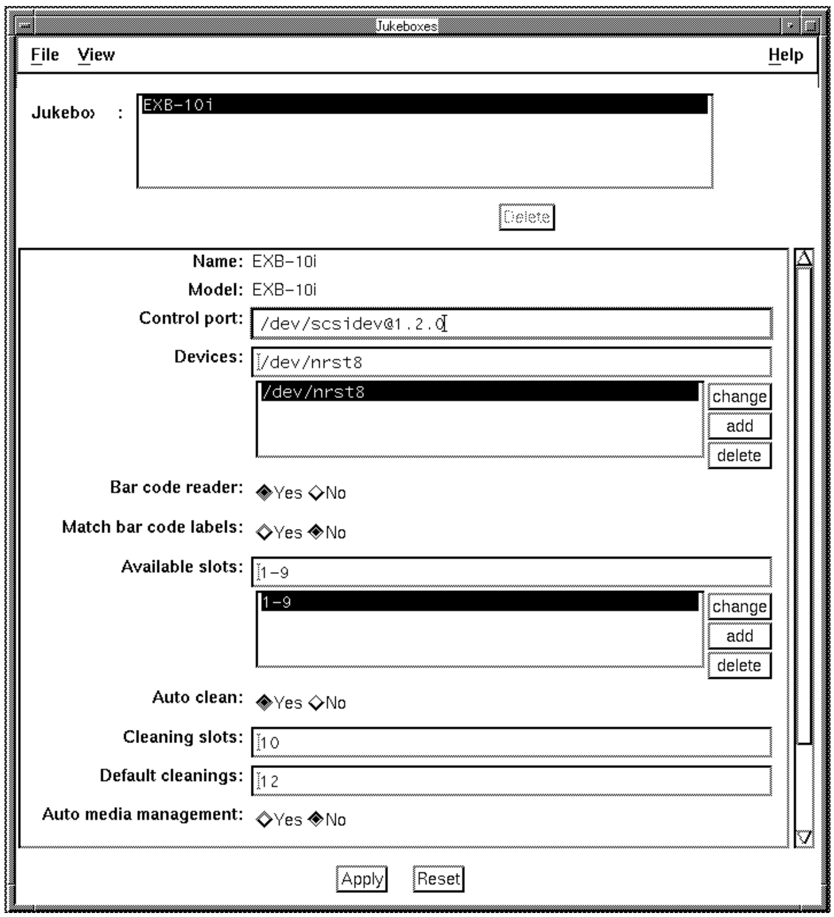

The following jukebox options appear in the Jukeboxes window. You determine all of these settings when you install and enable the Jukebox Software Module.

- Jukeboxes scrolling list – displays the names of the jukeboxes known to the Backup server.
- Name field – displays the name of the jukebox currently selected in the Jukeboxes scrolling list.
- Model choices – displays the model or type of the currently selected jukebox.

- Control port field – displays the system *pathname* of the device that controls the jukebox loading mechanism. The loading mechanism loads and unloads the backup volumes into and out of the device(s) in the jukebox.
- Devices field – displays the system *pathname* of the device that controls the media drive currently selected from the Devices scrolling list. The scrolling list displays the system pathnames for all of the devices located in the jukebox. Devices inside a jukebox, like single-media devices, are configured using the Devices window. Configure the jukebox devices before using the Jukeboxes window to complete the jukebox configuration.
  Use the change, add, and delete buttons to change, add, or delete the highlighted device *pathname* in the Devices field in the Jukeboxes window. See Chapter 5, "Navigating the Devices Window," for more information about the Devices window.

  Devices inside of a jukebox are designated by a (J) next to their *pathname* in the Devices display of the Backup Administrator window, as shown here.

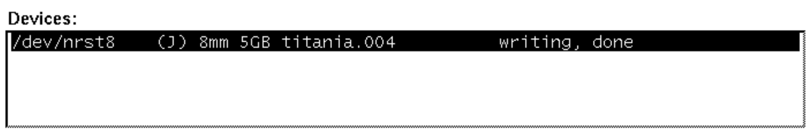

- Bar code reader choices – provides choices for labeling your backup volumes with bar code labels.

Note – For more information on bar code labels see the section "Labeling with Bar Codes" in this chapter.

- Match bar code labels choices – provides choices for matching the bar code label to the volume label.

Caution – If your jukebox device does not have a bar code reader, select No for the Bar code reader and Match bar code labels choices.

- Available slots field – displays the available slot range currently selected in the scrolling list. Backup uses any backup volumes in this range to back up the data from the clients on the network. The Available slots correspond to the numbered slots available for holding backup volumes inside the jukebox. Typically, the slots are identified by a range of numbers. Entering just a single number, for example “10,” restricts access to just slot “10.” Entering the number “10” is not the same as entering “1 through 10.”

- Auto clean choices – use to enable or disable the automatic cleaning function. When Auto clean is enabled and Cleaning interval is defined, Backup automatically cleans the devices in the jukebox and sends a “Device cleaned” notification, indicating the cleaning process is complete. Auto clean is disabled, by default.
  When you disable Auto clean, Backup ignores the value of Cleaning slots and uses the slots for backup volumes. Backup does not automatically clean the devices in the jukebox. If the Cleaning interval field contains a value, Backup sends a “Device cleaning required” notification message, indicating the device needs cleaning. Disable Auto clean to manually load a cleaning cartridge in the jukebox.

- Cleaning slots field – reserves a specified slot or range of slots in the jukebox for the cleaning cartridge(s). Backup automatically selects the correct slots for cleaning when you configure the jukebox.
  A single positive integer reserves the specified slot; two positive integers separated by a hyphen reserve a specified range of slots. If this field is blank or Auto clean is disabled, no slot is reserved for a cleaning cartridge.

- Default cleanings field– determines the number of times Backup uses a new cleaning cartridge mounted in a jukebox. Backup automatically selects the correct number of default cleanings when you first configure the jukebox. When the cartridge has been used the specified number of times, Backup sends a “Cleaning cartridge expired” notification message, indicating the cartridge is used up and needs replacing.

- Auto media management choices – gives Backup exclusive control over the media loaded in a jukebox. If you select Yes, Backup automatically labels, mounts, and writes to volumes it considers unlabeled. If you select No, Backup ignores the unlabeled media and does not consider it for backup.

▼ Adding More Jukebox Devices

To add more devices to a jukebox, you must first configure the Backup server to recognize the added devices using the Devices window. Then you must enter the *pathnames* for the jukebox devices in the Jukeboxes window so Backup will recognize that these devices belong to the jukebox.

Caution – After you add a device *pathname* to the Devices window, you must also add it to the Devices field in the Jukeboxes window. Backup must associate the device in the jukebox with the jukebox name. See Chapter 5, "Managing Media and Backup Devices," for a description of the Devices window.

To add a device to a jukebox, follow these steps:

1. **Select Devices from the Media pull-down menu.**
2. **Click the Create button in the Devices window.**

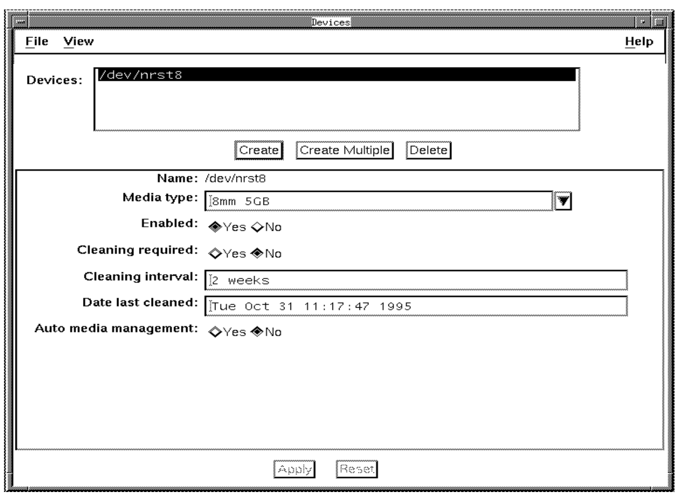

3. **Enter the jukebox device *pathname* in the Name field.**
4. **Select the appropriate Media type choice for the device you are adding to the jukebox.**
   By default, Yes is selected for the Enabled choices.
5. **Enter a value for Cleaning interval. This field accepts: 2days, 2 days, 2weeks, 2 weeks (any combination of numerals followed by day, days, month, or months, with or without a space between).**
   Backup keeps track of the time elapsed since the last cleaning and determines when cleaning is required. A blank value in Cleaning interval disables the automatic cleaning function.
6. **Click the Apply button.**

Next, enter the jukebox devices in the Jukeboxes window:

1. Select Jukeboxes from the Media menu to open the Jukeboxes window.

2. Enter the jukebox device *pathname* in the Devices field.

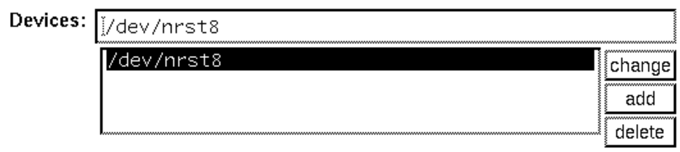

3. Click the add button to the right of the Devices scrolling list.

4. Click the Apply button.

Note – If a jukebox has more than one device, the device *pathnames* must be listed in the same order as their physical location in the jukebox. If you are unsure of their order, refer to the jukebox hardware manual.

For example:

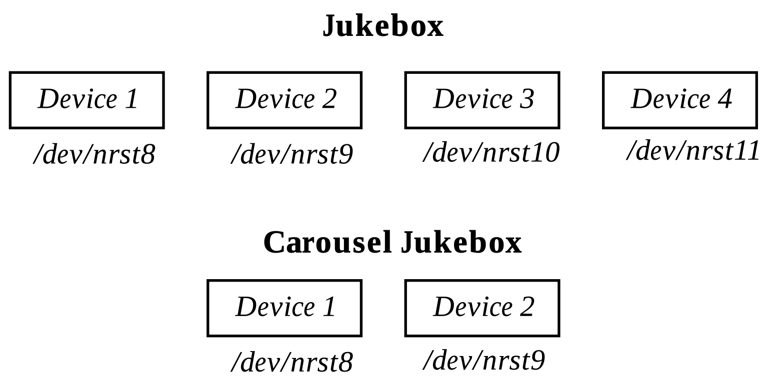

Note – For an alternate method of determining the names of devices within a jukebox, see “Testing the Device Driver Installation” in the *Solstice Backup 4.2 Installation and Maintenance Guide.*

▼ Changing Available Slots

The Available slots feature allows you to control exactly which backup volumes Backup uses for backing up data. Backup uses all of the volumes in the jukebox for recovers. However, you can control which volumes Backup automatically selects for backups by designating a range of available slots in the jukebox.

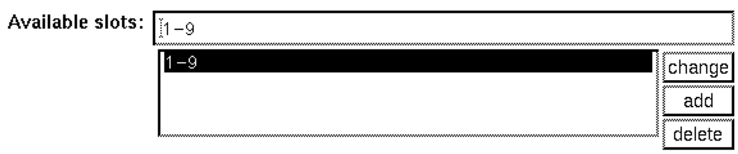

For example, you may designate slots 1 through 5 for your Backup backups in a jukebox that contains 10 slots. The entries can be a range of slot numbers or a single slot number.

1. **Enter the range of available slots into the Available slots field.**
2. **Click the change button.**
   The slot numbers appear in the Available slots scrolling list.
3. **Click the Apply button.**

With two-sided media, the number of available slots is always one-half the number of labels (or sides of the optical disks). For example, if you have 32 optical disks, labeled *atlas.001.a* to *atlas.032.b*, you have a total of 64 labels (and 64 sides). However, the number of available slots is 32.

Caution – Make sure you have placed backup volumes in all the available slots of the jukebox so Backup can proceed uninterrupted with an automatic backup.

Cleaning Jukebox Devices

Maintaining a backup device in good working order requires periodic cleaning. Backup provides automatic cleaning of devices located in a jukebox; Backup does not support automatic cleaning for stand-alone devices.

Note – The *Solstice Backup Installation and Product Notes* contain a list of currently supported jukeboxes for which Backup automatically recognizes cleaning cartridges.

Use both the Jukeboxes and Devices windows to make the necessary selections for automatically cleaning your jukebox devices.

The choices specific to the jukebox appear in the Jukeboxes dialog box where you enable and disable the automatic cleaning feature. Select the cleaning slot(s) for the cartridge(s).

