Solstice Backup 4.2 Installation and Maintenance Guide

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

Part No: 802-6103-10
Revision A, April 1996

SunSoft
A Sun Microsystems, Inc. Business

Contents

Preface

The *Solstice Backup for Solaris 4.2 Installation and Maintenance Guide* contains information for installing Backup's network storage management software on a UNIX server and UNIX clients. If you have a heterogenous network and back up a variety of client machines, you also need to purchase a ClientPak.

For instructions to configure your backup environment for scheduled backups, refer to the Solstice Backup for Solaris Administrator's Guide. To learn how to use Backup's backup and recover windows for manual backups, refer to the Solstice Backup for Solaris User's Guide. If you would like more technical information about the Backup commands, refer to the online manual pages after you have installed Backup.

Note – In this document, the terms Solstice Backup and NetWorker refer to the same product. The name NetWorker appears in several GUI screens.

Audience

This manual is intended for system administrators who are responsible for installing software and maintaining the servers and clients on a network.

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

Solstice Backup for Solaris™ is network storage management software for heterogeneous client/server computer environments. Backup ™ is available in several client/server configurations, to suit the needs of any networked environment. Whether you have a single fileserver, or a network of hundreds of systems, there is a Backup product to fit your requirements.

If you would like a "hands-off" solution to your backup needs, the Solstice Backup jukebox Software Module provides automatic, unattended network-wide backups.

With the addition of the Solstice Archive Application and the Solstice Hierarchical Storage Management (HSM) Application, the Solstice Backup family of products is a complete network storage management solution.

Backup automates the day-to-day process of backing up every computer on the network, thus protecting every system from file loss. Backup simplifies the management of backup media, gives notices of Backup events, and is easy to operate and administer through an X Window System graphical user interface (GUI).

The intuitive graphical user interface helps users easily recover lost files by viewing a history of the files they have backed up, choosing the ones they want to recover, and initiating a request to recover from the Backup server.

1

Product Overview

Sun Microsystems offers four base versions of its data protection software, designed to meet the varying needs of our customers. You can easily upgrade to more powerful and feature-rich versions as your environment changes and grows. In other words, if you are backing up a single system and decide to back up additional systems on your network, you may easily do so.

Choose a solution that best fits your environment with the knowledge that you can always easily move to more performance and functionality at a later date. The following list describes the four base Backup products:

- Solstice Backup Single Server
  - backs up a single fileserver to a single backup device
  - no network backup support for clients
  - storage management uses preconfigured settings only
  - bundled with the Solaris WorkGroup and Enterprise Servers only
  - not available as a separate, orderable product
- Solstice Backup, Server Edition
  - backs up a single fileserver to multiple backup devices
  - provides preconfigured settings and allows you to create your own configurations
  - includes Solstice Backup Turbo functionality
  - supports the Jukebox Software Module(s) (all sizes)
  - supports the Solstice Archive Application
  - supports the Solstice Hierarchical Storage Management (HSM) Application
  - supports the Solstice Database Module for Oracle (DMO) Application
  - supports the Solstice Backup Simple Network Management Protocol (SNMP) Module
  - allows media cloning
- Solstice Backup, Network Edition
  - backs up a heterogenous network of systems to one or two backup devices
  - provides preconfigured settings and allows you to create your own configurations
  - includes support for (10) clients with the option to purchase additional client connections
- Solstice Backup, plus Solstice Backup Turbo
  - backs up a heterogenous network of systems to a maximum of sixteen backup devices concurrently

- provides preconfigured settings and allows you to create your own configurations
- supports the Jukebox Software Module(s) (all sizes)
- supports the Solstice Archive Application
- supports the Solstice Hierarchical Storage Management (HSM) Application
- supports the Solstice Database Module for Oracle (DMO) Application
- supports the Solstice Backup Simple Network Management Protocol (SNMP) Module
- allows media cloning

Product Matrix

The following table provides a high level comparison of the key capabilities supported by the Solstice Backup 4.2 base products:

Capability	Backup Single Server	Backup Network Edition (aka: Backup)	Backup Server Edition	Network Edition Plus Turbo
Backup Server and clients	Server only	Yes	Server only	Yes
Backup heterogeneous (Unix, PC, Netware, Nt, Mac) clients over the network	No	Yes, comes with support to backup (1) clients; options available to add more clients	No	Yes, comes with support to backup (1) clients; options available to add more clients
# of backup devices supported	(1) single tape device local to the server	(2) tape devices (non-automated) local to the server	Up to (16) tape devices and jukeboxes can be used to backup concurrently	Up to (16) tape devices and jukeboxes can be used to backup concurrently
Scheduling	Preconfigured schedules only	Advanced Customizable	Advanced Customizable	Advanced Customizable
Media Pools - means to organize backups	Preconfigured pools only	Advanced Customizable	Advanced Customizable	Advanced Customizable

Capability	Backup Single Server	Backup Network Edition (aka: Backup)	Backup Server Edition	Network Edition Plus Turbo
Separately Orderable	No, only available co-packaged w/ Solaris WG and ENT servers	Yes, available through all Sun Channels	Yes, available through all Sun Channels	Yes, available through all Sun Channels
Upgrade functionality	Purchase Network or Server Edition; Does not support other options or applications	Purchase Turbo other options and applications	Upgradable to Network Edition Plus Turbo	Purchase other options and applications

The following illustration summarizes the base products described here. As the number of systems on your network grows and your need for more performance, power, and functionality grows, Sun has a data storage management solution for you. The arrows indicate the upgrade path you can take. Your databases and collection of backup volumes are fully transferable to and interchangeable with each Sun product.

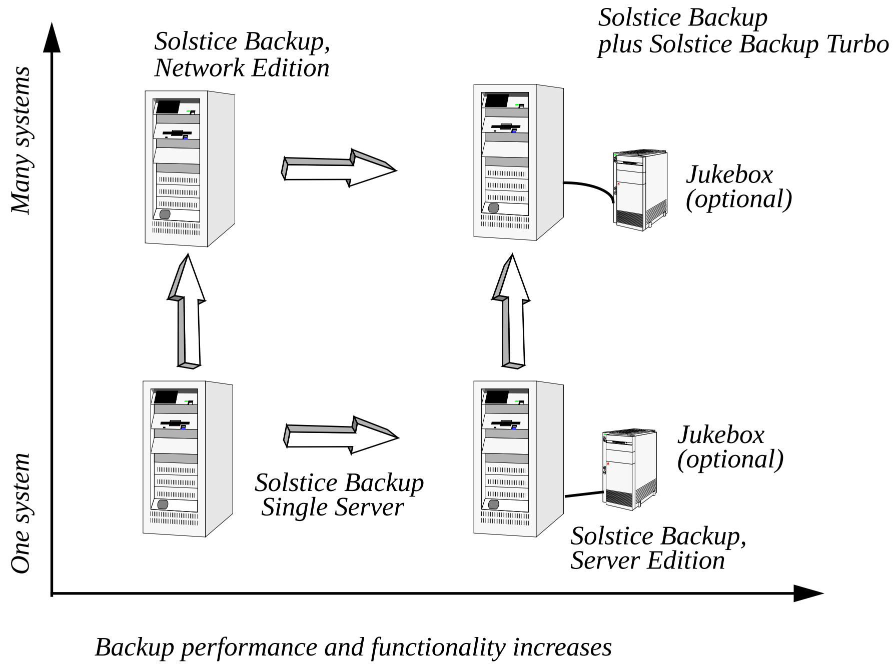

The Solstice Backup plus Solstice Backup Turbo option as well as the Server Edition provide the foundation and support for the more advanced storage management applications such as Solstice Archive, Solstice HSM and the Database Module for Oracle.

Solstice Storage Suite

In addition to base products, options, and advanced applications, Solstice Storage Suite is now available. It provides an integrated, comprehensive solution that automates the entire storage management process and creates a centralized storage center for distributed networks. It is comprised of Solstice Backup, Network Edition, Solstice Backup Turbo, Solstice Archive, Solstice HSM, Solstice Backup SNMP Module, and a 1-16 slot Jukebox module.

Solstice Backup maintains two databases of information: a *file index* and a *media index*. The file index stores an entry for every file backed up by Solstice Backup. The media index tracks the backup volumes. Backup *volumes* are storage media, as in magnetic tape or optical disk.

The *media manager* works with the two indexes to track where the backups are located on the backup volumes, making it easy to recover lost files. Whether you need to recover an entire filesystem or a single file, Backup tracks the location of the files on the backup volumes and requests the ones you need by name. You do not have to think about which backups are on which backup volumes. All you need to do to recover the required files is mount the requested backup volume in a Backup server device. Loading and mounting volumes is automatic if you are using a jukebox.

The Backup Client/Server Model

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

A Backup server is a system equipped with one or more backup devices. The server responsibilities include the following:

- writing files from many client systems to backup volumes
- maintaining an index of information about all the files backed up and on which volumes they are stored
- allowing the clients to browse the file index
- filling requests from clients to back up or recover files
- initiating automatic network-wide backups according to a specific schedule
- giving notices about Backup events

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

There is no restriction on the number of client connections Backup supports on a single backup server. However, the number of client connections you can add to your backup server successfully is affected by a wide variety of network factors.

Several factors can affect your Backup backups:

- number of client connections
- amount of data stored on each client machine
- performance/speed of server
- frequency of backups
- percentage of full backups versus incrementals
- network bandwidth
- traffic patterns

We recommend that you add clients to a single Backup server based on the performance of your backups.

The list of clients that Backup supports is constantly growing. Contact your local customer support office for the current list of supported client systems.

The clients on the network have software installed that allows the following privileges:

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

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

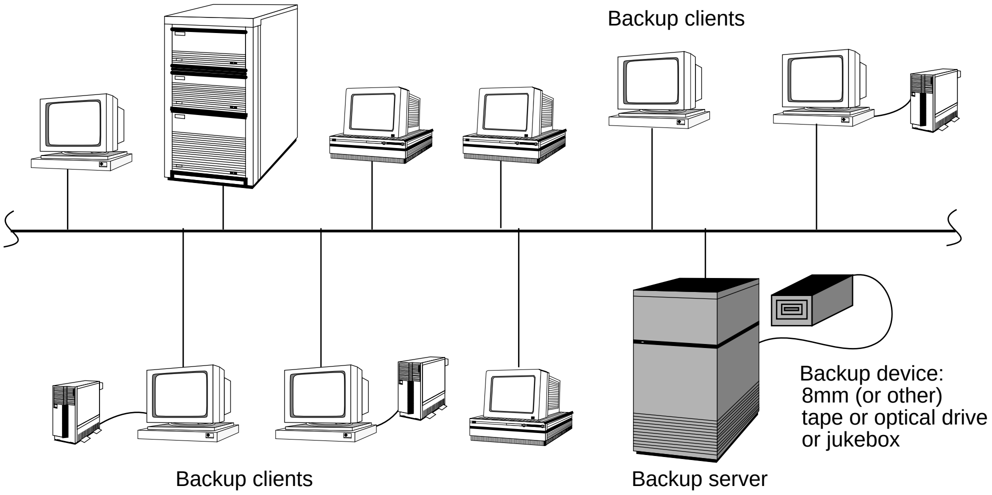

Backing up Clients

Backup backs up all the machines on a network. Your network-wide backups will be more efficient and easier to maintain because Backup has the following features.

- Preconfigured backup schedules suitable for most small-to-medium sized networks. You can easily edit these schedules using Backup's X Window System interface. You may create multi-level backups spanning any time period – weekly, monthly, bi-monthly, or quarterly.
- A sophisticated media manager that prompts you when a volume is needed for backup or file recovery. Backup sends a message through its window-based status monitor or electronic mail.
- The ability to back up to multiple premounted devices or to a jukebox for unattended backups.
- Filesystems or individual files being backed up can span multiple backup volumes or numerous filesystems and files can be saved to the same backup volume.
- The online backup feature allows you to perform backups while server and client systems remain in operation.
- Backups from many clients can take place in parallel. This feature keeps a steady stream of data supplied to the backup device or devices so they can operate at full speed. The parallel backup option also ensures that no single client can monopolize the server and prevent other clients from accessing the server. With TurboPak, Backup can back up to several devices concurrently, thus optimizing server performance even more.
- The ability to separate the media containing your full backups from media containing nonfull backups.
- The ability to sort data during backup to preselected volume pools.

▼ How to Recover Lost Files

Backup's index and media managers work together to expedite the recovery of lost files. The basic procedure for recovering files is as follows.

1. **Users browse the online index of backed-up files, using the graphical user interface, to identify the files they want to recover.**

2. **The media manager determines which backup volumes are needed to recover the files and requests only those volumes needed for recovery.**

3. **After a volume is loaded into the backup device, the files are recovered to the client.**

Note – The optional Solstice Jukebox Software Module automates the tape loading operation for fast recoveries.

SunOS 4.1.x System Installation 2

This chapter describes the procedure for installing Backup for Solaris on a SunOS 4.1.x system. Read through the procedure to become familiar with it, then decide how to organize your network of servers and clients. For example, if your network has only one system with an attached backup device, that system must be the Backup server. Backup supports a variety of media types, including 4mm, 8mm, optical disk, and digital linear tape. The list of supported devices continually grows. To obtain the latest list of supported devices, take advantage of our 24-hour FaxWorker interactive FAX server. The phone number is 415-812-6156. Follow the instructions supplied by FaxWorker, or simply request document 1905, the *NetWorker Compatibility Guide*.