The functions specific to the devices located in the jukebox appear the in the Devices window where you are notified when a device needs cleaning, the date the device was last cleaned, and decide how often a device should be cleaned.

Backup only cleans devices before mounting or after unmounting a backup volume in a device to ensure that there will not be any interference with other jukebox operations.

Cleaning cartridge support provides the following notification messages to keep you informed about cleaning cartridge operations:

Notification Message	Meaning
Device cleaning required	Auto Clean is disabled and the device needs to be cleaned.
Device cleaned	Auto Clean is enabled and the device has been cleaned.
Cleaning cartridge required	Auto Clean is enabled and there are no usable cleaning cartridges available.
Cleaning cartridge expired	The cleaning cartridge has been used the specified number of times.

Note – Check the documentation from your jukebox manufacturer for recommendations on the frequency and cleaning methods for your jukebox devices.

Automatically Managing Media

Auto media management further automates the backup procedure by freeing the administrator from the task of prelabeling volumes before backup. When Auto media management is enabled in the Jukebox dialog box, Backup assumes that the backup volumes loaded in the jukebox may be handled exclusively as Backup media. Backup volumes that appear unlabeled are considered blank and are automatically labeled, mounted, and overwritten with new data. Backup only uses backup volumes that it considers unlabeled if it cannot locate a writable volume or a volume ready for recycling.

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

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

Caution – Be careful about sharing Backup volumes between different backup devices if you have Auto media management enabled. You could potentially relabel and overwrite valuable data.

The Auto media management feature is located in both the Jukebox and Devices windows. For jukeboxes, you only need to enable Auto media management in the Jukeboxes window. Backup does not allow you to enable Auto media management in the Devices window if the device is located in a jukebox.

Note – If you back up to a stand-alone device and want to use the auto media management feature, see Chapter 5, “Managing Media and Backup Devices.”

When Auto media management is enabled, Backup assumes that the jukebox is exclusive to Backup and handles the backup volumes loaded in the device as Backup media. Do not enable this feature if the jukebox or device is a shared resource because Backup assumes it can label and overwrite any media it considers unlabeled. However, if the jukebox is a shared device, you may use the available slot range to restrict the volumes Backup uses for backup.

If you place or store other backup volumes you do not want Backup to use in the jukebox, place them in slots that are not included in the available slot range.

Operating Jukeboxes

This section contains information for performing jukebox operations, using other commands located in the Media pull-down menu. You may need to perform some of these tasks on a daily basis and others less frequently, depending on a variety of factors. For example, the frequency with which you perform these tasks depends upon the number of backup volumes you label at one time, how often you change the contents of your jukebox, and how often you need to manually mount volumes for backup.

Jukebox Mounting Dialog Box

To mount a backup volume in a jukebox, you must first make sure the correct jukebox device is selected in the Backup Administrator window. This is especially important if you have the Backup server configured to use more than one device.

Note – Although jukebox devices are identified by a (J) next to the *pathname* in the Backup Administrator window, it is a good idea to place an adhesive label on the jukebox identifying its device *pathname(s)*. This is to remind you which device *pathname(s)* belong to the jukebox. This practice is especially helpful when you are using more than one jukebox.

Open the Jukebox Mounting dialog box by selecting a jukebox from the Devices display in the Backup Administrator window, then click the Mount speedbar button or select the Mount command from the Media pull-down menu.

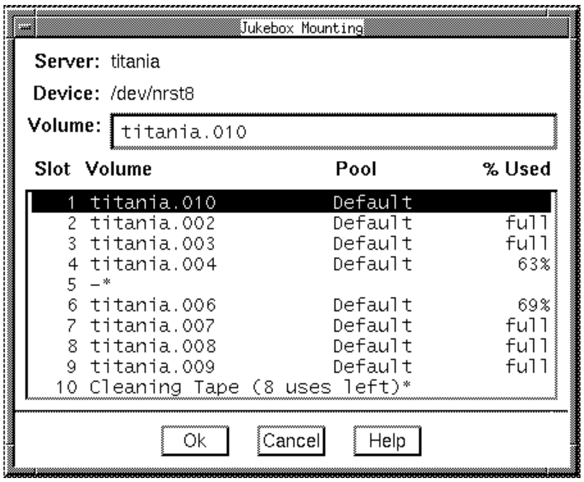

The Jukebox Mounting dialog box displays the name of the Backup server and the system pathname of the selected jukebox device. The Volume field displays the name of the currently selected backup volume. The scrolling list displays the slot numbers (Slot), the label of the volume stored in the slot (Volume), the pool (Pool) to which the volume belongs, and the percentage of the capacity filled on the backup volume (% Used). Use this information to determine which volumes to mount in the jukebox. You may not want to mount full volumes or volumes near capacity. To mount the selected backup volume in the jukebox, first select it from the list, then click the Ok button.

Jukebox Labeling Dialog Box

Use the Jukebox Labeling dialog box to make the necessary choices for labeling the backup volumes in a jukebox. Select a backup device located in a jukebox from the Devices display in the Backup Administrator window. Then, click the Label toolbar button or select Label in the Media menu to open the Jukebox Labeling dialog box.

The Jukebox Labeling dialog box appears, as shown.

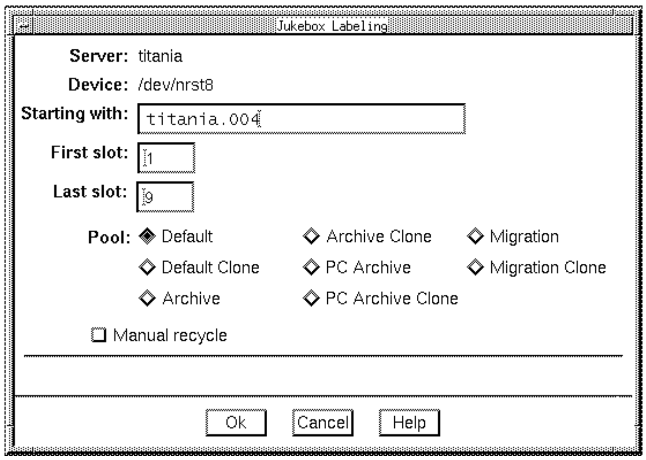

The Jukebox Labeling dialog box displays the name of the Backup server and the system *pathname* of the selected jukebox device. Backup starts labeling the media in the jukebox with the label displayed in the Starting with field.

First slot and Last slot fields display the range of slots containing backup volumes that Backup labels according to the label template selected in the Pools window. Use the First slot and Last slot fields to control the range of slots in the jukebox for labeling or updating backup volumes.

▼ Starting a Labeling Session

Backup uses the label template currently selected in the Pools window to label backup volumes loaded in the jukebox. Choose a label template from the Pools window or create your own before starting a labeling session.

Note – It takes Backup some time to label the volumes in the jukebox. You may want to perform the labeling process at a time when you do not need to back up or recover files.

To label the backup volumes in a jukebox, follow these steps.

1. **Select the jukebox device in the Devices display in the Backup Administrator window.**
   A jukebox device is identified by a (J) next to its *pathname*.

2. **Open the Jukebox Labeling dialog box by selecting the device in the Devices display, and then selecting the Label command from the Media pull-down menu or clicking the Label speedbar button.**

3. **Override the range of slots specified by the label template, if desired, by entering slot numbers into the First slot and Last slot fields in the Jukebox Labeling dialog box.**

4. **Select a pool from the Pool choices for the volumes being labeled in the jukebox.**
   Backup uses the label template associated with the selected pool to label the volumes.

Note – Only the pools enabled in the Pools window appear in the Jukebox Labeling dialog box.

---

5. **Select the Manual recycle toggle button only if you want to manually recycle the backup volume.**
   If you select Manual recycle, any previous browse and retention policies chosen for the volume and its save sets will be overridden and the volume can only be recycled manually. Use this option if you plan on backing up important data to either cloned or noncloned volumes, and you want to decide when the data and the volumes are ready for recycling.

6. **Click the Ok button to begin the labeling process.**
   A confirmation message appears, as shown.

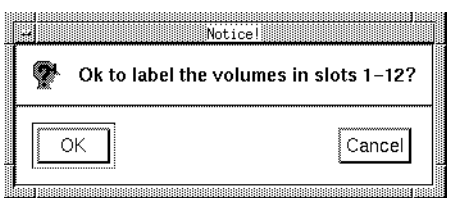

When a valid Backup label already exists on the media you are trying to label, a confirmation message appears. The confirmation is important because it keeps you from accidentally relabeling the media. When a backup volume is relabeled, its contents under the previous label cannot be recovered. When you select OK to confirm, the volumes in the slots are loaded, labeled, and unloaded.

Note – If you label a specific range of backup volumes, the Starting with name must match the label template. If you label a single backup volume, you can use any name – it does not have to match the label template. The First and Last slots must be identical when labeling a single volume.

The following examples illustrate labeling a specific range and a single volume.

First example:

If you previously labeled six tapes and need to label six more, your Jukebox Labeling dialog box might look like this.

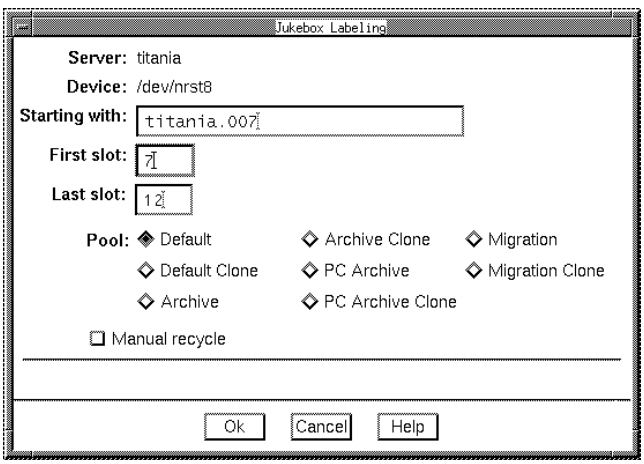

Caution – If you remove backup volumes from the jukebox, remember there may still be one or more volumes left in the jukebox devices. Make sure you unmount them by selecting the jukebox devices and selecting Unmount from the Media pull-down menu in the Backup Administrator window. This prevents you from accidentally relabeling these volumes when you reload the jukebox with new volumes.

Second example:

If you label a single backup volume, the First slot and Last slot must be identical, as shown.

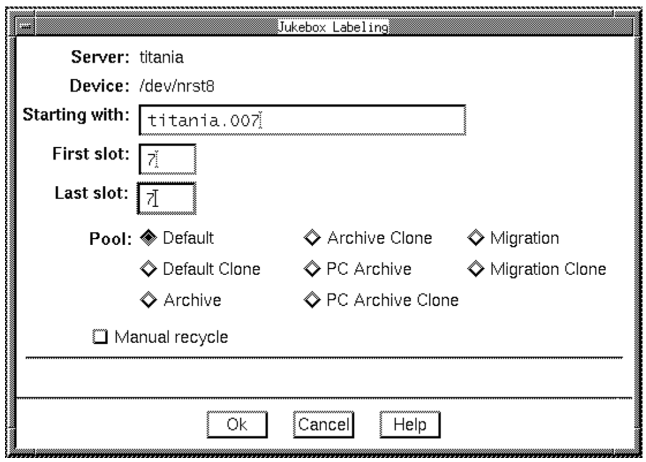

The Jukebox Labeling dialog box gives you control over volume labeling. You can use a label template to generate names as needed or you can specify a single unique name.

Note – For more information on labeling volumes see “Using Label Templates” in Chapter 7.

▼ Labeling with Bar Codes

Labeling media with bar code labels provides two distinct advantages: it significantly speeds up volume inventory and provides improved labeling accuracy.

Inventorying with bar code labels is fast and efficient because you do not have to load the volumes into the device. Instead, the jukebox scans the external bar code labels with an infrared light while the volumes remain in their slots. Inventorying with bar code labels greatly reduces the time it takes to locate a volume or determine the contents of a carousel cartridge.

Bar code labels also provide greater labeling accuracy because the labels are attached to the media prior to being loaded and scanned in the jukebox. Once the jukebox scans the bar code label, Backup records and tracks the label in the media index.

Caution – Backup uses bar code labels for only one reason: to inventory backup volumes. Backup uses the volume label (normally created with a label template) for identifying the volumes required for backup and recovery. However, Backup displays both the bar code label and the volume label in the Pending display of the Backup Administrator window when requesting a volume. You may also open the Volumes window to view both the volume label and its associated bar code label.

A backup volume must have a volume label, but does not require a bar code label. Use the Jukeboxes window and the Jukebox Labeling dialog box to recognize bar code labels on your volumes.

To label your Backup volumes with bar code labels, follow these steps:

1. **Properly apply the bar code labels to your backup volumes.**
2. **Place the volumes with the bar code labels in the jukebox.**
3. **Open the Jukeboxes window by selecting Jukeboxes from the Media pull-down menu.**

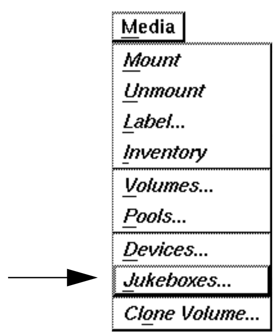

The Jukeboxes window appears, as shown.

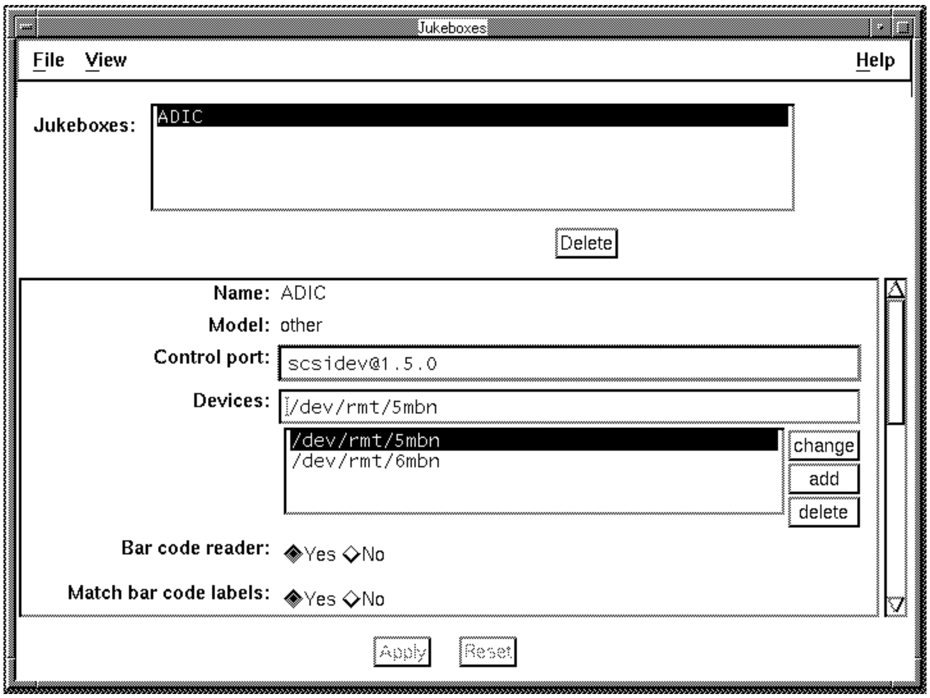

4. **Click the Yes button to enable the Bar code reader.**

5. **Click the Yes button in the Match bar code labels choices if you want the volume label (normally created using a label template) to match the bar code label.**
   If you do not want to override the label created by the label template with the bar code label, click No.

   In either case, you will have two labels: the volume label (which may or may not match the bar code label depending upon your bar code choices in the Jukeboxes window) and the bar code label.

6. **Open the Jukebox Labeling dialog box by selecting Label from the Media pull-down menu.**
   If you chose Yes for the Match bar code labels choices in the Jukeboxes window, the Starting with field in the Jukebox Labeling dialog box displays Bar code label.

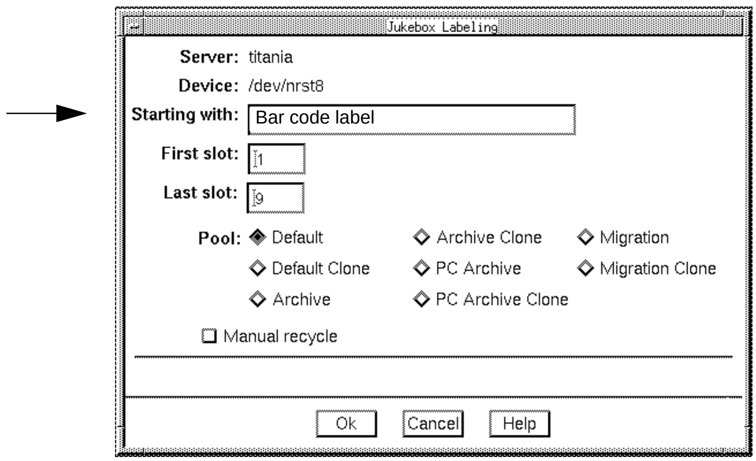

If you chose No for the Match bar code labels choices in the Jukeboxes window, the Starting with field in the Jukebox Labeling dialog box displays a label template name.

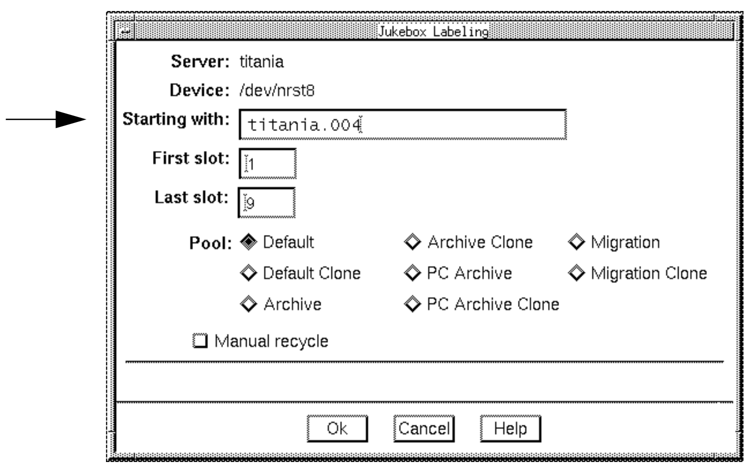

Once you select your bar code label choices and label the backup volume, both label names for the volume appear in the Volumes window.

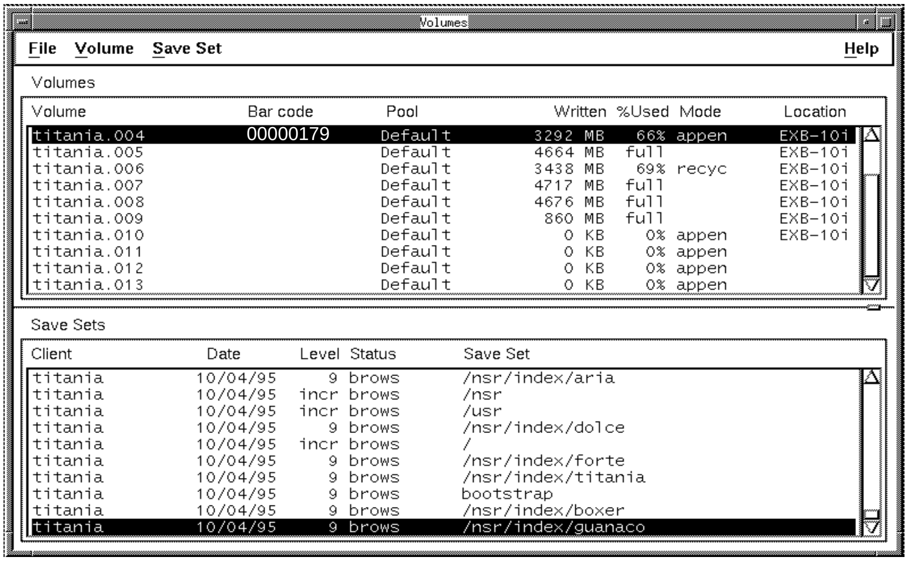

▼ Understanding Labeling with Bar Codes

The volume label and the bar code label are both pieces of information that Backup uses to identify backup volumes. Both labels are recorded in the media index. However, the volume label is also recorded internally on the media. Backup uses bar code labels for inventorying volumes and uses the volume labels to identify the volumes needed for backup and recovery. You can require the volume and bar code label to match, by making the appropriate selections for bar code labeling in the Jukeboxes window.

Caution – Do not use identical bar code labels for any of your Backup backup volumes. Using identical labels defeats the purpose of using bar code labels, which is to speed up the inventory process and ensure labeling accuracy.

If you try to label a second volume with an identical bar code label, and you have enabled Matching bar code labels in the Jukeboxes window, Backup displays an error message and does not allow you to label the second volume. The error message identifies the jukebox slots containing the two volumes with identical labels and the bar code number. To correct the problem, you must either apply a different label and begin the labeling process again, or you must disable Match bar code labels in the Jukeboxes window while labeling the second volume.

Note – If you choose not to match the volume label to the bar code label, you should create and attach volume labels to the outside of your media.

If you choose to have the labels not match, we suggest you label the volumes in this manner:

1. **Attach all of the bar code labels to the media, then load the volumes in the jukebox.**
2. **Begin the labeling process by clicking Yes for the Bar code reader choice, and clicking No for the Match bar code labels choice in the Jukeboxes window.**
   Backup uses the next available label from the label template for the volume name. Backup labels the volumes and records both labels in the media index.
3. **Open the Volumes window, after Backup completes the labeling process, to determine the volume label and bar code label for each volume.**
   Create your own volume labels to attach to the backup volumes.

   Use the information in the Volumes window to match the correct volume labels to the bar code labels. This is the easiest way to make sure that you attach the correct volume labels to the volumes with bar code labels.

Note – If you click Yes for the bar code choices and forget to attach a bar code label, you receive an error message informing you that there is no bar code label for that volume.

It is not necessary to label existing backup volumes with bar code labels if they are stored in a vault or offsite for long periods at a time. This is because you will not inventory these volumes often, if ever. However, if you have volumes you use often for recovery or for overwriting with new data, it is beneficial to label them with bar code labels. Labeling these volumes with bar codes saves you hours of time inventorying your backup volumes.

If you decide to use bar code labels on your existing backup volumes, you must first apply the bar code labels to the existing volumes. Then, load and mount each volume individually so Backup can match up the bar code label with the existing volume label.

Note – You may purchase a variety of bar code labels from a third-party vendor. You can choose numeric, alphanumeric, or a special combination of numbers and characters to meet your labeling needs. You can even order bar code labels that match your current volume labeling scheme.

If you label your volumes with the server name and an extension such as "001," order a range of labels starting with "*server_name*.001" and ending with "*server_nam*e.100." Labeling instructions for bar code labels are usually provided with your jukebox hardware documentation. If you have questions about bar code labels, contact the hardware manufacturer.

Using a consistent labeling scheme helps you better organize and track your backup volumes. It also aids the inventory process if all the volumes use bar code labels rather than just a limited number of them.

When Backup relabels volumes automatically, it reuses the original volume label name. You can only change the label name if you relabel the volumes manually. Backup scans the bar code label during the labeling process and updates the media index with the new volume name and its associated bar code label.

Jukebox Inventory Dialog Box

Use the Jukebox Inventory dialog box to enter the range of slots you want Backup to inventory in your jukebox. Select Inventory from the Media pull-down menu to open the Jukebox Inventory dialog box.

The Jukebox Inventory dialog box appears, as shown.

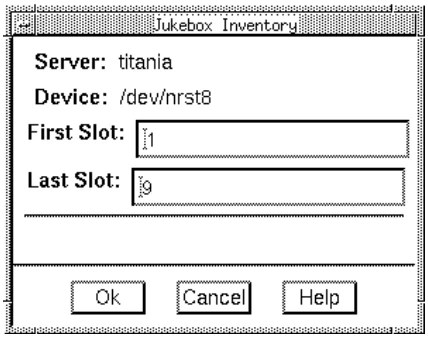

▼ Inventorying Backup Volumes

To inventory backup volumes in a jukebox when volumes have been moved or added, follow these steps:

1. **Load the new backup volumes into the jukebox.**
2. **Select the jukebox device in the Devices panel of the Backup Administrator window. Jukebox devices are marked with a (J).**
3. **Open the Jukebox Inventory dialog box by selecting Inventory from the Media pull-down menu.**
4. **Enter the range of slots you want to inventory in the First and Last Slot fields.**
5. **Click Ok to begin the inventory.**
   The Jukebox Inventory dialog box displays a message showing the progress of the inventory operation. Since Backup mounts each backup volume, reads its label, and unmounts it, the inventory process may take some time to complete.