Installation Overview

The distribution compact disc, read-only memory (CD-ROM) contains all of the Solstice Backup software:

- Backup Administrator, Backup, and Recover programs
- upgrade to Backup Turbo
- support for additional client connections
- optional Solstice Jukebox Software Module
- optional Solstice Archive Application
- optional Hierarchical Storage Management (HSM) Application
- optional Simple Network Management Protocol (SNMP) Module
- electronic versions of the Backup documentation set for UNIX in html format

Your distribution CD-ROM contains the Backup software for a server and clients of the same hardware platform. After installing the software on your server, use the same CD-ROM to install the client software on your Backup clients.

The clients may access the Backup software over the network or have it installed locally on their hard disks. If you install the software locally, you need to extract and install the software on each client system.

The HTML files of the Backup documentation for UNIX are provided for your convenience. You can install them at the same time as Backup or later. If you decide you would not use them or you have limited disk space, you do not have to copy them to a directory.

Once installed, you may use Backup for 30 days. After 30 days, if you have not already done so, you must purchase the appropriate enablers for the Backup products you want or the software will time-out.

Once enabled, you must *register* the Backup software as soon as possible. If you do not register Backup, it will time-out 45 days from the date you enabled it. You will not be able to use Backup to back up any more data until you register and authorize the software.

In summary, you must complete these tasks to install and use Backup.

- Extract and install the Backup software.
- Enable and register the Backup products. See Chapter 4, “Enabling and Registering Backup in this manual.

Server System and Installation Requirements

The software for a Backup server and clients running SunOS is in the *sun4* directory on the CD-ROM. Use the SunOS version 4.1.3 or later operating system with Backup.

First create a directory named `nsr_extract` and extract the software from the CD-ROM into this directory. To conserve disk space, you should delete the `nsr_extract` directory and its contents after the installation is complete.

You also need to choose a directory location on your SunOS server for the Backup programs, online indexes, and man pages. Backup chooses the following directories if you accept the default locations.

Data	Location	Space Needed
Software extraction (temporary)	/usr/tmp	39.1 Megabytes
Backup software	/usr/etc	23 Megabytes
online indexes	/usr/nsr	depends on data quantity
man pages	/usr/man	679 Kilobytes

You also need to choose a directory for the HTML files, if you decide to install them. Backup does not automatically choose a directory for the HTML files. The HTML files require approximately 9 megabytes of space.

Backup requires the following prerequisites to install the software:

- familiarity with the following UNIX commands: `su`, `mkdir`, `cd`, `tar`, `rsh`, and `mt`
- *root* privileges

To install Backup on a SunOS server, you need the following:

- A directory for installing the Backup software. Make sure there is enough space in the default installation directory. If you want to install the software in a directory other than the suggested one, be prepared to provide the directory *pathname* during the installation.

Caution – The UNIX PATH variable for the user *root* on the Backup server and on each Backup client must contain the directory where the Backup executables reside.

- A directory on the server large enough for the Backup indexes (usually `/usr/nsr`). The installation script checks for space and suggests one or more locations for the indexes.
- A directory with 679 kilobytes of disk space for the Backup online manual pages (for example, `/usr/man`).
- The system *pathname* of the device used for extracting the Backup software. For example, `/cdrom` for a system running SunOS 4.1.x.

- The system *pathname* of at least one *nonrewinding* backup device used by the Backup server to back up and recover files. For example, enter `/dev/nrst8` at the prompt for a system running SunOS 4.1.x.

If you are using an optical jukebox to back up and recover data, use the raw name of the device. For example, `/dev/rsd1c` instead of `/dev/sd1c`.

Completing a Remote Extraction Using NFS

If you do not have a CD-ROM drive on a machine where you would like to install Backup software, then the following steps will outline how to perform this over NFS.

In the following example `cdserver` is the name of the machine on which the Backup software CD-ROM is mounted, and `nwhost` is the machine that will be installing the software. In the following example it is also assumed that the CD-ROM is mounted as `/cdrom/sbu_4_2`.

1. **For a SunOS 4.1.x based cdserver add the following line to the** `/etc/exports` **file:**

```
# /cdrom/sbu_4_2
```

2. **On a SunOS 4.1.x cdserver run:**

```
cdserver# exportfs /cdrom/sbu_4_2
```

3. **On a SunOS 4.1.x nwclient add the following line to** `/etc/fstab`:

```
cdserver:/cdrom/sbu_4_2 /remotecd nfs hard,ro,intr 0 0
```

4. **Create the mount point:**

```
nwclient# mkdir /remotecd
```

5. **Mount the remote CD-ROM.**

```
nwclient# mount /remotecd
```

Then in the following sections, any path containing `/cdrom/sbu_4_2` should be replaced with `/remotecd`.

Extracting the Hypertext Markup Language (HTML) 2.0 Files

Documentation for the Backup product is also available in HTML format in the `/cdrom/sbu_4_2/Manuals/html` directory. To view this documentation, use an HTML 2.0 compliant web browser such as Netscape Navigator™.

Note – A web browser is not provided with the Backup product.

Choose Open File from the File menu in your web browser and type in the following path.

```
/cdrom/sbu_4_2/Manuals/html/index.html
```

The HTML index page is displayed. Click on the desired product name to view the documentation for that product.

For faster access to the Backup documentation in HTML format, copy the directory where the HTML files are stored to a directory in the local file system, as shown in the example.

```
# cp -r /cdrom/sbu_4_2/Manuals/html /local-path
```

Subsequently, access the HTML index page using the local file sytem path with your web browser instead of using the path to the files stored on the CD-ROM.

Installing the Server Software

To install Backup on a SunOS system, choose a directory with enough space for the software, or use the defaults provided by the installation script. The installation script asks you for the directory that will hold the Backup programs and for the device name(s) for the backup device(s) used for backing up and recovering files.

Install the online man pages if you want additional information about Backup. The online man pages require 679 kilobytes of disk space.

Files Modified during Installation

The Backup installation script, `nsr_ize`, modifies certain files on your SunOS system while installing Backup.

The following table lists the files modified by Backup:

Backup Server	**Modified Files**
SunOS system	/etc/rc.local /etc/rpc /etc/syslog.conf

Note – Save a copy of the original files before you install Backup.

Before starting the installation, change to the directory where the CD is mounted, for example, `/cdrom`.

Caution – If you are upgrading from Solstice Backup Single Server to Backup, refer to the *Solstice Backup 4.2 Product and Installation Notes* for detailed instructions. If you are updating from a previous version of Backup, you must remove the old version before installing the new software. The *Solstice Backup 4.2 Product and Installation Notes* also provides these instructions.

Use the following command at the shell prompt to install the software on the server:

```
# ./nsr_ize -i -s
```

The parts of the installation script that prompt you for information are shown in the following section. You must specify a directory for installing the software if you do not want to use the default location.

Running the SunOS Server Installation Script

This section includes instructions for installing the server software. If you are backing up to a jukebox and want to install both the server and driver software at the same time, see the section “Running the Server and Driver Installation Script” in this chapter.

Throughout the installation script, press the return key to accept the default response. Use the following command at the shell prompt to install the software on the server:

```
# ./nsr_ize -i -s
```

The installation script appears, as shown.

```
Backup(TM) - Release 4.2
Copyright (c) 1990-1996, Legato Systems, Inc. All rights reserved.
This product includes software developed by the University of
California, Berkeley and its contributors.

nsr_ize is about to install Backup server software on machine
`boxer'.

Install the Backup man pages [yes]? [Return]

Directory into which the Backup man pages
should be installed [/usr/man]? [Return]
        Installing Backup man pages into /usr/man

Install the sun4 Backup programs [yes]? [Return]
```

```
Directory where sun4 Backup programs should be installed [/usr/etc]?
[Return]
        Installing sun4 Backup programs into /usr/etc
        Creating /usr/etc/nsr_man

To set up a Backup server, you need to supply a directory with enough
free space to maintain all the on-line save file indexing and media
management information.

To set up a Backup client, you need to supply a directory for the
nsrexecd state file.

Below is a list of some of the filesystems, with their free space,
which you might consider:

/export                     :  150602
/usr                        :   66031
/                           :    8471

Directory to use for client and server information [/export]?
[Return]

Enter the tape or disk device(s) that are going to be used by the
Backup server. Use the no-rewind name for each tape device If you do
not choose a device a default device will be created for you.
(i.e.,use /dev/nrst8 instead of /dev/rst8).

Enter device name ([Return] if no more): /dev/nrst8
Select the device type for "/dev/nrst8".
The possible types are:
        8mm 5GB - 8mm double density tape.
        8mm     - 8mm tape.
        4mm     - 4mm DAT tape.
        himt    - 1/2 inch magnetic tape.
        vhs     - VHS Format Tape.
        3480    - 1/2 inch cartridge tape.
        qic     - 1/4 inch data cartridge.
        optical - Optical Disk.
        dlt     - Digital Linear Tape.
        vhs     - VHS Format Tape.

Select device type for /dev/nrst8:
(one of "8mm 4mm himt 3480 qic optical dlt 8mm 5GB")[8mm 5GB]? 8mm

Enter device name ([Return] if no more): [Return]
```

```
The nsrexecd program restricts access to a select set of Backup
servers.

Please enter the names of each computer running a Backup server that
will back up this computer, one name at a time. If a computer has
more than one network interface, please enter each interface's name
(one at a time). By default, only this server will be allowed access.
Enter the value 'all' to allow all servers to back up this computer.

Enter the first Backup server's name [no more]: all

Start Backup daemons at end of install [yes]? [Return]
        Directory /export/nsr, does not exist.
        Creating directory /export/nsr.
        Installing indexes in /export/nsr
        nsr-izing system files
        Modifying /etc/rpc
        Modifying /etc/syslog.conf
        Restarting syslog daemon
        nsr-izing system files
        Modifying /etc/rc.local
        Completing Installation
```

At this point in the script Backup asks you if you want to install and load the device driver. If you back up to a jukebox and need to load the device driver at this time, see the section “Running the Server and Driver Installation Script” in this chapter.

```
Install the Backup device drivers [yes]? [Return]
        * * * Loading Legato SCSI Interface Driver
module loaded; id = 9
        Starting Backup daemons

Backup successfully installed on `boxer'!
```

Caution – Enter the name of the directory where you installed Backup in the UNIX PATH variable for *root* on the Backup server and on each Backup client. Be sure to edit the UNIX PATH for the appropriate shell for your system (for example, if you are using `csh`, edit the `.cshrc` file; if you are using either the `Korn` or `Bourne shell`, edit the path of the `.profile` file).

Configuring the Clients to Recognize the Server

Configure the SunOS clients on your network to recognize the Backup server by defining an alias called “nsrhost” for your Backup server on each client.

For networks not using the Network Information Service (NIS) or Domain Name System (DNS), add an alias of “nsrhost” next to the name of the server machine in the `/etc/hosts` file for each client that backs up to Backup.

This section also shows you how to use NIS to configure the clients to recognize the Backup server. If you use DNS, refer to the appropriate documentation for a SunOS system.

Note – If you purchased Backup for more than one server, identify only one server as the `nsrhost` alias in the `/etc/hosts` file for each client. The server with the `nsrhost` alias is considered the primary Backup server.

▼ How to Change the Configuration for Each Client

If you *do not* have NIS (or DNS) on the network, log on to each client as *root* and edit the `/etc/hosts` file, adding the name “nsrhost” next to the Backup server name. For example:

1. **Log on to the client as *root*.**

```
jupiter#
```

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

```
137.69.1.1          venus
137.69.1.2          jupiter
137.69.1.3          mars nsrhost
137.69.1.4          mercury
```

In this example, the client named *jupiter* is backed up by the Backup server named *mars*.

3. **Repeat for each client that backs up to the Backup server.**
   The Backup clients can now access the server for backup.

▼ How to Change the Configuration Using NIS

If you have NIS on the network, follow these steps:

1. **Find your NIS master by using the** `ypwhich` **command:**

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

   In this example, *venus* is the master NIS host.

2. **Log on to the master NIS host as *root*:**

```
venus#
```

3. **Edit the** `/etc/hosts` **file, adding "nsrhost" as an alias next to the name of the Backup server:**

Note – If you have more than one Backup server, you can only use one server as the NIS master, usually your primary server.

```
137.69.1.1          venus
137.69.1.2          jupiter
137.69.1.3          mars nsrhost
137.69.1.4          mercury
```

4. **Change directory to** `/var/yp` **(or** `/etc/yp`**) and use the** `make` **command:**

```
venus# cd /var/yp
venus# make
```

   The Backup server can now access the clients for backup.

For more information about how clients bind to Backup servers, refer to the nsr manual page.

Client Installation Requirements

If you have clients of the same hardware platform as the Backup server, use the same software distribution media to install Backup on the clients. Instructions for installing the client software on like clients are provided in this section. For clients with different hardware platforms, you need to purchase and install the client software separately. Contact your local customer support office for information about the ClientPak software packages.

In order for Backup clients to use the Backup server to back up and recover files, they must be able to use the Backup software. There are two ways a client can access the Backup software.

- Clients have NFS-mounted a directory where the Backup programs are located.
- Clients have the Backup programs installed on their local disks.

Caution – In order to back up a Backup client (running Backup 4.0.2 or higher) over the network, the nsrexecd daemon must be running on the client. To enable nsrexecd, the installation script nsr_ize must be run on the client, which is described in detail in the section "Installing the Client over NFS" on page 25 in this chapter.

For more information about the access control policies for the Backup commands, refer to "Restricting Client Access" in Chapter 3 of the *Solstice Backup 4.2 Administration Guide* or the nsr man page.

If a client is not set up to share programs over the network, then you must install the Backup software directly on that client. Use the same Backup software distribution CD-ROM you used to install the software on the server.