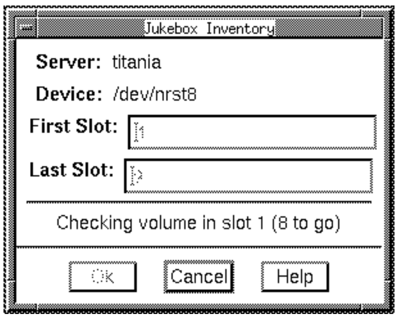

You may cancel the inventory operation after Backup completes inventorying the current volume by clicking Cancel in the Jukebox Inventory dialog box. Click Cancel to close the window when the inventory is complete.

After completing an inventory, Backup registers the contents of the jukebox and then proceeds with its network-wide backup and recover services.

▼ Disabling the Element Status Feature

Your jukebox may support the element status feature. This feature, along with the bar code label feature, speeds up the inventory process by eliminating the need to load and read volume labels.

Backup uses the element status feature to determine whether or not a slot contains a volume. If the volume has a bar code label, Backup reads it without loading the volume.

If you perform an inventory on a jukebox containing volumes without bar code labels, Backup must load the volume to read the internal label. In this case, the element status feature is not beneficial.

Note – Check the documentation from your jukebox manufacturer for information on features for your jukebox.

If you are not using bar code labels and want to disable the element status option for your jukebox, follow these steps:

1. **Select Jukeboxes from the Media pull-down menu in the Backup Administrator window.**

2. **Select Details from the View pull-down menu in the Jukeboxes window.**
3. **Use the scroll bar to view the Jukebox options field.**
4. **Enter the following option in the Jukebox options field, then click add to add it to the Jukebox options list and disable the element status option:**

```
!element_status
```

The option is added to the Jukebox options list, as shown. Click Apply to save the new jukebox settings, or click Reset to restore the previous settings.

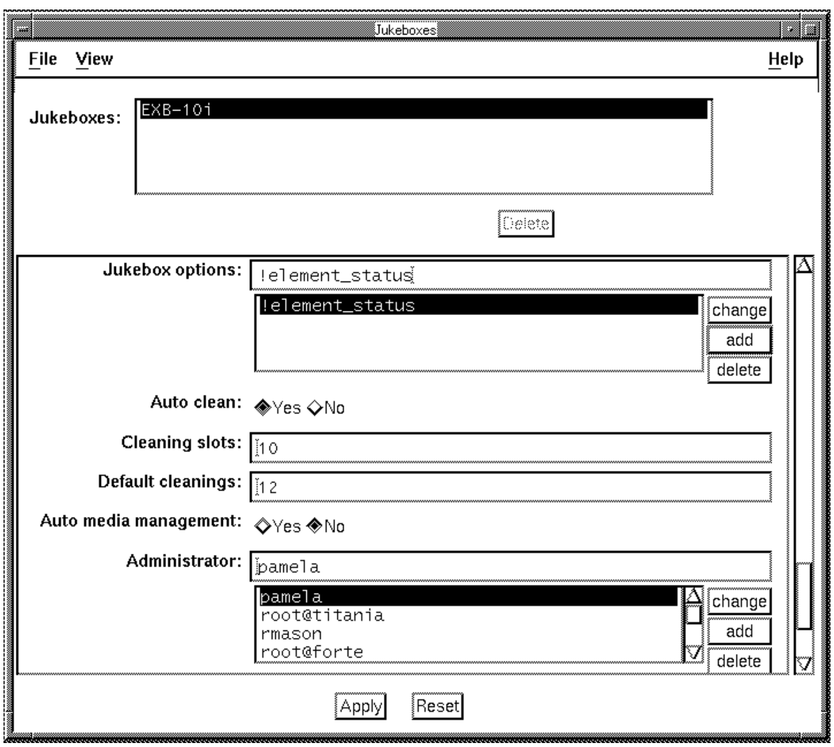

To restore the element status feature, select the setting from the Jukebox options list and click delete, then click Apply.

Understanding the Inventory Process

When Backup labels the contents of a jukebox, it registers the location of the backup volumes in the jukebox slots when it assigns the volume label. As long as you do not change the volumes in the jukebox after labeling them, Backup can access the volumes because each volume label is assigned to a specific slot.

However, if you change the contents of the jukebox without performing the labeling process, or if you move volumes into new slots, you must inform Backup that the jukebox now holds a different set of labeled volumes or that the volumes are in a different order. *This is called taking inventory.*

When you inventory the volumes in the jukebox, Backup reads the label of each volume and records its slot number. For example, if you have more than one job pack for a jukebox, you must inventory each time you remove one job pack and load another one into the jukebox, if you do not label the volumes in the new job pack.

Caution – Backup provides the ability for reading bar code labels to speed up the inventory process. We strongly recommend that you use bar code labels if you have a large number of backup volumes or change your jukebox contents often. See “Labeling with Bar Codes” in this chapter for more information.

Checking Jukebox Notifications

Backup uses electronic mail to send notices about Backup events. The software uses the three Tape mount request notices to inform the system administrator that the jukebox needs attention. The notification Tape mount request 1 is undefined so you can create your own notification message.

The following situations require attention:

- the backup volumes in the jukebox are 90% full
- the jukebox needs more backup volumes to continue
- the jukebox has a mechanical problem
- the jukebox device needs cleaning
- the cleaning cartridge in the jukebox needs replacing

Select Notifications from the Customize pull-down menu to open the Notifications dialog box. The Notifications dialog box contains tape mount and device cleaning request notifications pertaining to jukebox operations.

Note – See “Setting up Event Notification” in Chapter 7 for more information.

Backup automatically mounts the correct volume so that the recovery proceeds as long as the volume is loaded in the jukebox. If Backup requires backup volumes for a recovery that are not loaded in the jukebox, you receive a notice in the Pending window of the Backup Administrator program.

Caution – After you correct a jukebox problem, you may need to mount a backup volume before continuing to back up or recover files. Check the Pending window in the Backup Administrator program for messages.

Suggestions for Operating Jukeboxes

This section provides additional suggestions to help you use your jukebox and backup media effectively and reliably.

Recycling versus Adding More Backup Volumes

Backup can save files on volumes marked appen (appendable) in the Volumes dialog box. If the volumes inside of the jukebox are marked full, they cannot receive additional backups.

You can do one of the following with volumes marked full.

- If you need to keep the volumes for long-term storage, remove the full volumes and replace them with new media.
- If you do not need the data on the full volumes, you can manually change the mode to “recyc” in the Volumes dialog box and Backup will overwrite the data with new backups, yet maintain the existing labels. This is the only instance where you do not need to relabel a volume to make it eligible to be overwritten with new data.

Note – The mode of a volume is automatically changed to recyclable when all of the save sets on the volume have passed the time period specified by its retention policy.

There are advantages to both recycling media and adding more media to a pool. By recycling, you re-use the same backup volumes and do not add new volumes to the pool. However, the media may wear out over time and have a higher failure rate.

On the other hand, if your site requires that you maintain the backups in storage for a certain period of time, then you may have to add more media to the volume pool instead of recycling it. For example, a jukebox may need new backup volumes every three months if the company policy is to maintain the backups for a year. In this case, you have to keep adding new media to the pool until you can recycle the backup volumes that contain expired or old backups.

Using the Jukebox Cartridge Access Port

A Cartridge Access Port (CAP) allows you to deposit and withdraw media cartridges in a jukebox without opening the door to the jukebox. Each time you open the door of a jukebox to add or remove media, you invalidate the status of the jukebox. You then need to re-inventory the contents of the jukebox so that Backup can track the backup media. The inventory process can take a long time to complete.

This CAP feature is useful because you can add (deposit) and remove (withdraw) backup volumes in a jukebox without having to re-inventory the jukebox.

When you use the CAP to deposit or withdraw backup volumes, Backup does not automatically take inventory, read bar code labels, or locate empty slots in the jukebox. Use the Jukebox Inventory dialog box and Jukeboxes window for these tasks.

▼ Depositing a Cartridge

To deposit a cartridge using the CAP, follow these steps:

1. **Press the button on the front of the jukebox to move the cartridge holder forward and open the CAP.**

2. **Place the cartridge in the holder.**

3. **Press the button again to move the media into the jukebox and close the CAP.**

4. **At the Backup server, enter the** `nsrjb` **command at the system prompt, replacing *slot* with the slot number for the cartridge, and *volume* with the label name of the backup volume:**

```
nsrjb -d -S slot volume
```

If you do not specify a volume name with the `nsrjb` command, Backup marks the slot with a "-*" to indicate that the backup volume in that slot is unknown. You need to inventory the slot with this command:

```
nsrjb -I -S slot
```

To verify that the backup volume was deposited in the correct slot, click the Mount button in the Backup Administrator window.

▼ Withdrawing a Cartridge

To withdraw a backup volume from a specific slot in the jukebox, follow these steps:

1. **At the Backup server, enter the** `nsrjb` **command at the system prompt, replacing *slot* with the slot number for the cartridge, and *volume* with the label name of the backup volume:**

```
nsrjb -w {-S slot | volume | slot volume}
```

2. **When the media cartridge reaches the CAP, press the button to open the CAP.**

3. **Remove the backup volume and close the CAP.**

Using Pools with a Jukebox

If you plan to have full and nonfull backups, we recommend that you estimate the number of volumes needed for your full backups and assign them to the Full pool in the Jukebox Labeling dialog box. This ensures that your full backups are located in a consecutive range of slots in the jukebox, which makes it easy for you to remove all of the volumes at the same time.

Calibrating the Devices

Check the jukebox manufacturer‘s documentation for information about the frequency and method for calibrating the loading mechanism for the jukebox device.

Archive Application 9

This chapter provides information specific to configuring and operating Backup with the Archive Application. It also includes explanations that compare and contrast the different methods used for protecting data across a network.

The Backup archive server provides file archiving and retrieval services to a range of client machines. It is packaged as an add-on extension to existing Backup backup servers, and uses the same license mechanism as Backup.

The Backup archive client can be any machine on a network that employs archive services provided by the archive server. Clients may be enabled for backups, for archives, or for both.

Data archiving is the process of taking a snapshot of files or directories as they reside on primary media (usually disk) at a given point in time. The snapshot image is typically stored on removable media, such as tape or optical disc. Once the snapshot is safely stored on removable media, related files can optionally be deleted to conserve space on disk.

Navigating the Windows

The following sections explain the features and use of the various Backup Archive windows.

Clients Window

Before archiving can occur, you must configure Backup to recognize each archive client. Select the Clients Setup command from the Clients pull-down menu to open the Clients window.

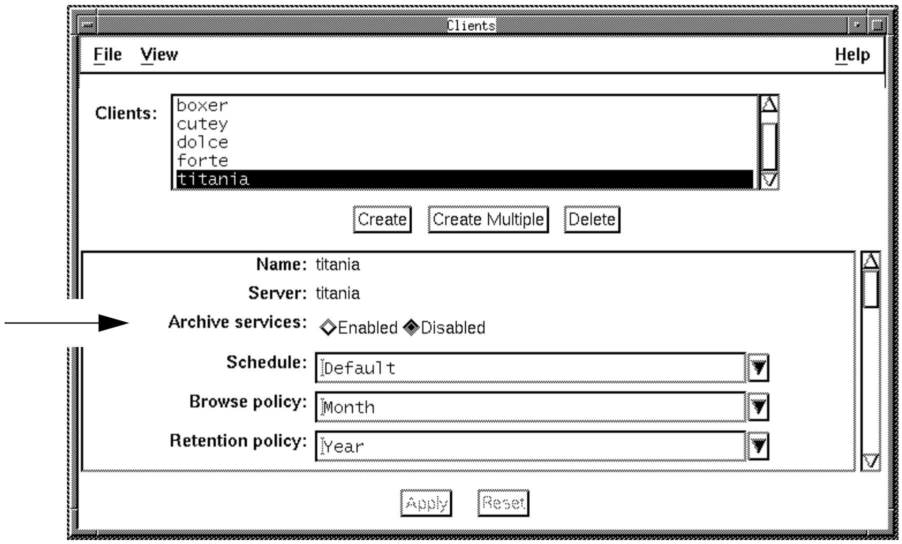

The Archive services choices enable or disable archives for the currently-selected client in the Clients list. To allow archives for the client, click the Enabled choice for Archive services. If archive services remain Disabled, the client will not be able to perform an archive.

Note – When you enable archive services for one client, you also enable other clients of the same name on that server.