To install the software on a SunOS client, you need the following:

- a directory with 15 megabytes of available disk space for installing the Backup client software
- a directory with 679 kilobytes of space for the online man pages

Installing the Client Software

This section describes the software installation for SunOS clients that have the same software and hardware as the server.

The installation script for a client is the same as the one for the server, except that you use the **-c** flag with the **nsr_ize** command.

Running the SunOS Client Installation Script

This installation script is used for installing the SunOS system software.

To install Backup locally on a SunOS client, enter the following command:

```
neptune# cd /cdrom-mount-point/SunOS
```

The installation script appears, as shown.

```
Backup(TM) - Release 4.2
Copyright (c) 1990-1996 Legato Systems, Inc. All rights reserved.
This product includes software developed by the University of
California, Berkeley and its contributors.

nsr_ize is about to install Backup client software on client machine
'neptune'.

Install the Backup man pages [yes]? [Return]

Directory into which the Backup man pages should be installed
[/usr/man]? [Return]
        Installing Backup man pages into /usr/man

Format new nroff copies of man pages (NOTE - this takes a while)
[yes]? [Return]

Install the sun4 Backup programs [yes]? [Return]

Directory where sun4 Backup programs
should be installed [/usr/bin]? [Return]
Installing sun4 Backup client programs into /usr/bin
        Creating /usr/bin/nsr_man
```

```
To set up a Backup server, you need to supply a directory with enough
free space to maintain all the on-line save file indexing and media
management information.

To set up a Backup client, you need to supply a directory for the
nsrexecd state file.

Below is a list of some of the filesystems, with their free space,
which you might consider:

/usr                        :   56404
/var                        :   22820
/                           :    1432

Directory to use for client and server information [/usr]? [Return]

The nsrexecd program restricts access to a select set of Backup
servers. Please enter the names of each computer running a Backup
server that will back up this computer, one name at a time. If a
computer has more than one network interface, please interface's
name (one at a time).

Enter the first Backup server's name [no more]: all
Allowing access to all Backup servers.

Start Backup daemons at end of install [yes]? [Return]
        Directory /usr/nsr, does not exist.
        Creating directory /usr/nsr.
        Installing Backup home directory in /usr/nsr
        nsr-izing system files
        nsr-izing system files
        Creating /etc/rc.nsr
        Modifying /etc/inittab
        Completing Installation
        Starting Backup daemons

Backup successfully installed on 'neptune'!
```

When you have completed installing the client software, remove the distribution media from the device and store it in a safe place.

Installing the Client over NFS

To NFS-mount Backup on a SunOS client, enter the following command:

```
neptune# nsr_ize -i -c
```

The NFS-mounted directory in which Backup is installed must be in *root's* PATH before running `nsr_ize`.

The installation script appears, as shown.

```
Backup(TM) - Release 4.2
Copyright (c) 1990-1996 Legato Systems, Inc. All rights reserved.
This product includes software developed by the University of
California, Berkeley and its contributors.

nsr_ize is about to install Backup client software on client machine
'neptune'.

Install Backup man pages [yes]? no
Install the sun4 Backup programs [yes]? no

Directory where the sun4 Backup programs are already installed
[/usr/etc]? /net/mars/usr/etc

The client-side nsrexecd program restricts access to a select set of
Backup servers. Please enter the names of each computer running a
Backup server that will back up this client, one name at a time. If
a computer has more than one network interface, please enter each
interface's name (one at a time).

Enter the first Backup server's name [no more]: mars
Enter the next Backup server's name [no more]: [Return]

* * * nsr-izing system files
        Modifying /etc/rc.local

Start Backup client daemon [yes]? [Return]
        * * * Starting Backup client daemon

Backup client successfully installed on 'neptune'
```

Once Backup is installed on both the server and the clients, refer to the *Solstice Backup 4.2 Administration Guide* for information on how to configure Backup for unattended backups. Refer to the *Solstice Backup 4.2 User Guide* to learn how to use the Backup graphical user interface to back up and recover files.

Installing the Optional Jukebox Software Module

This section contains instructions for installing the Jukebox Software Module after installing the server software, and instructions for installing the server and device driver software simultaneously.

In the following sections, the term jukebox is used to refer to a variety of backup devices, including the following: autoloader, carousel, library, near-line storage, datawheel, and autochanger.

There are four phases to installing the optional Solstice Jukebox Software Module on a SunOS system:

1. **Remove the old device driver, if necessary.**
2. **Install the device driver software.**
3. **Enable the Jukebox Software Module.**
4. **Register the Jukebox Software Module.**

Caution – **Backup** supports jukeboxes connected through SCSI or serial (RS-232) ports. If your jukebox is connected with SCSI, you must install a driver for the SCSI port.

If your jukebox is connected with a serial port, you do not need to install the driver. Simply skip the instructions in this guide regarding the driver installation. However, you do need to configure and connect your jukebox by following the hardware instructions shipped with your jukebox. You must also properly enable and register your Autochanger Software Module.

The next sections provide detailed information on installing and configuring the Jukebox Software Module for your SunOS Backup server. See Chapter 4, “Enabling and Registering Backup,” for information about enabling and registering the module.

Removing an Old Device Driver

If you are installing the Jukebox Software Module for the first time, skip this section. If you already have the Jukebox Software Module installed, you must first de-install the old driver before installing the new one.

To de-install the driver, become *root*, and change to the directory where the driver was originally installed, then enter the `deinstall` command:

```
# ./deinstall
```

Installing the Device Driver Software

The following is an overview of the procedure to complete the device driver software installation for a SunOS system:

1. **Log in as *root*.**
2. **Execute the installation script.**
3. **Reboot the system (optional).**

Running the Server and Driver Installation Script

The section includes instructions for installing both the server and device driver software. If you already extracted and installed the server software, see the section “How to Run the Driver Software Installation Script Only” on page 30

Throughout the installation script, press the return key to accept the default response.

Use the following command at the shell prompt to install the software on the server:

```
# ./nsr_ize -i -s
```

The parts of the installation script that prompt you for information follow. You must specify a directory for installing the software if you do not want to use the default location.

The installation script appears, as shown.

```
Backup(TM) - Release 4.2
Copyright (c) 1990-1996, Legato Systems, Inc. All rights reserved.
This product includes software developed by the University of
California, Berkeley and its contributors.
nsr_ize is about to install Backup server software on machine
`boxer'.

Install the Backup man pages [yes]? [Return]

Directory into which the Backup man pages
should be installed [/usr/man]? [Return]
        Installing Backup man pages into /usr/man

Install the sun4 Backup programs [yes]? [Return]

Directory where sun4 Backup programs should be installed [/usr/etc]?
[Return]
        Installing sun4 Backup programs into /usr/etc
        Creating /usr/etc/nsr_man

To set up a Backup server, you need to supply a directory with enough
free space to maintain all the
on-line save file indexing and media management information.

To set up a Backup client, you need to supply a directory for the
nsrexecd state file.

Below is a list of some of the filesystems, with their free space,
which you might consider:

/export                     :  150602
/usr                        :   66031
/                           :    8471

Directory to use for client and server information [/export/nsr]?
[Return]


Enter the tape or disk device(s) that are going to be used by the
Backup server. Use the no-rewind name for each tape device. If you
do not choose a device a default device will be created for you.
(i.e.,use /dev/nrst8 instead of /dev/rst8).

Enter device name ([Return] if no more): /dev/nrst8
```

```
Select the device type for "/dev/nrst8".
The possible types are:
        8mm 5GB - 8mm double density tape.
        8mm     - 8mm tape.
        4mm     - 4mm DAT tape.
        himt    - 1/2 inch magnetic tape.
        3480    - 1/2 inch cartridge tape.
        qic     - 1/4 inch data cartridge.
        optical - Optical Disk.
        dlt     - Digital Linear Tape.
        vhs     - VHS Format Tape.

Select device type for /dev/nrst8:
(one of "8mm 4mm himt 3480 qic optical dlt 8mm 5GB")[8mm 5GB]? 8mm

Enter device name ([Return] if no more): [Return]
The nsrexecd program restricts access to a select set of Backup
servers.

Please enter the names of each computer running a Backup server that
will back up this computer, one name at a time.  If a computer has
more than one network interface, please enter each interface's name
(one at a time).  By default, only this server will be allowed
access. Enter the value 'all' to allow all servers to back up this
computer.

Enter the first Backup server's name [no more]: all

Start Backup daemons at end of install [yes]? [Return]
        Directory /export/nsr, does not exist.
        Creating directory /export/nsr.
        Installing indexes in /export/nsr
        nsr-izing system files
        Modifying /etc/rpc
        Modifying /etc/syslog.conf
        Restarting syslog daemon
        nsr-izing system files
        Modifying /etc/rc.local
        Completing Installation

Install the Backup device drivers [yes]? [Return]
        * * * Loading Legato SCSI Interface Driver
module loaded; id = 9
        Starting Backup daemons

Backup successfully installed on `boxer'!
```

▼ How to Run the Driver Software Installation Script Only

Use this set of instructions if you want to install just the device driver software after already installing the server software. If you removed the `nsr_extract` directory where the files were originally installed, you have to re-extract the server software to install the driver.

Follow these steps to extract and install the driver software:

1. **Change to your existing `nsr` directory containing your active Backup files.**
   You want to copy the driver files residing in /*cdrom-mount-point*/`SunOS/drivers` to the `nsr` directory.

```
# cd /nsr
```

2. **Create a directory named `drivers` in the `nsr` directory to which you can copy the driver files from the CD.**

```
# mkdir drivers
```

3. **Change directory to the CD so you can copy the files over to the new `drivers` directory.**

```
# cd /cdrom-mount-point/SunOS/drivers
```

4. **Copy the driver files, including *install*, *deinstall*, and *uscidriver*, from *nsr_extract/sun4/drivers* to the new *drivers* directory.**

```
# cp -r -p * /nsr/drivers
```

5. **Change directories one more time back to */nsr/drivers*.**

```
# cd /nsr/drivers
```

6. **Begin the installation of just the device driver files by entering the install command.**

```
# ./install
```

The installation script appears, as shown.

```
Install the Backup device drivers [yes]? [Return]
       * * * Loading Legato SCSI Interface Driver
module loaded; id = 9
       Starting Backup daemons
Backup successfully installed on `boxer'
```

If you install additional jukeboxes later, you must enable and register each additional Jukebox Software Module you purchase. You only need to extract the Backup software to add an additional jukebox if you removed the software after the original installation.

▼ Testing the Device Driver Installation

After the driver loads successfully, test the driver installation.

Use the following procedure to verify that the driver loaded properly.

1. **Enter the following command:**

```
# /etc/LGTOuscsi/lusdebug 1
```

You should see the following response:

```
debug level was 0; is now 1
```

2. Enter the following command:

```
# /etc/LGTOuscsi/lusdebug 0
```

You should see the following response:

```
debug level was 1; is now 0
```

3. Enter the following command:

```
# /etc/LGTOuscsi/inquire
```

A list of SCSI devices attached to your server, if any, appears on your screen. Your list will probably look different than this one. If you attached your jukebox before installing the driver, your jukebox should appear in the list.

```
scsidev@0.0.0:FUJITSU M2263S-51201| Direct Access
scsidev@0.4.0:Quantum DLT4700| Sequential Access
scsidev@0.4.1:Quantum TZ Media Changer| Changer Device
```

▼ How to Configure a Jukebox

To configure a jukebox, you must run the `jb_config` program:

1. Become *root* on the Backup server.

```
% su root
Password:
#
```

2. Enter the `jb_config` command at the prompt:

```
# jb_config
```

3. **Backup displays the installation script and asks for your choices. The following example shows the responses for a SCSI Jukebox.**

```
1) Install an Autodetected SCSI Jukebox.
2) Install a Serial Jukebox.
3) Install an SJI Jukebox.
What kind of Jukebox are you installing? 1
These are the SCSI Jukeboxes currently attached to your system:
  1) scsidev@1.2.0: DLI Libra Series
  2) scsidev@0.2.1: Quantum DLT/Digital DLT
Which one do you want to install? 2
Installing an 'Quantum DLT/Digital DLT' jukebox.
Name you would like to assign to the jukebox device? dlt
Pathname of the control port for the jukebox device? [scsidev@0.2.1]
[Return]
Do you want automated device cleaning support enabled? (yes/no) n
Enter pathname of media drive 1 [/dev/nrst8]: ? [Return]
This media device has not been configured yet.  Please
select a media device type for /dev/nrst8.
a) himt
b) qic
c) 4mm
d) 8mm
e) 8mm 5GB
f) 3480
g) dlt
h) vhs
i)optical
Choice? g

Jukebox has been added successfully
```

If you enter **3**, “Install an SJI Jukebox,” at the first question, you see the following response:

```
What kind of Jukebox are you installing? 3
Enter the number corresponding to the type of jukebox
you are installing:
1) ADIC-1200c/ADIC-1200d           2) ADIC-VLS
3) ARC-DiamondBack                 4) Breece Hill
5) DLI Libra Series                6) Quantum DLT/Digital DLT
7) EXB-10e/EXB-10h                 8) EXB-10i
9) EXB-60                          10) EXB-120
11) EXB-210                        12) EXB-218
13) EXB-400 Series                 14) HP-C1553A/Surestore 12000e
15) Metrum (SCSI)                  16) Qualstar
17) Spectralogic                   18) STK-9704/Lago 340
```

```
19) STK-9708/Lago 380(SCSI)         20) IBM 7331/IBM 9427
21) ATL/Odetics SCSI                22) HP-Optical 630MB/1.3GB
23) other
Choice? 6
Installing an 'Quantum DLT/Digital DLT' jukebox.
Name you would like to assign to the jukebox device? dlt
Pathname of the control port for the jukebox device? [scsidev@1.2.0]
[Return]
Do you want automated device cleaning support enabled? (yes/no) n
Enter pathname of media drive 1 [/dev/nrst8]: ? [Return]
This media device has not been configured yet.  Please
select a media device type for /dev/nrst8.
a) himt
b) qic
c) 4mm
d) 8mm
e) 8mm 5GB
f) 3480
g) dlt
h) vhs
i) optical
Choice? c
```

Testing the Jukebox Connection

To test the jukebox connection, run the `jbexercise` program with two pieces of “scratch” media loaded in the first and last slots of the jukebox. The drives must be empty, with their doors open.