Find the Archive users list at the bottom of the scrolling section of the Clients window. To allow users on the client to perform manual archives, enter their username into the Archive users list. To schedule an archive request of an entire workstation, root (or equivalent) must be on the Archive users list for that client, or root@client must be in the Administrator list for the server.

Note – For a complete description of the Clients window, see Chapter 3, “Configuring and Monitoring Clients.”

Archive Requests Window

To schedule an archive on a one-time basis, choose Archive Requests from the Customize menu.

The Archive Requests window appears, as shown.

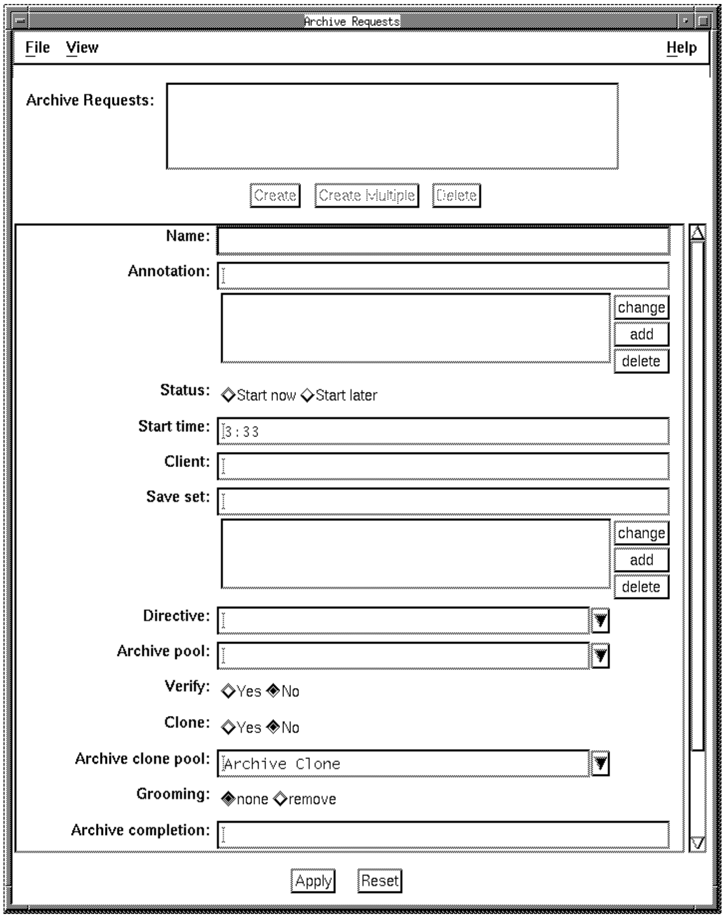

The Archive Requests scrolling list displays archives previously requested on this server. You can add new archive requests and delete old ones using the Create and Delete buttons.

The scrolling panel displays the following information:

- Name – displays the archive name.

- Annotation – a mixed-case comment string, limited to 1024 characters, which you provide for every archive request.
- Status – indicates when to begin the archive. Start now means archiving begins when you click the Apply button. Start later means archiving begins as indicated in the Start time field.
- Start time – gives the next time to begin, based on a 24-hour clock.
- Client – displays the hostname for the archive client. To request an archive for the server, enter its name in the Client field.
- Save set – specifies pathnames of directories or files to archive.
- Directive – specifies the backup method, usually Unix standard or Unix with compression.
- Archive pool – specifies the volume pool to which archives should be sorted. The default volume pool is Archive.
- Verify – indicates whether or not to automatically check the integrity of data archived on tape.
- Clone – indicates whether or not to automatically clone the archive volumes for extra security.
- Archive clone pool – specifies the volume pool for archive clones. The default archive clone pool is Archive Clone.
- Grooming – indicates whether or not client files and directories should be removed after archiving is complete.

Caution – If you use directives that include instructions to skip files, do not enable the Grooming option. Grooming occurs after a file has been archived. If a file is skipped, it cannot be groomed and will cause the archive to fail.

- Archive completion – contains optional command to execute after archiving is complete, for example `/usr/ucb/mail`.

Archive Request Control Window

To see the results of an archive request, or to start or disable a scheduled archive, choose Archive Request Control from the Server pull-down menu.

The Archive Request Control window appears, as shown.

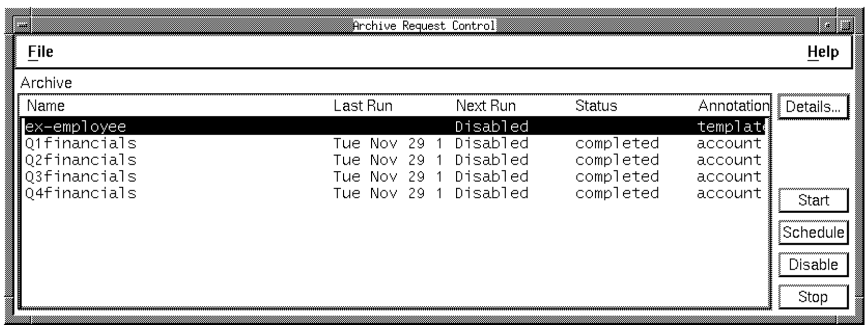

The most-recently-run archive request is initially highlighted in the Archive list. Usually it is Disabled, since scheduled archives only run once. To see details about when this archive completed and how successfully it ran, click the Details button.

If you want details about, or control over, a different archive request, select the one you want from the Archive list.

To initiate an archive request immediately, click the Start button. This has the same effect as selecting a Status of Start now for that archive request in the Archive Requests window.

To reschedule an archive request, click the Schedule button and enter a new starting time. This has the same effect as selecting Start later and specifying a new Start time in the Archive Requests window.

To disable an archive request that you scheduled here or in the Archive Requests window, click the Disable button. To halt an archive in progress, click the Stop button.

Archive Request Details

To see the progress of a recent archive (the one currently selected), click the Details button inside the Archive Request Control window.

The Archive Request Details window appears, as shown.

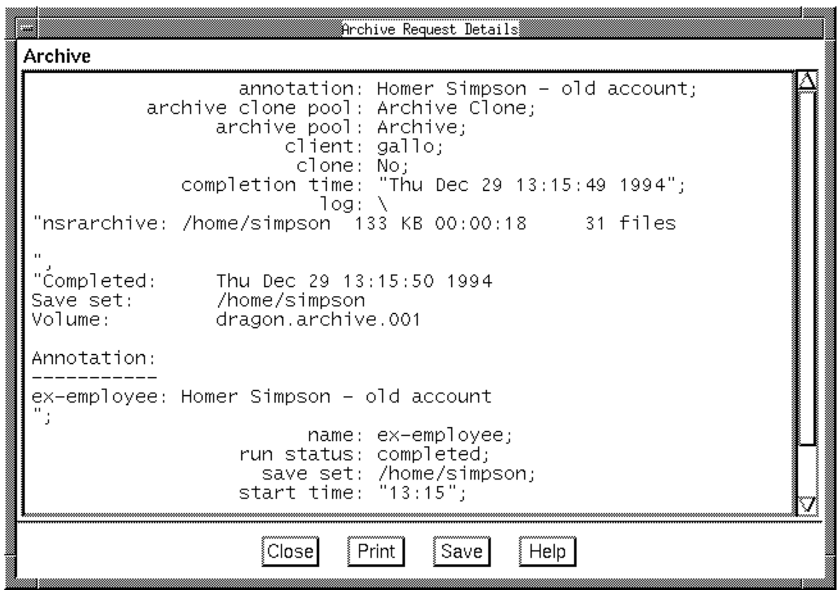

This window provides information about the completion of an archive request. Some of the information in this window, such as annotation, archive clone pool, archive pool, client, clone, name, save set, start time, status, and verify, also appears in the Archive Requests window.

The completion time field displays when the archive finished. Its duration is the difference between this and start time. The log field shows messages generated by the archive. The run status field shows the outcome of the archive request, either completed, failed, or partial.

Note – For more information about failed archives, see the log file, usually `/nsr/logs/daemon.log` on the server.

Archive Example

Suppose you must shut down and remove the workstation (and *hostname*) of someone who has left the company. It would be wise to archive the system data first, in case its filesystems contain essential files you need to access later.

This section gives an example of how to schedule and run an archive request.

▼ Creating an Archive Client

Only registered archive clients can use the archive facility. To create an archive client, follow these steps.

1. **Choose Client Setup from the Clients menu. The Clients window appears, as shown on page 236.**
2. **Click the Create button; the Clients window changes.**
3. **Enter the hostname of the workstation in the client Name field.**
4. **Click Enabled after Archive services to allow archives for this client.**
5. **If you want to permit users on the archive client machine to use archive and retrieve, scroll to the bottom of the Clients window and add their user names to the Archive users field.**
   (It is unlikely that you would want to allow manual archives and retrieves on a workstation about to be shut down.)

Machine *hostname* is now a registered archive client. However, an archive will not take place until you request one.

Note – If you want to allow archives on the server, make sure that archives are enabled for the server as a client of itself.

▼ Making an Archive Request

Valid archive users may request archives manually using the `nwarchive` command. However, a manual archive often takes a long time. To avoid overloading a busy network with an archive request during the day, schedule it late at night when the network has less traffic.

To make an archive request, follow these steps:

1. **Choose Archive Requests from the Customize menu. The Archive Requests window appears, as shown on page 237.**
2. **Click the Create button; the Archive Requests window changes.**
3. **Enter the Name you want to assign to this archive request, and a brief Annotation to remind you of the purpose for the archive.**
4. **Click Start later for the Status choice to schedule this archive for that night.**
5. **Enter the Start time you want, or accept the default starting time of 3:33 a.m.**
6. **Enter the archive client machine *hostname* in the Client field, and the *pathname(s)* you want to archive in the Save set field.**
7. **Specify a custom Archive pool, or accept the default Archive volume pool.**
8. **Click Yes for the Verify choice to check that the archived data was saved correctly.**
9. **If you want to make a duplicate copy of this archive volume, click Yes for the Clone choice.**
   Accept the default Archive Clone pool.
10. **To remove files and directories from disk after archiving them to tape, click remove for the Grooming choice.**
11. **To be notified when the archive completes, enter a command into the Archive completion field, for example:**

```
/usr/ucb/mail -s archive_request admin@titania
```

12. **Click the Apply button to activate your changes.**

You have now requested an archive of a client machine *hostname* to begin that night at 3:33 a.m.

Checking the Archive Request

The next morning, you should check the outcome of the archive. If you set up an Archive completion notice, look for an e-mail message containing a log of the archive request.

To check details of the archive, follow these steps:

1. **Choose Archive Request Control from the Server menu. The Archive Request Control window appears, as shown on page 239.**
2. **Click the Details button. The Archive Request Details window appears, as shown on page 240.**
3. **If this window shows the archive completed successfully, you can safely reconfigure the ex-employee's machine.**
   If the archive failed, you can reschedule it (see the following section.)

You may use the Archive Request Control window to start, schedule, disable, or stop another archive.

Rescheduling the Archive Request

Suppose that the archive of *hostname* did not complete last night because, for example, an energy-conscious employee turned off the computer. You decide to reschedule for the next night.

To reschedule the archive, follow these steps:

1. **Choose Archive Request Control from the Server menu. The Archive Request Control window appears, as shown on page 239. Make sure that the archive request you want is highlighted.**
2. **Click the Schedule button. The Archive Request Schedule window appears, as shown.**

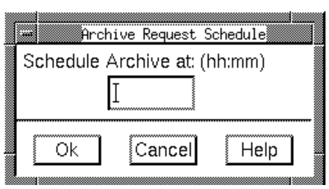

3. **Enter a new starting time in the Schedule Archive field, using the 24-hour clock, and click Ok.**

The archive request executes again that night, at the time you specified. If you change your mind and want to discontinue the archive, click the Disable button in the Archive Request Control window.

▼ Clone and Verify

Backup contains two preconfigured volume pools for use with archiving: Archive and Archive Clone. If you want to create a new volume pool for archives, follow these steps:

1. **Choose Pools from the Media menu.**
2. **Click the Create button in the Pools window.**
3. **Fill out all the fields according to your needs.**

Caution – Make sure the Pool type is set to Archive, and the Store index entries field is set to No. These two traits distinguish archive pools from backup pools.

4. **Click the Apply button to activate the new volume pool.**

Now when you schedule a new archive request, you may use the new archive volume pool you created. If you choose to clone archive data, you should also create a new archive clone pool. Backup will write archive data only to an archive volume and archive clone data only to an archive clone volume.