When running `jbexercise`, specify the control port and the device type. The control port for SCSI jukebox models is typically */dev/scsidev@x.x.x*. The exact control port pathname is shown when running `jb_config`. For example, the following command runs the `jbexercise` program on the DLT jukebox:

```
# jbexercise -c /dev/scsidev@1.2.0 -m Quantum DLT/Digital DLT
```

Refer to the `jbexercise` man page for additional information.

Enabling the Jukebox Software Module

See Chapter 4, “Enabling and Registering Backup” for instructions on enabling and registering the autochanger module.

Starting the Backup Administrator Program

By installing Backup, you configured the server to provide backup and recover services. To bring up the Backup Administrator window, enter the `nwadmin` command at the system prompt:

```
# nwadmin &
```

If you have trouble starting Backup, one of the following may be the cause:

- The Backup daemons may not be running properly.
- The DISPLAY variable on the machine you are using may not be set correctly.

See Appendix B, “Troubleshooting” for information on checking and starting the Backup daemons and on setting the DISPLAY variable correctly.

Removing the SunOS Backup Software

This section describes how to remove the Backup software for the server, client, and driver.

Removing the Server and Device Driver Software

If you need to remove the server and/or device driver software from your SunOS system, first make sure you are in the same directory where it was originally installed, then follow the steps below.

Enter the following command at the system prompt:

```
boxer# nsr_ize -r -s
```

The following script appears; enter the appropriate responses:

```
Backup(TM) - Release 4.2
Copyright (c) 1990-1996 Legato Systems, Inc. All rights reserved.
This product includes software developed by the University of
California, Berkeley and its contributors.

nsr_ize is about to remove Backup from machine `boxer'.

Shutdown any currently running Backup programs [yes]? [Return]
        * * * Killing Backup daemons

Remove the Backup man pages [yes]? [Return]
Directory from which the Backup man pages
should be removed [/usr/man]? [Return]
        Removing Backup man pages from /usr/man

Remove the Backup programs [yes]? [Return]
Directory from which the sun4 Backup programs
should be removed [/usr/etc]? [Return]
        Removing sun4 Backup programs from /usr/etc

Warning: if you answer [yes] to the following question the entire
contents of the /nsr directory will be removed. This includes the
online client index files and server media index file. You should
only answer [yes] if you want to completely remove Backup from the
system.

Do you want to remove Backup configuration and database files [no]?
[Yes]
       Removing /nsr/res/nsr.res (Backup server configuration) file
        Removing /nsr/res/nsrjb.res (Backup jukebox configuration)
file
       Removing /nsr/res/nsrla.res (Backup nsrexecd configuration)
file
        de nsr-izing system files
        Modifying /etc/rpc
        Modifying /etc/syslog.conf
        Restarting syslog daemon
        Modifying /etc/rc.local
        Removing Backup directories

Remove the Backup device drivers [yes]? [Yes]
        Removing RAP resource directory /var/rap
        Removing Backup logging directory /nsr/logs
Removing all on-line index and volume information

Backup successfully removed from `boxer'.
```

Removing Only the Drivers

If you want to remove only the drivers, and keep the Backup software, change directories to where the drivers are located and run the **deinstall** command. For example:

```
# cd /nsr/drivers
# ./deinstall
```

Removing the Client Software

Follow the instructions in this sample script to remove the SunOS client software.

```
neptune# ./nsr_ize -r -c
Backup(TM) - Release 4.2
Copyright (c) 1990-1996 Legato Systems, Inc. All rights reserved.
This product includes software developed by the University of
California, Berkeley and its contributors.

nsr_ize is about to remove Backup client software from client machine
`neptune'.

The following Backup programs are currently running and must be
shutdown before continuing:
 19174      -   0:00 nsrexecd

Shutdown currently running Backup programs [yes]? [Return]
        * * * Killing Backup daemons

Remove the Backup man pages [yes]? [Return]

Directory from which the Backup man pages
should be removed [/usr/man]? [Return]
        Removing Backup man pages from /usr/man

Remove the Backup programs [yes]? [Return]

Directory from which the sun4 Backup programs
should be removed [/usr/bin]? [Return]
Removing sun4 Backup client programs from /usr/bin
        de nsr-izing system files
        Removing /etc/rc.nsr
        Modifying /etc/inittab
        Removing Backup directories
        Removing RAP resource directory /var/rap
        Removing Backup logging directory /nsr/logs
Removing all on-line index and volume information

Backup successfully removed from ‘neptune’
```

Solaris 2.x System Installation 3

This chapter describes the procedure for installing Backup on a Solaris 2.x system. Read through the procedure to become familiar with it, then decide how to organize your network of servers and clients. For example, if your network has only one system with an attached backup device, that system must be the Backup server. Backup supports a variety of media types, including 4mm, 8mm, optical disk, and digital linear tape. The list of supported devices continually grows. To obtain the latest list of supported devices, take advantage of our 24-hour FaxWorker interactive FAX server. The phone number is 415-812-6156. Follow the instructions supplied by FaxWorker, or simply request document 1905, the NetWorker *Compatibility Guide*.

Installation Overview

The distribution compact disc, read-only memory (CD-ROM) contains all of the Backup software:

- Backup Administrator, Backup, and Recover programs
- upgrade to Backup TurboPak
- support for additional client connections
- optional Backup Autochanger Software Module
- optional Backup Archive Application
- optional Hierarchical Storage Management (HSM) Application
- optional Simple Network Management Protocol (SNMP) Module
- electronic versions of the Backup documentation set for UNIX in AnswerBook and HTML formats.

Your distribution CD-ROM contains the Backup software for a server and clients of the same hardware platform. After installing the software on your server, use the same distribution CD-ROM to install the client software on your Backup clients.

The clients may access the Backup software over the network or have it installed locally on their hard disks. If you install the software locally, you need to extract and install the software on each client system.

The electronic versions of the Backup documentation for UNIX are provided for your convenience. You can install them at the same time as Backup or later. If you decide you would not use them or you have limited disk space, you do not have to copy them to a directory.

Once installed, you may use Backup for 30 days. After 30 days, if you have not already done so, you must purchase the appropriate enablers for the SunSoft products you want or the software will time-out.

Once enabled, you must *register* the Backup software as soon as possible. If you do not register Backup, it will time-out 45 days from the date you enabled it and you will not be able to use it to back up any more data.

In summary, you must complete these tasks to install and use Backup.

- Install the Backup software.
- Enable and register the Backup products. See Chapter 4, “Enabling and Registering Backup,” in this manual.

Server System and Installation Requirements

The software for a Backup server and clients running Solaris is in the *Solaris* directory on the CD-ROM. Installation of the AnswerBook and html files is optional; they are provided as a convenience.

Note – Sun recommends that you use the Solaris version 2.3 or later operating system with Backup. Solaris 2.3 is equivalent to SunOS 5.3.

You also need to choose a directory location on your Solaris server for the Backup programs. Backup chooses the following directories if you accept the default locations.

Backup Server	Locations	Space Needed
Solaris system	/usr/bin/nsr /usr/sbin/nsr /usr/lib/nsr /usr/man	23 Megabytes 22 Megabytes 2 Megabytes 679 Kilobytes

To install Backup on a Solaris server, you need the following:

Caution – The UNIX PATH variable for the user *root* on the Backup server and on each Backup client must contain the directory where the Backup executables reside. By default, this is `/usr/sbin/nsr` and `/usr/bin/nsr`.

- A directory on the server large enough for the Backup indexes (usually `/nsr`). The installation script checks for space and suggests one or more locations for the indexes.
- A directory with 679 kilobytes of disk space for the Backup online man pages (for example, `/usr/man`).
- The installation prompts you for the system *pathname* of at least one *nonrewinding* backup device used by the Backup server to back up and recover files. For example, `/dev/rmt/0mbn` for a system running Solaris 2.3.

  If you are using an optical jukebox to back up and recover data, you should use the raw name of the device. For example, */dev/rdsk/c0t1d0s2*.

Extracting the Hypertext Markup Language (HTML) 2.0 Files

Documentation for the Backup product is also available in HTML format in the `/cdrom/sbu_4_2/Manuals/html` directory. To view this documentation, use an HTML 2.0 compliant web browser such as Netscape Navigator™.

Note – A web browser is not provided with the Backup product.

Choose Open File from the File menu in your web browser and type in the following path.

```
/cdrom/sbu_4_2/Manuals/html/index.html
```

The HTML index page is displayed. Click on the desired product name to view the documentation for that product.

For faster access to the Backup documentation in HTML format, copy the directory where the HTML files are stored to a directory in the local file system, as shown in the example.

```
# cp -r /cdrom/sbu_4_2/Manuals/html /local-path
```

Subsequently, access the HTML index page using the local file sytem path with your web browser instead of using the path to the files stored on the CD-ROM.

Installing the Server Software

Install the Backup software by using the `pkgadd` command. During the `pkgadd` installation the Backup programs are installed in the directories `/usr/bin/nsr`, `/usr/sbin/nsr`, and `/usr/lib/nsr`. The installation script asks you to provide the name of the backup device to be used for backing up and recovering Backup files.

Caution – If you are upgrading from Single Server to Backup, refer to the *Solstice Backup 4.2 Product and Installation Notes*, for detailed instructions. If you are updating from a previous version of Backup, you must remove the old version before installing the new software. The *Solstice Backup 4.2 Product and Installation Notes*, also provides these instructions.

Install the online man pages if you want additional information about Backup. The man pages require 679 KB of disk space.

Install the AnwerBook and HTML files containing the Backup documentation set for UNIX if you want to be able to view the documentation online.

Note – The software distribution CD-ROM contains SunOS and Solaris versions of Backup. The Solaris version is located in the *Solaris* directory on the CD-ROM.

Files Modified during Installation

The Backup installation script modifies certain files when installing Backup with pkgadd.

The table below lists files modified by Backup:

Backup Server	Modified Files
Solaris system	/etc/rpc /etc/syslog.conf

Note – Save a copy of the original files before you install Backup.

Running the Solaris pkgadd Server Installation

Use pkgadd to install the Backup software on your Solaris system.

1. **Insert the Backup distribution CD-ROM into the CD-ROM drive. If you do not have a CD-ROM attached to your system, refer to the instructions in "How to Install the Client over NFS" on page 53.**

2. **Become *root* on the server.**

```
% su root
Password:
#
```

3. **Mount the CD-ROM.**

Note – To install the server, you must choose *both* the client and server packages. If you want to install all of the available Backup software at one time, including the device driver and man pages, press the [Return] key for the default "all" when the server installation script asks you to select a package to install.

4. Enter the following command at the system prompt; assuming `cdrom/sbu_4_2` **is the** *mount-point* **of the CD-ROM device you are using.**

```
# pkgadd -d /cdrom/sbu_4_2/Solaris/
```

The parts of the installation script that prompt you for a response are shown. In the following example, all the packages are installed: the server, client, man pages, and device driver.

Note – If you enter "all" for the name of the first Backup server, the current server can be backed up by all other servers. If the names of one or more servers are entered (including the current server), the current server will only be backed up by those servers.

The following packages are available:

How to Install the Packages for the Man pages, AnswerBook, and Server

```
# pkgadd -d .
The following packages are available:
1  SUNWsbuc  Solstice Backup (Backup/Recover) Client (sparc) 4.2
2  SUNWsbum  Solstice Backup (Backup/Recover) Man (sparc) 4.2
3  SUNWsbus1 Solstice Backup (Backup/Recover) Server (sparc) 4.2
4  SUNWsbus2 Solstice Backup (Backup/Recover) Device
Drivers(sparc)4.2
Select package(s) you wish to process (or ‘all’ to process all
packages). (default: all) [?,??,q]: 1
```

How to Add the Man Pages

```
# cd /cdrom/sbu_4_2/Solaris/sparc
```

```
# pkgadd -d . SUNWsbum
```

Note – For x86 installations, replace “sparc” in the following code example with “x86”.

```
Processing package instance <SUNWsbum> from
</cdrom/sbu_4_2/Solaris/sparc>

Solstice Backup (Backup/Recover) Man
(sparc) 4.2
SunSoft, Inc.
Using </usr> as the package base directory.
## Processing package information.
## Processing system information.
## Verifying disk space requirements.
## Checking for conflicts with packages already installed.
## Checking for setuid/setgid programs.

This package contains scripts which will be executed with super-user
permission during the process of installing this package.

Do you want to continue with the installation of <SUNWsbum> [y,n,?] y

Installing Solstice Backup (Backup/Recover) Man as <SUNWsbum>
```

```
## Installing part 1 of 1.
/usr/lib/nsr/man_ize
/usr/share/man/man3/getdate.3
/usr/share/man/man5/mm_data.5
...
/usr/share/man/man8/tapeexercise.8
/usr/share/man/man8/uasm.8
[ verifying class <none> ]
## Executing postinstall script.

Backup(TM) - Release 4.2
Copyright (c) 1990-1995, Legato Systems, Inc.  All rights reserved.
This product includes software developed by the University of
California, Berkeley and its contributors.


Modifying /usr/man/man.cf

Installation of <SUNWsbum> was successful.
```

▼ How to Add the AnswerBook

♦ **To add the AnswerBook online documentation, type the following:**

```
# cd /cdrom/sbu_4_2/Solaris/sparc
# pkgadd -d . SUNWsbuAB
```

Note – For x86 installations, replace "sparc" in the following code example with "x86".

Viewing AnswerBook Online Documentation

To view the AnswerBook online documentation, type:

```
% /usr/openwin/bin/answerbook
```

▼ How to Add the Server Packages

```
# pkgadd -d . SUNWsbus1

Processing package instance <SUNWsbus1> from
</cdrom/sbu_4_2/Solaris/sparc>

Solstice Backup (Backup/Recover) Server
(sparc) 4.2
SunSoft, Inc.
Backup(TM) - Release 4.2
Copyright (c) 1990-1995, Legato Systems, Inc.  All rights reserved.
This product includes software developed by the University of
California, Berkeley and its contributors.

Enter the tape or disk device(s) that are going to be used by the
Backup server. Use the no-rewind name for each tape device If you do
not choose a device a default device will be created for you.  (i.e.,
use /dev/rmt/0mbn instead of /dev/rmt/0mb).


Enter device name ([Return] if no more): /dev/rmt/0bn
Select the device type for “/dev/rmt/0bn”.
The possible types are:
8mm 5GB- 8mm double density tape.
8mm - 8mm tape.
4mm- 4mm DAT tape.
himt- 1/2 inch magnetic tape.
3480- 1/2 inch cartridge tape.
qic- 1/4 inch data cartridge.
optical- Optical Disk.
dlt- Digital Linear Tape.
vhs- VHS Format Tape.

Select device type for /dev/rmt/0bn:
(one of “8mm 4mm himt 3480 qic optical dlt 8mm 5GB”) [8mm 5GB]?
[Return]


Enter device name ([Return] if no more):[Return]

Start Backup daemons at end of install [yes]? [Return]
Using </usr> as the package base directory.
```

```
## Processing package information.
## Processing system information.
## Verifying package dependencies.
## Verifying disk space requirements.
## Checking for conflicts with packages already installed.
## Checking for setuid/setgid programs.

The following files are being installed with setuid and/or setgid
permissions:
  /usr/sbin/nsr/savegrp <setuid root>

Do you want to install these as setuid/setgid files [y,n,?,q] y

This package contains scripts which will be executed with super-user
permission during the process of installing this package.

Do you want to continue with the installation of <SUNWsbus1> [y,n,?]
y

Installing Solstice Backup (Backup/Recover) Server as <SUNWsbus1>

## Installing part 1 of 1.
/usr/lib/nsr/nsrindexasm
/usr/lib/nsr/nsrmmdbasm
/usr/sbin/nsr/ansrd
...
/usr/sbin/nsr/scanner
/usr/sbin/nsr/tapeexercise
[ verifying class <none> ]
/usr/sbin/nsr/recoverindex <linked pathname>
## Executing postinstall script.
Modifying /etc/rpc
Modifying /etc/syslog.conf
Restarting syslog daemon
Completing Installation
Starting Backup daemons
Backup successfully installed on `waif'!

Installation of <SUNWsbus1> was successful.
```

Configuring the Clients to Recognize the Server

You must configure the Solaris clients on the network to recognize the Backup server. Configure the clients by defining an alias called “nsrhost” for your Backup server on each client.

For networks not using the Network Information Service (NIS) or Domain Name System (DNS), add an alias of “nsrhost” next to the name of the server machine in the `/etc/hosts` file for each client backing up to Backup.

This section also shows you how to use NIS to configure the clients to recognize the Backup server. If you are using DNS, refer to the appropriate documentation for a Solaris system.

Note – If you have purchased Backup for more than one server, identify only one server as the `nsrhost` alias *in the* `/etc/hosts` file for each client. The server with the `nsrhost` alias is considered the primary Backup server.

▼ How to Change the Configuration for Each Client

If you *do not* have NIS (or DNS) on the network, log on to each client as *root* and edit the `/etc/hosts` file, adding the name “nsrhost” next to the Backup server name. For example:

1. **Log on to the client as *root*.**

```
jupiter#
```

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

```
137.69.1.1                    venus
137.69.1.2                    jupiter
137.69.1.3                    mars nsrhost
137.69.1.4                    mercury
```

In this example, the client named *jupiter* is backed up by the Backup server named *mars*.

3. **Repeat for each client that backs up to the Backup server.**

The Backup clients can now access the server for backup.

▼ How to Change the Configuration Using NIS

If you have NIS on the network, follow these steps:

1. **Find your NIS master by using the ypwhich command:**

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

In this example, *venus* is the master NIS host.

2. **Log on to the master NIS host as *root*.**

3. **Edit the */etc/hosts* file, adding “nsrhost” as an alias next to the name of the Backup server.**

Note – If you have more than one Backup server, you can only use one server as the NIS master, usually your primary server.

```
137.69.1.1                    venus
137.69.1.2                    jupiter
137.69.1.3                    mars nsrhost
137.69.1.4                    mercury
```

4. **Change directory to `/var/yp` (or `/etc/yp`) and use the make command:**

```
venus# cd /var/yp
venus# make
```

The Backup server can now access the clients for backup.

For more information about how clients bind to Backup servers, refer to the **nsr** manual page.

Client Installation Requirements

If you have clients of the same hardware platform as the Backup server, use the same software distribution media to install Backup software on the clients. Instructions for installing the client software on like clients are provided in this section. For clients with different hardware platforms, you need to purchase and install the client software separately. For detailed information about our services, support policies, and software subscriptions, contact your local customer support office for programs and their availability.

In order for Backup clients to use the Backup server to back up and recover files, they must be able to access the Backup software. There are two ways a client can access the Backup software:

- Clients have NFS-mounted a directory where the Backup programs are located.
- Clients have the Backup programs installed on their local disks.

Caution – In order to back up a Backup client (version 4.0.2 or higher) over the network, the `nsrexecd` daemon must be running on the client. The `pkgadd` installation starts `nsrexecd`.

For more information about the access control policies for the Backup commands, refer to “Restricting Client Access” in Chapter 3 of the *Solstice Backup 4.2 Administration Guide* or the `nsr` man page.

If a client is not set up to share programs over the network, you must install the Backup software directly on that client. Use the same Backup software distribution media you used to install the software on the server.

Installing the Client Software

This section describes the software installation for Solaris clients that have the same software and hardware as the server.

Note – When the client installation script asks you to select a package to install, be sure you select the client software. Do not press the [Return] key for the default “all.”

Before you begin the installation, mount the CD-ROM in your usual manner.

When you complete the client software installation, remove the CD-ROM from the drive and store it in a safe place.

▼ Running the Solaris pkgadd Client Installation

```
# cd /cdrom/sbu_4_2/Solaris/sparc

# pkgadd -d . SUNWsbuc

Processing package instance <SUNWsbuc> from
</cdrom/sbu_4_2/Solaris/sparc>
Solstice Backup (Backup/Recover) Client (sparc) 4.2
SunSoft, Inc.
Backup(TM) - Release 4.2
Copyright (c) 1990-1995, Legato Systems, Inc.  All rights reserved.
This product includes software developed by the University of
California, Berkeley and its contributors.
To set up a Backup server, you need to supply a directory with enough
free space to maintain all the on-line save file indexing and media
management information.
To set up a Backup client, you need to supply a directory for the
nsrexecd state file.
Below is a list of some of the filesystems, with their free space,
which you might consider:
/                              :   53068
/usr                           :   29208
Directory to use for client and server information [/nsr]?
The nsrexecd program restricts access to a select set of Backup
servers.  Please enter the names of each computer running a Backup
server that will back up this computer, one name at a time.  If a
computer has more than one network interface, please enter each
interface's name (one at a time).
Enter the first Backup server's name [no more]: waif
Enter the next Backup server's name [no more]:
Start Backup daemons at end of install [yes]?
Using </usr> as the package base directory.
## Processing package information.
## Processing system information.
## Verifying disk space requirements.
## Checking for conflicts with packages already installed.
## Checking for setuid/setgid programs.
The following files are being installed with setuid and/or setgid
permissions:
/usr/bin/nsr/recover <setuid root>
/usr/bin/nsr/save <setuid root>
```

```
/usr/sbin/nsr/nsrarchive <setuid root>
/usr/sbin/nsr/nsrretrieve <setuid root>
Do you want to install these as setuid/setgid files [y,n,?,q] y
This package contains scripts which will be executed with super-user
permission during the process of installing this package.
Do you want to continue with the installation of <SUNWsbuc> [y,n,?] y
Installing Solstice Backup (Backup/Recover) Client as <SUNWsbuc>
## Installing part 1 of 1.
/usr/bin/nsr/Backup
/usr/bin/nsr/nsr.help
/usr/bin/nsr/nsrwatch
...
/usr/sbin/nsr/nsrretrieve
/usr/sbin/nsr/savefs
[ verifying class <none> ]
/usr/bin/nsr/nwadmin <linked pathname>
/usr/bin/nsr/nwbackup <linked pathname>
/usr/sbin/nsr/save <linked pathname>
## Executing postinstall script.
 Directory /nsr, does not exist.
 Creating directory /nsr.
 nsr-izing system files
 nsr-izing system files
 Creating /etc/init.d/Backup
 Creating /etc/rc2.d/S95Backup
 Creating /etc/rc2.d/K05Backup
 Completing Installation
 Starting Backup daemons
Backup successfully installed on `waif'!
Installation of <SUNWsbuc> was successful.
```

▼ How to Install the Client over NFS

To install Backup over NFS on a Solaris client follow these steps:

1. **Place the distribution CD-ROM in the drive on a Solaris server.**

2. **Log on as *root* on the server. At the system prompt enter:**

```
# pkgadd -s nfsdir -d /cdrom/Backup_4_2/Solaris/sparc
```

   where *nfsdir* is the name of a shared (exported) NFS directory. You copy Backup client software from the distribution CD-ROM to an NFS directory on the server.

3. **From the installation script select** SUNWsbuc **to copy just the Backup client software.**
   You may also select SUNWsbum to copy reference manual pages.

4. **Log on as *root* on the client where you wish to install the Backup client software.**
   Remote mount the NFS directory *nfsdir* from the server. At the system prompt enter:

```
# pkgadd -d nfsdir
```

5. **From the installation script select** SUNWsbuc **and** SUNWsbum **to install the Backup client software and the manual pages, if applicable.**

6. **Answer the remaining installation script questions.**
   See the section “Running the Solaris pkgadd Client Installation” for the appropriate responses. Finally, unmount *nfsdir.*

Once Backup is installed on both the server and the clients, refer to the *Solstice Backup 4.2 Administration Guide* for information on how to configure Backup for unattended backups. Refer to the *Solstice Backup 4.2 User Guide* to learn how to use the Backup graphical user interface to back up and recover files.

Completing a Remote Extraction Using NFS

If you do not have a CD-ROM drive on a machine where you would like to install Backup software, then the following steps will outline how to perform this over NFS.

In the following example `cdserver` is the name of the machine on which the Backup software CD-ROM is mounted, and `nwhost` is the machine that will be installing the software. In the following example it is also assumed that the CD-ROM is mounted as `/cdrom/sbu_4_2`.

1. **For a Solaris 2.x based cdserver add the following line to the** `/etc/dfs/dfstab` **file:**

```
# share -F nfs -o ro /cdrom/sbu_4_2
```

2. **On a Solaris 2.x cdserver run:**

```
 cdserver# shareall
```

On a Solaris 2.x nwclient add the following line to `/etc/vfstab`:

```
cdserver:/cdrom/sbu_4_2 - /remotecd nfs - no ro
```

3. **Create the mount point:**

```
nwclient# mkdir /remotecd
```

4. **Mount the remote CD-ROM.**

```
nwclient# mount /remotecd
```

Then in the following sections, any path containing `/cdrom/sbu_4_2` should be replaced with `/remotecd`.

How to Install the Optional Jukebox Software Module

This section contains instructions for installing the Jukebox Software Module after installing the server software as well as instructions for installing the server and device driver software simultaneously.

In the following sections, the term jukebox is used to refer to a variety of backup devices, including the following: autoloader, carousel, library, near-line storage, datawheel, and autochanger.

You may have deleted the autochanger software after installing the Backup server software. If you add the Jukebox Software Module to a Solaris system later, you may need to use *pkgadd* to install the software from the CD-ROM. There are four phases to installing the optional Backup Jukebox Software Module:

1. **Remove the old device driver, if necessary.**
2. **Install the device driver software.**
3. **Enable the Jukebox Software Module.**
4. **Register the Jukebox Software Module.**

Caution – Backup supports jukeboxes connected by SCSI or serial (RS-232) ports. If your jukebox is connected with SCSI, you must install a driver for the SCSI port.

If your jukebox is connected through a serial port, you do not need to install the driver. Simply skip the instructions in this guide regarding the driver installation. However, you do need to configure and connect your jukebox by following the hardware instructions which were shipped with your jukebox. You must also properly enable and register your Jukebox Software Module.

See Chapter 4, “Enabling and Registering Backup” for information about enabling and registering the jukebox module.

If you install additional jukeboxes later, you must enable and register each additional Jukebox Software Module you purchase. You only need to extract the Backup software to add an additional jukebox if you removed the software after the original installation.

Follow the instructions in the next sections to install, test, and configure the jukebox module on a Solaris Backup server.

Removing an Old Device Driver

If you already have a driver and are installing a new one, you need to remove the old device driver from your Backup server first. See the section "Removing the Driver Software" on page 66 for detailed instructions.