If you have already made an archive and want to make a clone of it, follow these steps:

1. **Choose Clone from the Save set menu.**
2. **Enter criteria for locating save sets in the Save Set Clone window.**

Note – Click the More button and enter "Archive" in the Pool field.

3. **Click the Query button to see save sets matching your criteria.**
4. **Select the save set you wish to clone and click the Clone button.**
5. **Click the Start button in the Save Set Status Clone window to activate the save set clone.**

To verify data already archived on an archive volume, you have two alternatives:

- Clone the archive data by choosing Clone from the Save set menu. During cloning, the original archive data will be verified as the save set is copied from one volume to another. When you are done, you may re-use the cloned volume.
- Determine the save set ID of the archived data, for example by searching the Backup Archive Retrieve window. Then run the following command at the system prompt, substituting the save set ID of the archived data for *ssid* (this is how Backup verifies archived data):

```
# nsrretrieve -n -S ssid
```

Both alternatives verify the integrity of the data, but do not actually compare archived data with data on disk.

Chapter 5 and Chapter 6 in this manual provide more information about volume pools and save set cloning.

▼ Archiving Shortcut

To schedule and run an archive request, follow this general procedure:

1. **Create and enable an archive client using the Clients window.**

2. **Using the Archive Requests window, fill in all or most of the fields with your preferences.**
   Name and Client are mandatory, as is the Start time if you Start later. Save set defaults to empty, Directive and Archive pool to the default, Verify and Clone to No, and Grooming to none.

3. **Check the archive status in the Archive Request Details window.**

4. **To reschedule an archive, bring up the Archive Request Schedule window.**

   a. **To discontinue a scheduled archive request, click the Disable button.**

   b. **To start and stop an archive request (for example, to test it), click the Start and Stop buttons.**

Understanding the Archive Feature

Archive save sets are similar to backup save sets. The principal difference is that archive save sets have no expiration date. Also, archives are always full – there are no levels of differential saves, or incremental saves.

Note – Archives are not recorded in the online file index, so they are not affected by the browse policy. This feature helps conserve disk space.

Retrieve is similar to recover, except that it works with archive save sets instead of backup save sets. Since archived files are not recorded in the file index, the user interface for retrieve is based on save sets, rather than on a directory hierarchy.

Users on the Backup administrator list have permission to configure archive services. These users, and users registered on the Archive users list in the Clients window, have permission to use the archiving and retrieval facilities. Registered users may archive any file for which they have read permission.

Anyone can browse archive save sets – that is, look at information in the media database. However, you may only retrieve files that you own, unless you are *superuser*, in which case you may retrieve any file.

Note – If you want to overwrite an archive tape, first make sure nobody will ever need the data again. Then simply relabel the volume, as you would a backup volume.

Archive Functions

The following describes the three categories of archive functions.

Data Archiving

Data archiving can be performed by end users or by the system administrator. Registered users can perform manual archives that start right away. System administrators can perform manual archives or schedule an archive to take place anytime during the next 24 hours. For example, the best time to perform a large archive might be in the middle of the night.

Both users and administrators can request an extra copy of their archive save set, called a *clone*. Backup Archive employs the volume pools feature to separate backup volumes from archive volumes and archive volumes from archive clone volumes.

After archiving is complete, users and administrators are given the option of deleting archived files and directories or leaving them in place. This option is called *grooming*. Grooming helps conserve disk space after a project is finished.

This chapter describes the Archive Application using the Backup GUI. For information on archiving using the command line interface, see the `nsrarchive` man page.

Data Verification

Since archived files are often deleted from the system, Backup provides an extra measure of security to make sure archived data is correct. Backup verifies data in two ways.

- media verification – Backup checks the archive volume to ensure it is writable and contains no bad spots.
- data verification – Backup reads data from the archive volume as if doing a retrieve, but does not actually write any archived data back to disk.

If a volume is suspect, or if there are problems with the data on the tape, Backup issues a warning and suspends grooming.

If you decide to groom files, we recommend you also select Verify or Clone to avoid deleting improperly archived files.

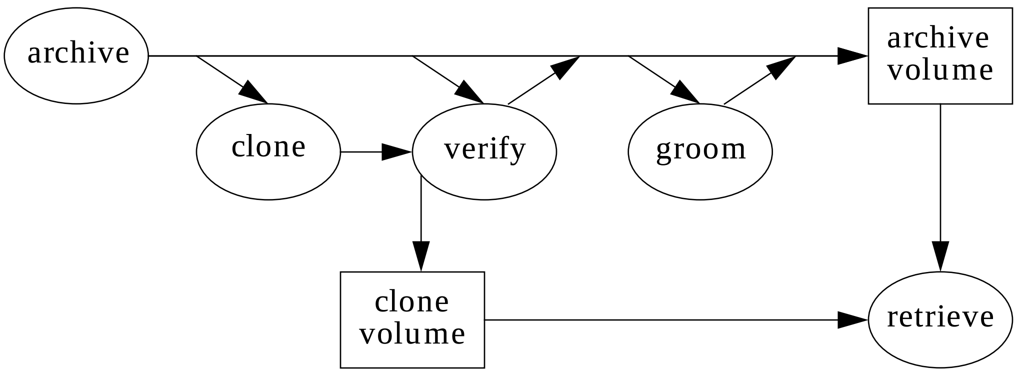

Data Retrieval

When you use Backup Retrieve, the Backup Retrieve window displays archived save sets for the selected server, listed by client name. You can only retrieve a save set if you have administrator or archive user privileges for that server, are the owner of files in the save set, or are root.

It is possible to search for specific archives and to alter the sort order of archive save sets in the viewing list. See Chapter 5, “Archiving and Retrieving Files” in the User’s Guide for more details on retrieving archived files.

When the user picks an archive save set to retrieve, and the administrator ensures that the relevant archive volume (or a clone of that volume) is mounted on the Backup server, retrieval can begin.

Retrieved save sets can be relocated, renamed, or allowed to overwrite existing files of the same name, as with the Backup recover feature.

Methods for Protecting Data

Data backup is the process of storing copies of files and directories from local disk onto removable media, usually tapes. These copies can be recovered in case the original files are lost or damaged. The system administrator usually schedules backups on a daily basis. Any new files, or files that changed since the last backup, are copied to tape so they can be restored on disk if necessary.

Archiving is normally performed on data associated with specific projects, rather than on an entire system. Unlike data backup, end users usually archive their files on an as-needed basis, so a network-wide archiving policy is not needed. Archives, unlike backups, are not associated with a level (full, differential, or incremental).

When users archive project data, they can choose to automatically delete the files from the system disk to conserve space. In this case, archived files need to be placed on long-lasting archive media.

Hierarchical storage management (HSM) is a data management strategy where data is automatically migrated from one storage medium to another, based on a set of rules. The rule most often employed is *access rate* – the longer a file is inactive, the more likely it is to migrate.

Storage hierarchy is usually governed by the cost of storage for each media. The benefit of HSM is that it provides users with a seemingly infinite storage capacity, at the lowest possible cost.

The principal goals of backup, archiving, and HSM are as follows:

- The goal of backup is to protect data against accidental loss or damage. Backups should be reliable and efficient.
- The goal of archiving is to conserve online storage space. Storage media must be durable, safe, and reliable.
- The goal of HSM is to conserve network storage resources. Migration and recall must be automatic and reliable.

Hierarchical Storage Management Application 10

In addition to backup and archive solutions, Backup has an optional application that provides hierarchical storage management (HSM). This chapter describes HSM and provides instructions for using the optional Backup HSM application.

The Backup HSM server provides file migration and recall services to a range of client machines. HSM is packaged as an optional extension to existing Backup backup servers, and uses the same license mechanism as Backup.

The Backup HSM client can be any machine on a network that uses file migration and recall services provided by an HSM server. Clients may be enabled for backups, for archives, for HSM, or for all three.

The HSM application provides a way to automatically move data between your local disk and other storage locations. This flexibility allows you to reduce cost of storage by moving data from more expensive, faster media to less-expensive, slower media and back again. This process is called *migrating* and *recalling*. Migrating is the process of moving files from a client to a remote migration storage device; recalling is the process of moving files from the remote storage device back to the original location on the client. Storing seldom-used files on remote storage devices frees local disk space for current activities with the stored files still readily available when needed. Except for a relatively longer access time for migrated files, the entire migration and recall process is transparent to the user.

Navigating the Windows

The following sections explain the features and how to configure and monitor the Backup HSM Application. When you enable HSM on the server, the Backup Administrator program does not change, but the migration windows and features become available.

You use two commands in the Backup Administrator window to configure and monitor HSM: Migration Setup and Migration Control.

- You define HSM parameters in the Migration window, which you access through the Migration Setup command in the Clients pull-down menu.
- You monitor HSM activities using the Migration Control window, which you access through the Migration Control option in the Server pull-down menu.

The HSM software is installed on the client as part of the Backup software. When it is enabled on the server, it is automatically enabled for all machines who are clients of that server. However, client side setup of HSM requires two additional steps.

- Make sure the `nsrexecd` program is running on the client.

Note – When you install Backup using `pkgadd` or `nsr_ize -i -c`, `nsrexecd` is automatically run. If you use `rsh` to install Backup, you have to start `nsrexecd` manually.

- Fill the fields in the Migration window.

Migration Window

The Migration Setup command in the Clients pull-down menu provides access to the Migration window. The Migration window contains all the configurable or selectable items relating to HSM services.

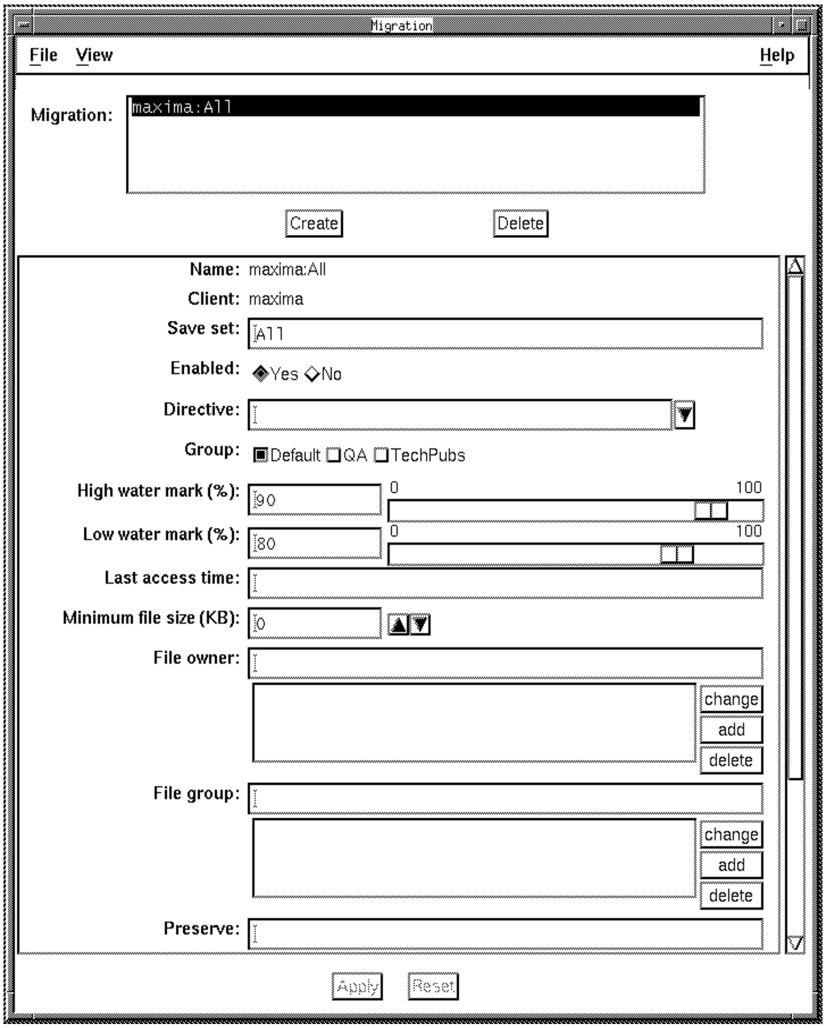

The Migration window contains fields for the criteria that define HSM rules. Clients and save sets meeting these criteria are available for pre-migration. Pre-migration copies the file to the storage location leaving the original on the client's local disk. When Backup backs up a group containing migration save sets, pre-migration occurs. This happens automatically and is not controlled by high and low-water marks. Backup requires that files be pre-migrated as part of a group. When high-water mark is reached or the local disk is full, migration occurs; the pre-migrated file is quickly deleted from the disk, leaving a stub which contains information about the migrated file.