Installing the Device Driver Software

The following is an overview of the procedure to complete the device driver software installation on a Solaris system:

- log in as *root*
- use *pkgadd* to install the device driver software

Follow these steps to install the device driver software:

1. **Log in as *root* at the shell prompt:**

```
% su
password:
#
```

2. **To install the software, run** pkgadd **with the system *pathname* of the directory on the CD-ROM.**

```
# cd /cdrom/sbu_4_2/solaris/sparc
# pkgadd  -d . SUNWsbus2
```

Enter your responses to the following questions provided by the installation script. Press [Return] after each response to accept the default values listed in the brackets.

▼ How to Add the Device Drivers

```
# cd /cdrom/sbu_4_2/Solaris/sparc
# pkgadd -d . SUNWsbus2

Processing package instance <SUNWsbus2> from
</cdrom/sbu_4_2/Solaris/sparc>
```

```
Solstice Backup (Backup/Recover) Device Drivers
(sparc) 4.2
# @(#)copyright.sh 1.1 95/09/10 Copyright (c) 1995, Legato Systems,
Inc.

Copyright (c) 1990-1995, Legato Systems, Inc.
All Rights Reserved.
Using </> as the package base directory.
## Processing package information.
## Processing system information.
## Verifying package dependencies.
## Verifying disk space requirements.
## Checking for conflicts with packages already installed.
## Checking for setuid/setgid programs.

This package contains scripts which will be executed with super-user
permission during the process of installing this package.

Do you want to continue with the installation of <SUNWsbus2> [y,n,?]
y

Installing Solstice Backup (Backup/Recover) Device Drivers as
<SUNWsbus2>

## Executing preinstall script.
## Installing part 1 of 1.
/etc/LGTOuscsi/changers
/etc/LGTOuscsi/hpflip
/etc/LGTOuscsi/ielem
...
/usr/kernel/drv/lus
/usr/kernel/drv/lus.conf
[ verifying class <none> ]
## Executing postinstall script.

Installation of <SUNWsbus2> was successful.
```

Caution – Do *not* relocate the driver package (<*SUNWsbus2*>). You may relocate the other packages, but not the device driver.

How to Test the Device Driver Installation

After the driver loads successfully, test the driver installation.

Use the following procedure to verify that the driver loaded properly.

1. **Enter the following command:**

```
# /etc/LGTOuscsi/lusdebug 1
```

You should see the following response:

```
debug level was 0; is now 1
```

2. **Enter the following command:**

```
# /etc/LGTOuscsi/lusdebug 0
```

You should see the following response:

```
debug level was 1; is now 0
```

3. **Enter the following command:**

```
# /etc/LGTOuscsi/inquire
```

A list of SCSI devices attached to your server, if any, appears on your screen. Your list will probably look different than this one. If you attached your jukebox before installing the driver, your jukebox should appear in the list.

```
scsidev@0.0.0:FUJITSU M2263S-51201                          |
Direct Access
scsidev@0.4.0:Quantum DLT4700                               |
Sequential Access
scsidev@0.4.1:Quantum TZ Media Changer                      |
Changer Device
```

▼ How to Configure a Jukebox

To configure a jukebox you must run the `jb_config` program.

1. **Become *root* on the Backup server.**

```
% su root
Password:
#
```

2. **Enter the jb_config command at the prompt:**

```
# jb_config
```

3. **Backup displays the installation script and asks for your choices. The following example shows the responses for a SCSI Jukebox.**

```
1) Install an Autodetected SCSI Jukebox.
2) Install a Serial Jukebox.
3) Install an SJI Jukebox.

What kind of Jukebox are you installing? 1
```

```
These are the SCSI Jukeboxes currently attached to your system:
  1) scsidev@1.2.0: DLI Libra Series
  2) scsidev@0.2.1: Quantum DLT/Digital DLT
Which one do you want to install? 2
Installing an 'Quantum DLT/Digital DLT' jukebox.
Name you would like to assign to the jukebox device? dlt
Pathname of the control port for the jukebox device? [scsidev@0.2.1]
[Return]
Do you want automated device cleaning support enabled? (yes/no) n
Enter pathname of media drive 1 [/dev/nrst8]: ? [Return]
This media device has not been configured yet.  Please
select a media device type for /dev/nrst8.
a) himt
b) qic
c) 4mm
d) 8mm
e) 8mm 5GB
f) 3480
g) dlt
h) vhs
i)optical
Choice? g

Jukebox has been added successfully
```

If you enter **3**, “Install an SJI Jukebox,” at the first question, you see the following response:

```
What kind of Jukebox are you installing? 3
Enter the number corresponding to the type of jukebox
you are installing:
1) ADIC-1200c/ADIC-1200d     2) ADIC-VLS
3) ARC-DiamondBack           4) Breece Hill
5) DLI Libra Series          6) Quantum DLT/Digital DLT
7) EXB-10e/EXB-10h           8) EXB-10i
9) EXB-60                    10) EXB-120
11) EXB-210                  12) EXB-218
13) EXB-400 Series           14) HP-C1553A/Surestore 12000e
15) Metrum (SCSI)            16) Qualstar
17) Spectralogic             18) STK-9704/Lago 340
19) STK-9708/Lago 380(SCSI)  20) IBM 7331/IBM 9427
21) ATL/Odetics SCSI         22) HP-Optical 630MB/1.3GB
23) other
Choice? 6
Installing an 'Quantum DLT/Digital DLT' jukebox.
Name you would like to assign to the jukebox device? dlt
```

```
Pathname of the control port for the jukebox device? scsidev@1.2.0
[Return]
Do you want automated device cleaning support enabled? (yes/no) n
Enter pathname of media drive 1 [/dev/nrst8]: ? [Return]
This media device has not been configured yet.  Please
select a media device type for /dev/nrst8.
a) himt
b) qic
c) 4mm
d) 8mm
e) 8mm 5GB
f) 3480
g) dlt
h) vhs
i) optical
Choice? c

Jukebox has been added successfully.
```

Testing the Jukebox Connection

To test the jukebox connection, run the `jbexercise` program with two pieces of “scratch” media loaded in the first and last slots of the jukebox. The drives must be empty and have their doors open, if they have any.

When running `jbexercise`, specify the control port and the device type. The control port for SCSI jukebox models is typically */dev/scsidev@x.x.x*. The exact control port pathname is shown when running `jb_config`. For example, the following command runs the `jbexercise` program on the DLT jukebox:

```
# jbexercise -c /dev/scsidev@1.2.0 -m Quantum DLT/Digital DLT
```

Refer to the `jbexercise` man page for additional information.

Enabling the Jukebox Software Module

After installing the device driver, enable and register the autochanger module by following the instructions in Chapter 4, “Enabling and Registering Backup.”

Starting the Backup Administrator Program

By installing Backup, you have configured the server to provide backup and recover services. To bring up the Backup Administrator window, use the nwadmin command at the system prompt:

```
# nwadmin &
```

If you have trouble starting Backup:

- The Backup daemons may not be running properly.
- The DISPLAY variable on the machine you are using may not be set correctly.

See Appendix B, “Troubleshooting,” for information on checking and starting the Backup daemons and on setting the DISPLAY variable correctly.

Removing Solaris Backup Software

This section describes how to remove the server, client, and driver software from your Solaris system.

▼ Removing the Server Software

Use pkgrm to remove the server software on your Solaris system.

To remove the server software, use the command below:

```
# pkgrm SUNWsbus1
```

```
The following package is currently installed:
   SUNWsbus1       Solstice Backup (Backup/Recover) Server
                   (sparc) 4.2

Do you want to remove this package? y

## Removing installed package instance <SUNWsbus1>

This package contains scripts which will be executed with super-user
permission during the process of removing this package.
```

```
Do you want to continue with the removal of this package [y,n,?,q] y
## Verifying package dependencies.
## Processing package information.
## Executing preremove script.
de nsr-izing system files
Modifying /etc/rpc
Modifying /etc/syslog.conf
Restarting syslog daemon

To completely remove Backup from the system
you must remove the /nsr directory. /nsr
may be a symbolic link, if so you will also need to
remove the directory that it points to.

Warning: the /nsr directory contains the online
client file index files and the server media index file.
You should only remove it if you want to completely
remove Backup from the system.
Backup successfully removed for upgrade from `waif'.
## Removing pathnames in class <none>
/usr/sbin/nsr/tapeexercise
/usr/sbin/nsr/scanner
/usr/sbin/nsr/savegrp
...
/usr/lib/nsr/nsrmmdbasm
/usr/lib/nsr/nsrindexasm
## Updating system information.

Removal of <SUNWsbus1> was successful.
```

▼ Removing the Client Software

To remove the client software use the command below:

```
# pkgrm SUNWsbuc
The following package is currently installed:
   SUNWsbuc          Solstice Backup (Backup/Recover) Client
                     (sparc) 4.2

Do you want to remove this package? y

## Removing installed package instance <SUNWsbuc>

This package contains scripts which will be executed with super-
user permission during the process of removing this package.

Do you want to continue with the removal of this package [y,n,?,q]
y
## Verifying package dependencies.
## Processing package information.
## Executing preremove script.
    Removing /etc/init.d/Backup
    Removing /etc/rc2.d/S95Backup
    Removing /etc/rc2.d/K05Backup
Backup successfully removed from `waif'.
## Removing pathnames in class <none>
/usr/sbin/nsr/savefs
/usr/sbin/nsr/save
/usr/sbin/nsr/nsrretrieve
...
/usr/bin/nsr/nsr.help
/usr/bin/nsr/Backup
## Updating system information.

Removal of <SUNWsbuc> was successful.
```

Caution – If both the server and the client software are installed, you must remove the server software before removing the client software.

Removing the Driver Software

To remove the driver software use the following command:

```
# pkgrm SUNWsbus2
The following package is currently installed:
   SUNWsbus2      Solstice Backup (Backup/Recover) Device Drivers
                    (sparc) 4.2

Do you want to remove this package? y

## Removing installed package instance <SUNWsbus2>

This package contains scripts which will be executed with super-
user
permission during the process of removing this package.

Do you want to continue with the removal of this package [y,n,?,q]
y
## Verifying package dependencies.
## Processing package information.
## Executing preremove script.
## Removing pathnames in class <none>
/usr/kernel/drv/lus.conf
/usr/kernel/drv/lus
/etc/LGTOuscsi/writebuf
/etc/LGTOuscsi/tur
...
/etc/LGTOuscsi/hpflip
/etc/LGTOuscsi/changers
## Updating system information.


Removal of <SUNWsbus2> was successful.
```

Removing the Man Pages

To remove the man pages use the command below:

```
# pkgrm SUNWsbum
```

```
The following package is currently installed:
   SUNWsbum        Solstice Backup (Backup/Recover) Man
                   (sparc) 4.2

Do you want to remove this package? y

## Removing installed package instance <SUNWsbum>

This package contains scripts which will be executed with super-user
permission during the process of removing this package.

Do you want to continue with the removal of this package [y,n,?,q] y
## Verifying package dependencies.
## Processing package information.
## Executing preremove script.

Backup(TM) - Release 4.2
Copyright (c) 1990-1995, Legato Systems, Inc.  All rights reserved.
This product includes software developed by the University of
California, Berkeley and its contributors.


Modifying /usr/man/man.cf
## Removing pathnames in class <none>
/usr/share/man/man8/uasm.8
/usr/share/man/man8/tapeexercise.8
/usr/share/man/man8/scanner.8
...
/usr/share/man/man3/getdate.3
/usr/lib/nsr/man_ize
## Updating system information.


Removal of <SUNWsbum> was successful.
```

Enabling and Registering Backup 4

After you complete the installation process, read this chapter to properly enable and register your Backup products. If you enable a Backup product and do not properly register the software, it automatically disables itself 45 days after the date you enabled the software. In other words, you have 45 days to register your products after you enable them.

Following is an overview of the enabling and registering process:

1. **Install Backup and any additional client connections or optional modules you purchased.**

2. **Enter required company and product information in the Server window, including the product serial number.**
   This step is only necessary for the server products.

Note – Although Sun does not assign serial numbers, this field is required. Enter the number 1 in this field.

3. **Enable all installed Backup products, including optional modules, using the nsrcap command.**

4. **Locate, fill in, and email the registration information to license@Sun.Com. Alternatively, you may fax the Product Registration Form to the Sun Licensing Center at 1-317-364-7220. SunSoft will then send you authorization code(s) to permanently enable your Solstice product(s). If you have questions about the authorization code process, call the Sun**

Licensing Center at 1-800-872-4786, or 1-317-364-7216. Callers from Canada may call at 1-800-722-4786. Authorization codes for this product will only be issued from the U.S. Sun Licensing Center.

▼ How to Provide Company and Product Information

You only need to follow the instructions in this section if you are registering and enabling a Backup server product.

Fill in your company and product information in the Server window to complete the registration process. The information you supply in the Server window is automatically added to the registration form printed from the Registration window.

Follow these steps to enter your company and product information:

1. **Open the Server window by selecting the Server Setup command from the Server menu.**

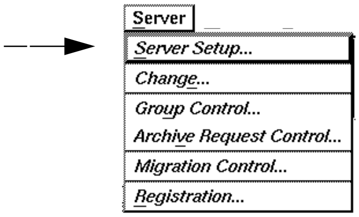

2. **Use the scroll bar on the right side of the window to view the fields where you enter your company and product information.**

3. **Enter the information in the appropriate fields. You must fill in the User name, Company, Phone, Purchase date, and Product serial number fields. If you do not fill in all five fields, you will not be able to print the registration form from the Registration window.**
   Your Server window will look similar to the following.

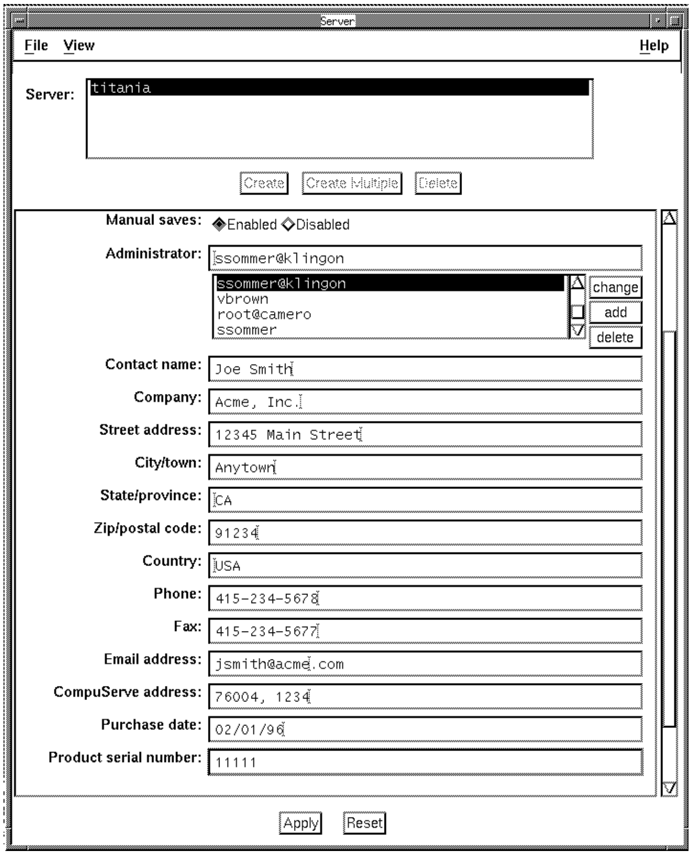

Server: titania

Manual saves: ◆ Enabled ◇ Disabled

Administrator: ssommer@klingon

ssommer@klingon
vbrown
root@camero
ssommer

Contact name: Joe Smith

Company: Acme, Inc.

Street address: 12345 Main Street

City/town: Anytown

State/province: CA

Zip/postal code: 91234

Country: USA

Phone: 415-234-5678

Fax: 415-234-5677

Email address: jsmith@acme.com

CompuServe address: 76004, 1234

Purchase date: 02/01/96

Product serial number: 11111

4. **Click the Apply button, once you have completed entering all of the necessary information.**

After completing this section you are ready to proceed with enabling your Backup products.

▼ How to Enable Backup Products

After you have entered the appropriate information in the Server window, follow the instructions in this section to properly enable your Backup products.

An *enabler* is an 18-digit code, printed on a paper certificate, and shipped with each Backup product, which you enter either using the `nsrcap` command at a system prompt or the Registration window from the Server menu. Once enabled, register your Backup product(s) as soon as possible. If you do not register your Backup product(s), they will "time-out" in 45 days and you will not be able to use Backup to back up your data.

The following steps show how to enable products from the command line:

1. **Find the Backup enabler code in your package.**

2. **Become *root* on the Backup server.**

```
% su root
Password:
#
```

3. **Use the `nsrcap` command at the system prompt to enter the enabler code.**

```
# nsrcap -v -c xxxxxx-xxxxxx-xxxxxx
```

substituting *xxxxxx* with your 18-digit code.

Note – If you have the Registration window open at the time you enable Backup, close it to refresh the window.

Once you have enabled a product, its enabler code will appear in the Registration window.

Enabling Jukeboxes

If you have more than one jukebox connected to the Backup server, and if you use the command line (**nsrcap** command) to enable the jukebox, Backup prompts you for the correct model.

If you enable the jukebox using the Registration window, Backup automatically enables the correct model without prompting you.

▼ How to Register Backup Products

After you have completed enabling your Backup productsl locate, fill in, and email the registration information to license@Sun.Com. Alternatively, you may fax the Product Registration Form to the Sun Licensing Center at 1-317-364-7220. SunSoft will then send you authorization code(s) to permanently enable your Solstice product(s). If you have questions about the authorization code process, call the Sun Licensing Center at 1-800-872-4786, or 1-317-364-7216. Callers from Canada may call at 1-800-722-4786. Authorization codes for this product will only be issued from the U.S. Sun Licensing Center.

To FAX or e-mail a copy of the output of the Registration window, follow these steps:

1. **Open the Registration window by selecting the Registration command in the Server menu.**

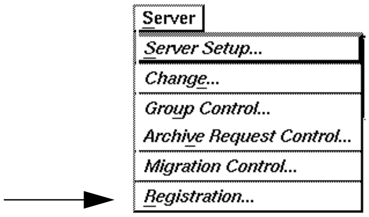

The Registration window is shown below. This is how it looks after you install and enable Backup. Notice there is an Expiration date and the Auth code field is empty (until you authorize the product).

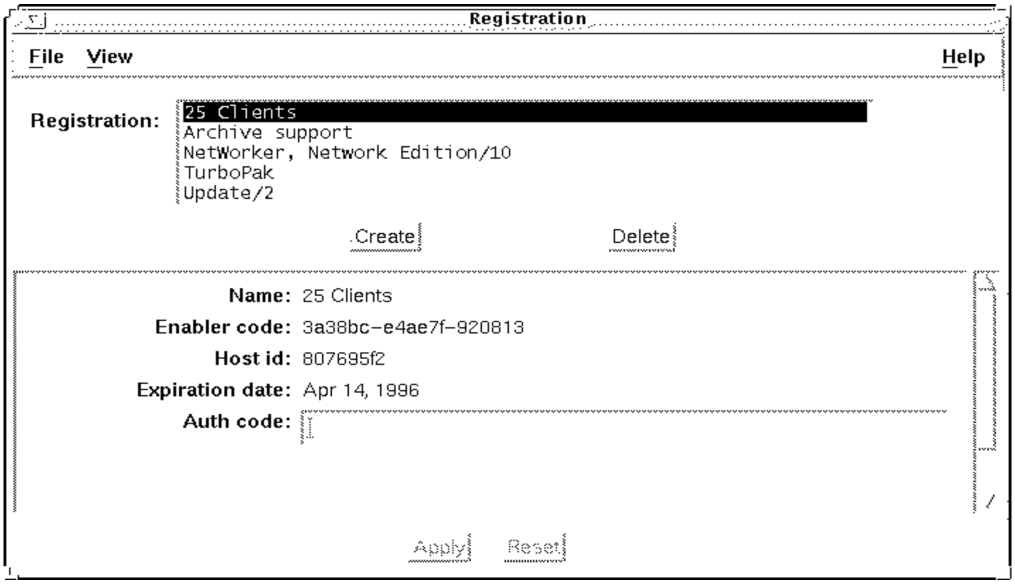

Registration:
25 Clients
Archive support
NetWorker, Network Edition/10
TurboPak
Update/2

Name: 25 Clients
Enabler code: 3a38bc-e4ae7f-920813
Host id: 807695f2
Expiration date: Apr 14, 1996
Auth code:

The Registration window lists the name of each enabled module, the unique enabler code, host id of the server, an expiration date, and an authorization code. The Auth code field is blank until you enter the authorization code sent to you by SunSoft.

When you select an item in the Registration scrolling list, the bottom half of the window changes to display the information pertaining to that item.

2. **Select Tabular from the View menu to display the tabular view of the window.**

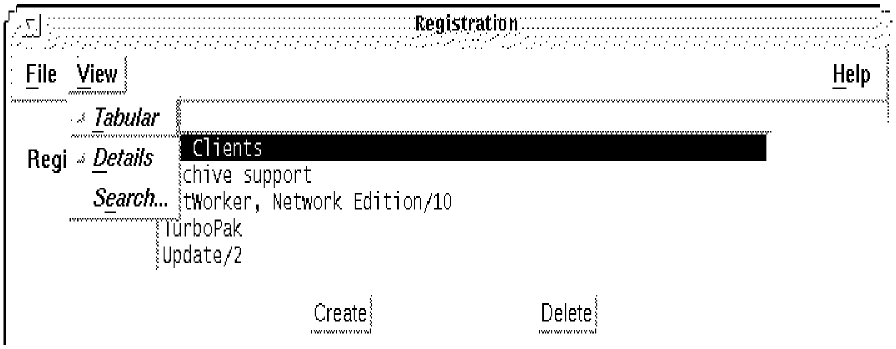

The Registration window changes to a tabular view, as shown.

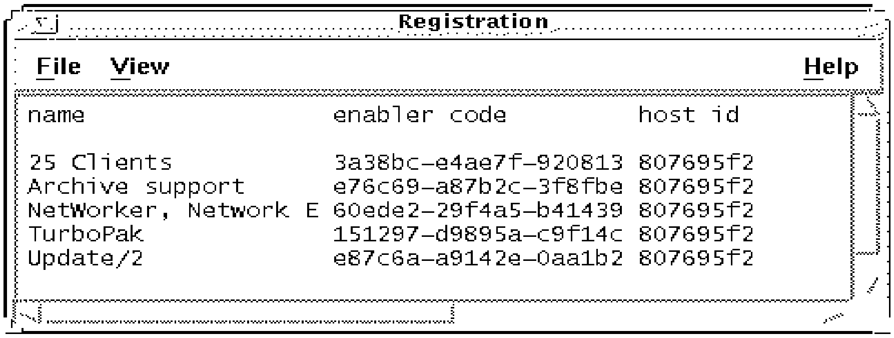

```
name                        enabler code           host id

25 Clients                  3a38bc-e4ae7f-920813 807695f2
Archive support             e76c69-a87b2c-3f8fbe 807695f2
NetWorker, Network E 60ede2-29f4a5-b41439 807695f2
TurboPak                    151297-d9895a-c9f14c 807695f2
Update/2                    e87c6a-a9142e-0aa1b2 807695f2
```

3. **Select Print from the File menu in the Tabular view of the Registration window, or copy and paste the information from the window into an e-mail message.**

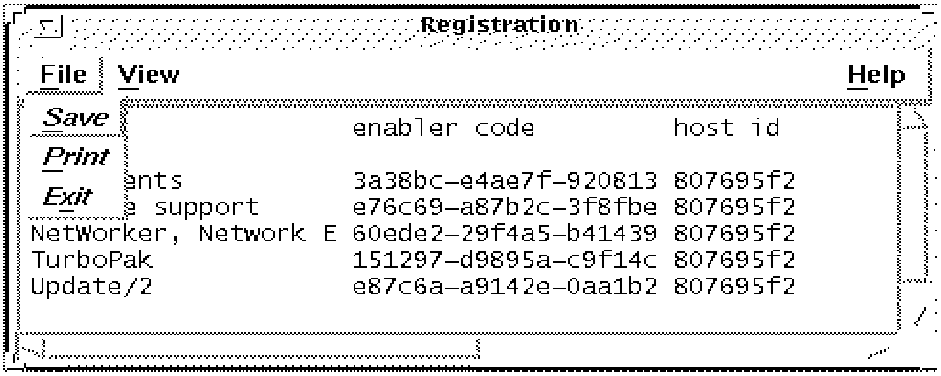

```
                           enabler code         host id
ents                       3a38bc-e4ae7f-920813 807695f2
e support                  e76c69-a87b2c-3f8fbe 807695f2
NetWorker, Network E       60ede2-29f4a5-b41439 807695f2
TurboPak                   151297-d9895a-c9f14c 807695f2
Update/2                   e87c6a-a9142e-0aa1b2 807695f2
```

Backup displays the Print Options dialog box, as shown.

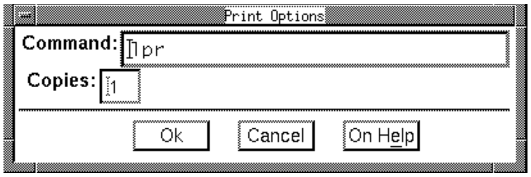

Command: lpr

Copies: 1

4. **Enter the appropriate UNIX command in the Command field to send a copy of the Registration information to a printer. Typically, this command is lpr or lp. Click the [Ok] button.**

Caution – If you did not previously enter the required information in the Server window, you will not be able to print the registration form.

5. **Select Normal from the View menu to restore the Registration window to its original display.**

6. **After you have completed enabling your Backup productsl locate, fill in, and email the registration information to license@Sun.Com. Alternatively, you may fax the Product Registration Form to the Sun Licensing Center at 1-317-364-7220. SunSoft will then send you authorization code(s) to permanently enable your Solstice product(s). If you have questions about the authorization code process, call the Sun**

Licensing Center at 1-800-872-4786, or 1-317-364-7216. Callers from Canada may call at 1-800-722-4786. Authorization codes for this product will only be issued from the U.S. Sun Licensing Center.

SunSoft will register the software and send you an authorization code for each Backup product you have purchased.

▼ How to Authorize Backup Products

Once you receive the authorization codes from SunSoft, follow these steps to authorize each Backup product:

1. **Open the Registration window by selecting the Registration command in the Server menu.**

2. **Highlight an item in the Registration scrolling list.**

3. **Enter the authorization code for that item into the Auth code field. Enter the number exactly as provided by SunSoft.**

4. **Click the Apply button.**
   The item is now registered. You see the message “Authorized – No expiration date” in the Expiration date field.