The Migration window contains the following configurable or selectable items.

- Migration scrolling list – displays the list of clients and save sets configured for HSM services.
- Create and Delete buttons – enable adding or deleting clients to the Migration scrolling list.
- Save set field – is where you enter the name of the save set you want considered for migration. You may also enter “All” or a UNIX shell pattern.
- Enabled buttons – enable and disable automatic migration services.
- Directive field – is where you enter a directive option, usually either Unix standard directives or Unix with compression directives.
- Group field – contains a selectable button for each group you configured in the Groups window from the Customize pull-down menu. Selecting a group makes the designated save sets in that group eligible for pre-migration.
- High water mark (%) field – contains a slide bar to set the value that, when reached, starts migration. This value is the percentage of disk space filled.
- Low water mark (%) field – contains a slide bar to set the value that, when reached, stops migration. This value is the percentage of disk space filled after migration.
- Last access time field – is where you enter the value that determines which files are eligible for migration. Files that have not been accessed since the entry in this field are considered eligible for migration. The format for this entry is a number followed by seconds, minutes, hours, days, weeks, months, or years, and must end with the word “ago.” For example, 3 days ago.
- Minimum file size (KB) field – specifies the minimum size for a file to be eligible for migration. Files smaller than this entry do not provide enough available disk space, after being replaced with a stub, to make it worth migrating them.
- File owner field – limits migration to those files belonging to the name(s) entered here. Leaving this field blank allows files for all owners to be migrated.
- File group field – limits migration to those files belonging to the UNIX group name(s) entered here. Leaving this field blank allows migration for all files in all UNIX groups.

- Preserve field – is where you enter the names of files that you do not want migrated. File names must be full *pathnames*. You can enter specific file names, patterns, or wild cards. Some system files are automatically preserved, even if not specifically listed here. These files include system files, shared libraries, and all executables and data files used by Backup. On both Solaris and SunOS systems, the following files are excluded from migration:
  - all files in the `/`, `/usr`, and `/var` filesystems
  - all files that end with `.so`
  - all files (executables and data files) used by Backup

  In addition, on Solaris systems, all files in the `/opt` filesystem are excluded from migration.

The Migration window provides you with the ability to customize HSM services for your unique requirements.

Migration Control Window

The Migration Control window provides status information about migration activities that occurred during the previous seven days.

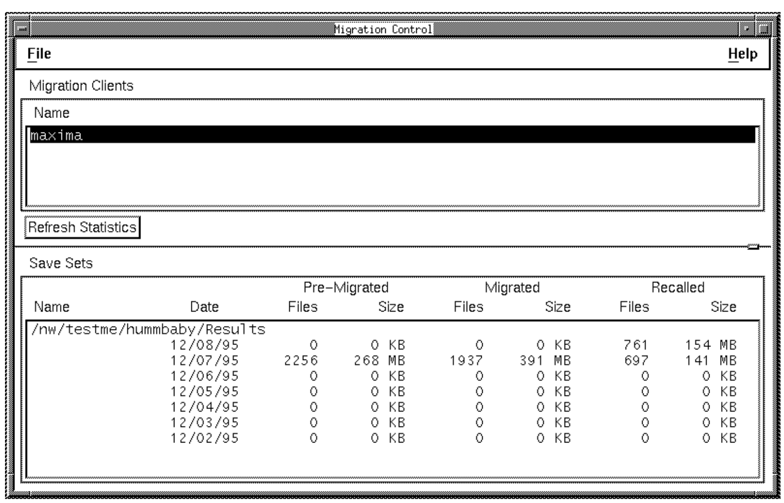

The Migration Control window displays the following:

- Migration Clients – displays a scrolling list of clients configured for HSM services.

- Refresh Statistics – clicking this button updates the Save Sets scrolling list to display current statistics.
- Save Sets – displays statistics for all migration activities that occurred in the last seven days.

You can also manually produce reports on HSM activities using command line instructions. See “Producing Reports on HSM Activity” in this chapter.

HSM Example

This section provides an example of configuring a client for pre-migration and migration services. It is intended to be a guide, not a procedure, since every site has different needs.

▼ Creating an HSM Client

This section provides detailed instructions for the steps that are unique to the HSM application. Other configuration tasks used for the Backup product, such as setting up schedules, groups, pools, and clone pools, are described elsewhere in this manual.

1. **Make sure the system that you want to receive migration services is configured in the Client Setup window as a client of the HSM server.**

Note – Detailed instructions for using the Client Setup window are provided in Chapter 3, “Configuring and Monitoring Clients.”

---

2. **Select Migration Setup from the Client pull-down menu. The Migration window appears.**
3. **Select the client you want to configure in the Migration scrolling list.**
   If the name of the client you want to configure does not appear in the Migration scrolling list, click the Create button and add the client in the Client field.
4. **Enter the names of the save set(s) you want considered for migration in the Save set field.**
   Enter “All” if you want all save sets considered for migration.
5. **If it is not selected, click the Yes button next to Enabled to allow HSM services for the selected client. (The default is Yes.)**

6. **Click the down-arrow at the right of the Directive field to display the directives choices.**

7. **Select the appropriate directive for your requirements.**

Note – Detailed information about directives is provided in Appendix A, "Managing the Backup Environment," in the *Solstice Backup 4.2 Installation and Maintenance Guide*.

---

8. **Click the appropriate group button to assign the client to a group. You set up new groups using the Groups command in the Customize pull-down menu.**
   Selecting a group makes designated files for that group eligible for pre-migration.

Next set the migration criteria for the client.

1. **Set the High water mark and Low water mark using the following format.**

2. **Set Last access time using the following format.**

```
xx yyy ago
```

   where xx = number of units and yyy = seconds, minutes, hours, days, weeks, months, or years; for example

```
3 days ago
```

3. **Click the arrows, or type in the text box, to set the Minimum file size field to the value you want considered for migration.**

4. **Enter the name of the owner(s) of the file(s) you want considered for migration in the File owner field and click Add.**
   If you want all owners allowed, leave this field blank. If you want all owners allowed except for *owner_name*, enter -*owner_name* in this field.

5. **Enter the name of the group with access to these files in the File group field and click Add.**
   If you want all groups allowed except for *group_name*, enter -*group_name* in this field.

6. **Enter the names of files you do not want migrated in the Preserve field and click Add.**
   These entries must be full pathnames and may contain UNIX shell wildcard characters. Certain system files are automatically excluded from migration.

7. **Click Apply to save your settings.**

8. **If required, you can now select a choice for Migration clone pool in the Group window.**
   The default selection, Migration Clone, already appears. For more information about configuring pools for backup groups, see Chapter 7, "Customizing Backups."

The Backup HSM application automatically migrates files meeting the defined criteria when the client filesystem reaches the value in the High water mark (%) field.

▼ Performing a Super-full

A super-full is a clone of the most recent full backup of a save set and of all migration save sets for the designated client filesystem.

For information about cloning, see Chapter 6, "Recovering and Cloning Save Sets" in this manual.

To create a super-full, use the following instructions.

1. **Log in as *root*.**

2. **Substituting the name of your client and save set, enter**

```
# nsrclone -c client_name -N save_set_name
```

A super-full provides a complete recover for a client filesystem in case of loss of data.

Producing Reports on HSM Activity

In addition to the information provided in the Migration Control window, there are several ways to produce reports detailing HSM activity. Refer to the man pages for details on using these commands.

- Use the `nsrinfo` command to list files in a save set.
- Use the `mminfo` command or the cloning browser to determine which save sets were migrated in the previous twenty-four hours.
- Use the `nsrmig -n` command to produce a report of files eligible for migration without actually migrating them.
- Use the `nsrpmig -n` command to produce a report of files eligible for pre-migration without actually pre-migrating them
- Configure the Migration Complete notification.

Understanding HSM

Hierarchical storage management (HSM) is complimentary to, and not a substitute for, backup and archiving. Backup, archive, and HSM all store data on media; however, each one has a specific purpose.

Backup is the process of storing copies of data on media. The original file is left in place. Backup is done on a regular schedule on the entire filesystem to ensure that data can be recovered if the original file is lost or damaged, a disk crashes, or in event of a disaster.

Archiving is the process of storing data that you no longer need to access regularly. Archive and retrieval are usually manual procedures. Users can initiate an archive whenever desired, usually at the end of a project. The media used for archiving must be safe and durable. After files are archived, they are usually removed from the local filesystem, thus freeing up disk space.

Hierarchical Storage Management (HSM) is a data management strategy where data is automatically migrated from one storage medium to another, based on a set of rules. The rule often employed is *access time* – the longer a file is inactive, the more likely it is to migrate. The file is removed from the local disk, leaving a stub (pointer to the new location). When you access the file, it is automatically returned to the original location. The migration and recall process is transparent to the user.

The principal goals of backup, archiving, and HSM are as follows:

- The goal of backup is to protect data against accidental loss or damage. Backups should be reliable and efficient.
- The goals of archiving are to provide extra protection and to conserve online storage space. Storage media must be durable, safe, and reliable.
- The goal of HSM is to conserve network storage resources. Migration and recall must be automatic and reliable.

A migration daemon searches the client filesystem for files which meet the administrator-defined rules. The criteria for these rules include file location, when last accessed, file size, file ownership, and amount of free space in the client filesystem.

Files meeting the specified criteria are *pre-migrated*. When Backup backs up a group containing migration save sets, pre-migration occurs. This happens automatically and is not controlled by high and low-water marks. When high-water mark is reached or the local disk is full, the pre-migrated file is quickly deleted from the disk, leaving a *stub*.

The stub is a UNIX symbolic link. The stub contains information about the file and serves two purposes:

- as a place holder for the migrated file, making it appear as though the file is still resident on the local disk
- as a pointer to the new location, allowing the HSM software to find the migrated file and recall it to the local disk

After a file migrates, and is replaced with a stub, the user can perform the same actions on it as on any other file in the filesystem. The file can be moved, renamed, and any other action that does not require read or write access.

When a user accesses a migrated file to read, write, or change attributes, Backup recalls the file to the location of the stub. After the recall starts, the file begins to open. The recall operation takes place concurrently with the opening of the file. Therefore, the user notices a delay in reading and writing until the file is completely recalled to its original position.

When a file is removed from the filesystem name space, it is also removed from the migration store.

Backup can perform “super-fulls” for HSM clients. A super-full is a clone (or copy) of the most recent full backup of a save set and a clone of all migration save sets. A super-full contains a complete backup of the data on the client and the data in the migrated store. This is very important if you ever need to recover from a disaster.

File Migration

File migration can be either scheduled, on-demand, or both, depending on the requirements of your system. All you have to do is define your criteria and assign the appropriate criteria to each client. Backup automatically migrates each client's files that meet those criteria. Backup automatically recalls a migrated file when a user accesses it.

Scheduled Migration

Scheduled migration is a "sweeping" operation determined by the assigned rules. Backup generates lists of files that are candidates for migration according to the criteria you define. The most frequently-used parameter to determine these candidates is access time. You can enable or disable migration for each client. Certain files are always excluded from migration. These files include system files, shared libraries, and all executables and data files used by Backup.

On both Solaris and SunOS systems, the following files are excluded from migration:

- all files in the */*, */usr*, and */var* filesystems
- all files that end with .so
- all files (executables and data files) used by Backup

In addition, on Solaris systems, all files in the */opt* filesystem are excluded from migration.

Backup automatically controls migration with rules you define. In this case, typical criteria for rules are high-water mark and low-water mark. When the capacity of a filesystem approaches the high-water mark, Backup HSM automatically migrates files until the low-water mark is reached.

On-demand Migration

You can perform on-demand, or manual, migration when the file system is full or is in the process of becoming full. Only those files that have been pre-migrated can be manually migrated. When you perform manual file migration, migrating large files provides the most benefit. See the `nsrpmig` and `nsrmig` man pages for additional information.

Simple Network Management Protocol Module — 11

This chapter provides instructions for using the Simple Network Management Protocol (SNMP) Module from your network management window. The SNMP Module is a separate optional add-on module for Solstice Backup. The SNMP Module allows communication of Backup event notifications to SNMP-compliant network management stations through the standard SNMP trap mechanism.

For detailed explanations of SNMP operation, refer to your network management software documentation.

The Backup SNMP option enables system administrators using network management software to do the following:

- receive Backup event notifications and status information through the network management window
- launch the Solstice Backup Administrator window from the Tools menu in the Domain Manager window or the Misc menu in the HP OpenView Network Node Manager window

---

Note – An SNMP trap is an unsolicited notification sent from the SNMP agent to the network manager's event manager. The SNMP daemon must be running on the Backup server and the network management station to send and receive SNMP traps.

Configuring Notifications

Icon(s) representing your Backup server(s) are present on your network management console. From this console using your network management software, you

- configure the manner of event trap notifications (for example, flashing icon or color change)
- create new SNMP notification schemes, through the Backup Administrator window, with different priorities and events
- track pending, alert, and other configured messages

Backup SNMP Defaults

The following table provides Backup specific SNMP information you may want to know.

SNMP Parameter	Backup Default
host-name	network management station name
community	public
enterprise object ID	160 (.1.3.6.1.4.1.160)
trap-type	1

▼ Configuring SunNet Manager

The following modification to the `/var/opt/snm/snmp.traps` file formats the trap messages and notifications from the Backup server to make them more readable.

1. **Use your favorite text editor to open `snmp.traps` for editing.**
2. **Add the following lines to the end of the file:**

```
enterprise 1.3.6.1.4.1.160
1 Backup_Trap
```

Using SunNet Manager

When you start SunNet Manager, the following window displays.

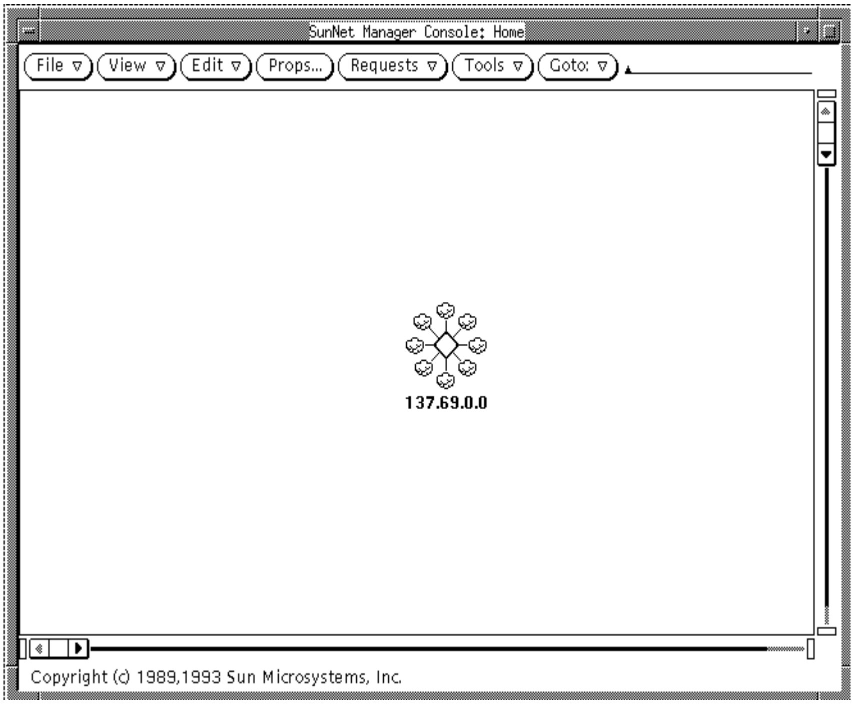

▼ Adding the Backup Option

Use the following instructions to add the Backup option to the SunNet Manager Tools menu.

1. Click the Tools pull-down menu.

2. **Select Customize; the SunNet Manager Console: Custom Tools window displays.**

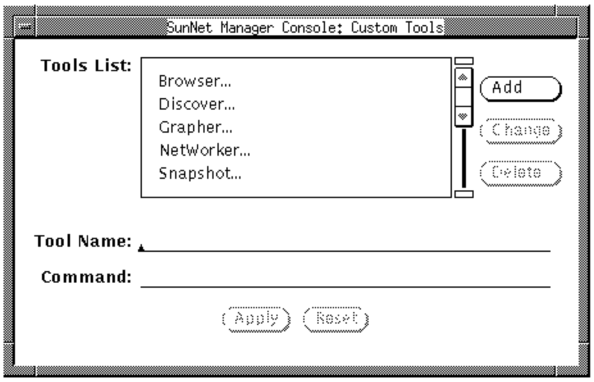

3. **Enter the following in the Tool Name field:**

```
Backup
```

Note – If your network management station is not running Backup, or if you need to access more than one Backup server, specify the name of your Backup server(s) in this field by entering the following:

```
Backup_server-name
```

If you are adding more than one Backup server to the Tools List, be sure to give each one a unique name.

4. **Enter the following in the Command field:**

```
nwadmin
```

Note – If your network management station is not running Backup, enter the following in the Command field to specify the name of your Backup server.

```
nwadmin -s server-name
```

5. **Click the Add button, and the Backup option appears in the Tools List.**

Configuring a Backup Server

Use the following instructions to configure your Backup server to send event and trap notifications to SunNet Manager.

1. **Double-click the icon in the center of the window to gain access to servers and clients on your network.**
2. **Expand the network to display the server you are going to configure.**
3. **Select the server you are going to configure.**
4. **Select Backup in the Tools menu; the Backup Administrator window appears. See Chapter 2, "Getting Started," for a complete description of the Backup Administrator window.**
5. **Click the Change Server speedbar button.**

6. **Enter the hostname of the server you plan to configure in the Server field and click Ok. In the Customize pull-down menu select Notifications; a window similar to the following appears.**

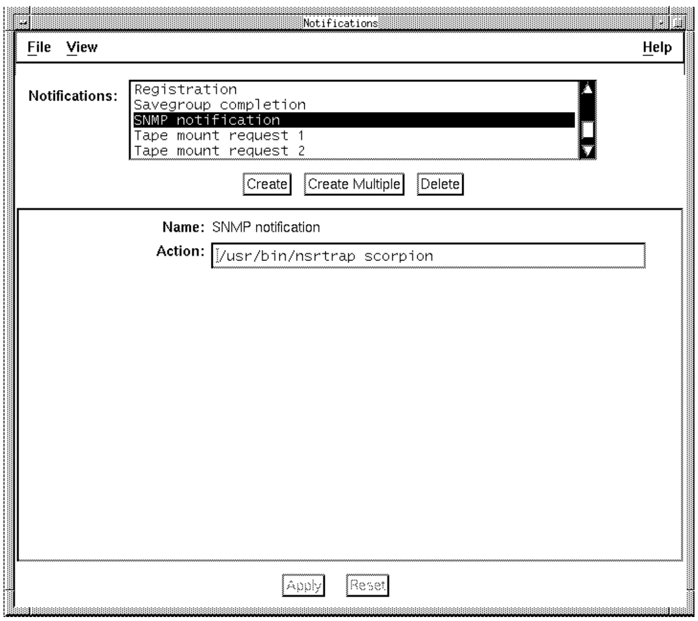

7. **In the Notifications scrolling list, select the required notifications.** Refer to the nsrtrap man page for information about creating notifications.

8. **Click the Create or the Create Multiple button to confirm your selection(s).**

9. **Click the Apply button to implement your notification.**

Customizing a Notification

Use the following instructions to customize notifications. You can set priorities, specify which types of events send traps, and which traps are sent to specific destinations.

1. In the View menu select the Details option.

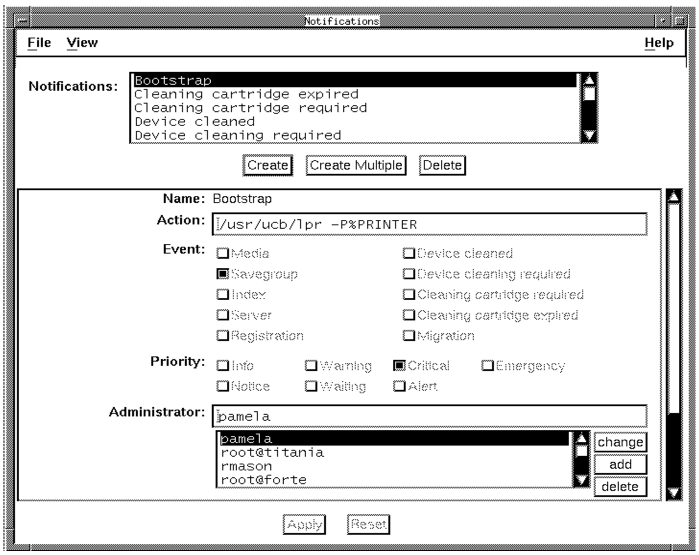

2. Click to highlight the boxes that represent the events and priorities that meet the needs of the notification displayed in the Action field.

Glossary

This glossary contains terms and definitions found in this manual. Most of the terms are specific to Backup products.

1-9

Intermediate backup levels. Each number represents a backup level. Lower levels back up more files.

8mm

Choice in the Devices window; represents eight-millimeter magnetic tape.

archive

The process by which Backup backs up directories or files 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.

Backup

The network-based software product to back up and recover filesystems.

Backup client

A machine that can access the backup and recover services from a Backup server.

Backup daemons

Daemons specific to the Backup environment.

Backup server

The machine on a network running the Backup software, containing the online indexes, and providing the backup and recover services to the clients on a network.

browse policy

The browse policy determines how long entries for your files will remain in the online file index.

client

A machine that accesses the Backup server to back up or recover files. Clients may be workstations, PCs, or fileservers with gigabytes of data.

clone

The process by which Backup makes an exact copy of saved data (save sets). Backup can clone individual save sets or entire backup volumes.

command line

The shell prompt, where you enter commands.

compressasm

The Backup directive used for compressing and decompressing files.

device

The backup device connected to the Backup 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 Backup 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.

heterogeneous

Heterogeneous networks are networks with systems of different platforms that interact meaningfully across the network.

himt

Choice of in the Devices window; represents half-inch magnetic tape.

HSM

Hierarchical storage management provides a way to automatically move data between a local disk and other storage locations to conserve network storage resources.

incremental (i)

A backup level in which only files that have changed since the last backup are backed up.

interoperability

The ability of software and hardware on multiple machines from multiple vendors to communicate meaningfully.

job pack

A tray or carousel that holds several backup volumes.

jukebox

A device which has the ability to move media among various components located in the device, including slots and media drives. Jukeboxes automate loading, labeling, and mounting of the media, during backups and recovers.

level [1-9]

A backup level that backs up files that have changed since the last backup of any lower level.

machine

Any computer, including file or compute servers, diskfull workstations, or diskless workstations.

media

Magnetic tape or optical disks used to back up files.

media index

A database of information maintained by Backup which tracks every backup volume.

media manager

The Backup component that tracks save sets to backup volumes.

migrate

The process of moving files meeting administrator-defined criteria from local disk to slower, less-expensive storage devices.

namespace

An application associated with an entry in the online file index.

notice

A response to a Backup event.

nsrhost

The logical *hostname* of the machine that is the Backup server.

online indexes

The databases located on the server that contain all the information pertaining to the client backups and backup volumes.

operator

The person who monitors the server status, loads backup volumes into the server devices, and otherwise executes day-to-day tasks using Backup.

override

A backup level that takes place instead of the scheduled one.

preconfigured

Existing selections or configurations for different Backup features.

pre-migration

Command used for making a copy of a file in preparation for migration.

print

Send data to a printer.

qic

A choice in the Devices window; represents quarter-inch cartridge tape.

recall

The process of moving a migrated file from a storage location back to the original location on local disk.

recover

The Backup command used to browse the server index and recover files from a backup volume to a client's disk.

retention policy

The retention policy determines how long entries will be retained in the media index and thus be recoverable.

recycle

A volume whose data has passed both its browse and retention policies and is available for relabeling.

save

The Backup 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 Backup.

save set ID

An internal identification number assigned to a save set by Backup.

scanner

The Backup command used to read a backup volume.

server

The machine on a network running the Backup software, containing the online index, and providing backup and recover services to the clients on a network.

shell prompt

The command line to which you enter UNIX commands.

skip (s)

A backup level in which files are skipped and not backed up.

SNMP

Simple NetWork Management Protocol

stub

A UNIX symbolic link that serves as a place holder for a migrated file and as a pointer to the new location of a migrated file.

super-full

A clone of the most recent full backup of a save set and of all migration save sets for the designated client filesystem.

super-user

A UNIX user with *root* privileges.

system administrator

The person normally responsible for installing, configuring, and maintaining Backup.

user

The person who can use Backup 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 Backup.

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

B

C

D

E

F

G

H

M

N

O

P

Q

R

S

T