   The Registration window looks similar to this:

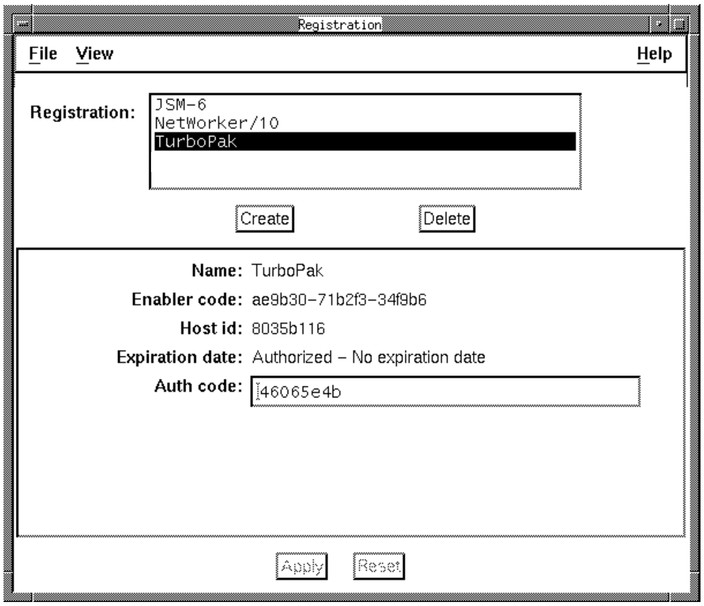

Registration: JSM-6
NetWorker/10
TurboPak

Name: TurboPak

Enabler code: ae9b30-71b2f3-34f9b6

Host id: 8035b116

Expiration date: Authorized – No expiration date

Auth code: 46065e4b

Repeat Step 2 through Step 4 for each Backup product that you installed.

Caution – Backup products automatically disable if you do not enter an authorization code for each one within 45 days from the date they were enabled.

You are now ready to take Backup for a “test drive.” Refer to Chapter 1, “Introduction,” and to Chapter 2, “Getting Started,” in the *Solstice Backup 4.2 Administration Guide* to begin using Backup.

Recovering from a Disk Crash 5

This chapter contains procedures to follow to prepare for and recover from a major disk crash. The four different types of disk crash situations in this chapter are described below.

- The primary disk, which contains the operating system and Backup binaries, is damaged. This can apply to a client system or a Backup server.
- The secondary disk, which contains other filesystems, is damaged. This can apply to a client system or a Backup server.
- The Backup server disk, which contains the online indexes (the */nsr* filesystem), is damaged. You have to recover the indexes before using Backup to recover any filesystems.
- The Backup fileserver was destroyed. You have to recover everything to a new Backup server.

If a primary disk suffers a head crash, you may need to replace the disk, boot from mini-root, format and partition the disk, re-install the operating system and Backup binaries, and then recover affected filesystems. In this case, before using Backup to recover the data to the disk, you should consult the system administration manuals you used to set up your fileserver for the first time.

If a secondary disk suffers a head crash, its recovery procedure is simpler, since you do not have to re-install the operating system and Backup binaries.

Preparing for a Crash

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

If the Backup server filesystem or disk that contains `/nsr` and /or any contents of `/nsr` linked to another location is destroyed, the recovery procedure involves an extra step – you must recover the online indexes for the server as well as the server filesystems. The `/nsr` directory on the server contains one index for each client, including an index for the server as a client of itself.

If your Backup server was destroyed (in a fire, for example) you need to replace it with another machine. You may do this as long as you do the following:

- name the replacement server with the same *hostname* as the original Backup server, if possible
- re-install Backup using the same directory locations for the online indexes as in the original installation
- re-register the new Backup server

Caution – Once you understand the procedure for a disaster recovery, make sure you have carefully thought of a disaster recovery plan for your site. If possible, you should test the ability to recover from a disaster at your site.

If you set up your network and enabled Backup for scheduled network-wide backups, you are well-prepared for a disaster. Every time Backup backs up a group of clients, it also backs up all the online indexes for those clients. This backup includes the file index, media database, and Backup configuration files for the server itself (which contain entries for the client indexes, should these require recovery). The media database, Backup configuration files, and part of the server index are saved to a special save set named *bootstrap*. The save set identification numbers (ssid) for all recent bootstraps are sent to a default printer, providing a hard copy for your records.

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

- Keep a file containing hard copies of the bootstrap records. Place these daily sheets of paper in a three-ring binder or a file folder.

- Make a hard copy record of the disks, partition sizes, and mount points for the server and any clients that have a local hard disk. This information makes a future recovery much smoother for you.
- Save your enabler code certificate and the authorization code for the Backup server.

File the Bootstrap Information

Backup sends a record of the bootstrap save sets to your default printer, so you have a piece of paper with the dates, locations, and save set id numbers needed for disaster recovery.

If you ever need to recover the server online indexes, the information on this piece of paper can save you a great deal of time. Save this information in a safe place.

The information sent to the printer looks similar to this:

```
August 20 03:30 1995 Backup bootstrap information Page 1

  date     time    level        ssid  file  record       volume
8/19/95  2:29:08       9  1148868949    56       0     mars.005
8/20/95  2:52:25       9  1148868985    77       0     mars.001
```

Backup prints all the bootstrap save sets for the past month. The bootstrap save set may span more than one backup volume. The file and record numbers are used to find the associated save set quickly.

You can also manually back up the Backup server indexes by using the `saveindex` command. Using this command also sends the bootstrap information to a printer. For example:

```
# saveindex -c server_name
```

To use the `saveindex` command you must be *root* on the Backup server.

File the Disk Information

Use the disk information command (df) to find out how the Backup server disks are partitioned and mounted. Use the appropriate operating system command to print disk partitioning information. For example, on a SunOS system use the df and dkinfo commands. Do the same for any Backup clients that have local hard disks. You should do this for each system Backup backs up, unless the systems are consistent in disk and filesystem layout. Print and file this information in case you ever have to recover from a disk crash.

For example, the df information looks similar to this:

```
% df
Filesystem   kbytes    used     avail    capacity Mounted on
/dev/sd0a    15087     7853     5725     58%        /
/dev/sd0e    361474    281710   43616    87%        /usr
/dev/sd1c    559860    362486   141388   72%        /usr/src
/dev/sd1a    818627    729462   7302     99%        /export
/dev/sd0d    28181     14747    10615    58%        /var
/dev/sd0c    559860    432409   71465    86%        /home/mars
/dev/sd2c    624263    513089   48747    91%        /nsr
```

The following dkinfo command examples give you information about how each disk is partitioned for a SunOS 4.x system:

```
% dkinfo sd0a
    SCSI CCS controller at addr f8800000, unit # 24
    1151 cylinders 9 heads 80 sectors/track
    33120 sectors (46 cyls)
    starting cylinder 0
% dkinfo sd0b
    1151 cylinders 9 heads 80 sectors/track
    197280 sectors (274 cyls)
    starting cylinder 46
```

The following `prtvtoc` command example gives you information about how each disk is partitioned for a Solaris 2.x system. The device name is the “raw” device corresponding to the device name used for the output from the `df` command.

```
# prtvtoc /dev/rdsk/c0t3ds0
```

```
                          First     Sector    Last      Mount
Partition  Tag  Flags     Sector    Count     Sector    Directory
0          0    00        0         410400    410399    /
1          0    00        410400    410400    820799
2          0    00        0         3116400   3116399
3          0    00        820800    1024800   1845599   /opt
6          0    00        1845600   1270800   3116399   /usr
```

If a disk is destroyed in a head crash, you will be able to restore it and recover the filesystems to their original state, using the hard copy information from these disk information commands. At a minimum, you need to have partitions large enough to hold all the recovered data.

Using the Save Set Recover Feature or the Recover Procedure

You have a choice between using the save set recover feature or the normal recovery procedure to recover filesystems after disk failure. This section describes the advantages of each method.

During backup, Backup multiplexes different filesystems simultaneously to the backup media. By recovering multiple filesystems at the same time, Backup has the opportunity to de-multiplex the filesystem save sets from the same backup volume in parallel, thus only reading each backup volume once. This is why you should try to recover multiple filesystems at the same time if it is practical for you to do so. You can do this by marking different save set points (using the save set recover feature) or different mount points (using the normal recover procedure).

Note – With release 4.1.2 or higher, the Backup server can recover several save sets from the same backup volume in parallel, eliminating the need to read the same backup volume several times during a recovery.

An advantage to using the save set recover feature is that you spend less time browsing and marking filesystems. With normal recovery, an entry for each individual file is accessed in the Backup file index to reconstruct an accurate view of the filesystem. It takes time to "pick" the most recent versions of files from the tape. With save set recover, individual file browsing is bypassed and entire save sets are recovered in one step. If the browse policy has expired, save set recover is the only way to recover a file using the GUI. Refer to Chapter 6, "Recovering and Cloning Save Sets," in the *Solstice Backup 4.2 Administration Guide.*

Caution – Any time you have to recover the primary disk (for example, *root*), you should do so in single-user mode from the system console, not from the X window system. Before starting this procedure, make sure all your filesystems are mounted.

There are two ways to use the save set recover feature:

- Run the nwadmin program and choose the Recover command from the Save Set menu. This opens the Save Set Recover window.
- Run `recover` with the `-S` *ssid* option from the system prompt. This is the way to recover if you are not using the X window system. Refer to the `mminfo` man page for instructions on how to find the save set(s) you want to recover. Refer to the `recover` man page for detailed instructions on using the `recover -S` command.

To use the Save Set Recover window for recovering an entire filesystem, you need to pay attention to the backup levels of the save sets you are marking for a recovery. For each save set you wish to recover, you need the last full backup followed by the most recent level backups of the save set to bring the system back to the state it was in before the system crash.

A level 0 backup is a full backup. Other levels after a full are represented by ascending numbers, or by the letter i for incremental saves. In other words, a level 3 backs up more data than a level 5. The following illustration shows you several backups over time. The bars represent the level backups since the last full backup.

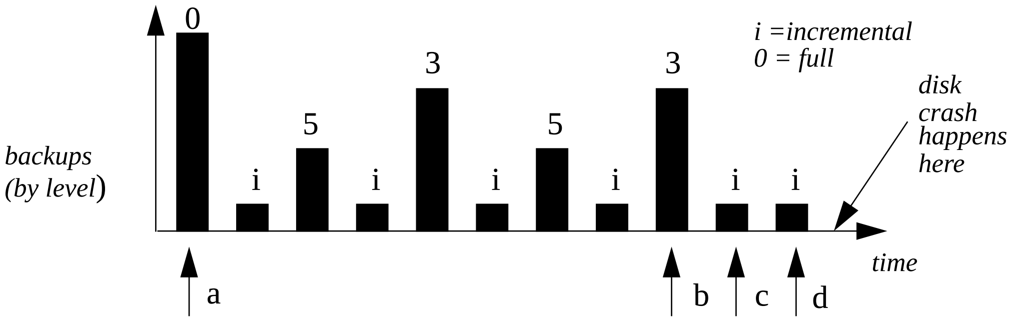

The arrows in the illustration point to the save sets you would need in order to recover your filesystems from the disk crash.

Note – By using save set recover, you may recover files that were deleted between backups. For example, if file F existed at time a, but was deleted prior to time b, file F will be recovered. This may require more disk space than is available if large files or a significant number of deleted files are recovered. Normal recover does not recover the deleted files, and the filesystem will be restored exactly as it was at time d with no superfluous files recovered.

For example, suppose you wish to recover the save set /usr shown in the Save Set Recover window example.

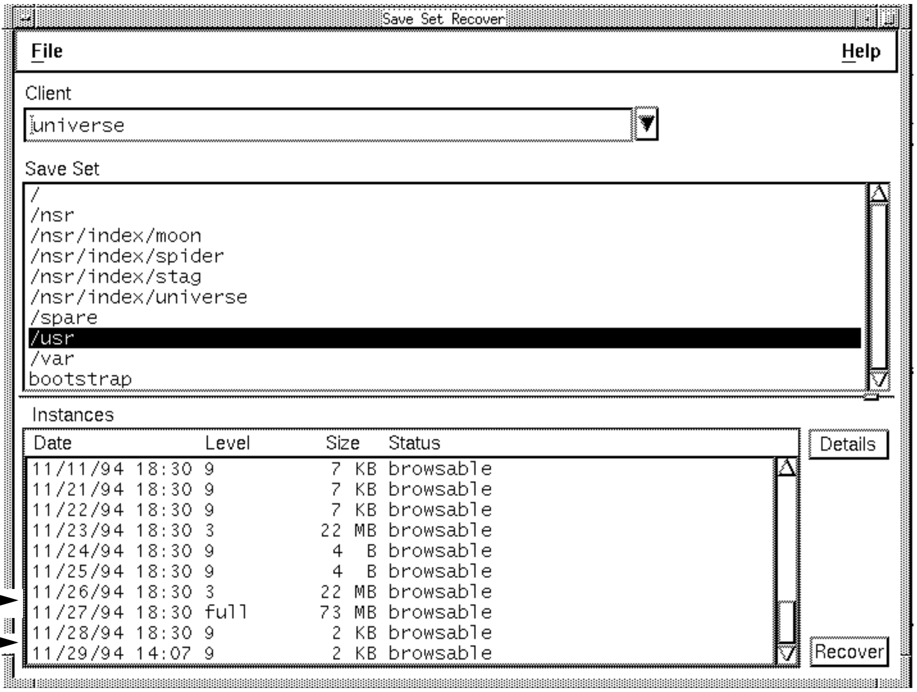

Client

universe

Save Set

```
/
/nsr
/nsr/index/moon
/nsr/index/spider
/nsr/index/stag
/nsr/index/universe
/spare
/usr
/var
bootstrap
```

Instances

Date	Level	Size	Status
11/11/94 18:30	9	7 KB	browsable
11/21/94 18:30	9	7 KB	browsable
11/22/94 18:30	9	7 KB	browsable
11/23/94 18:30	3	22 MB	browsable
11/24/94 18:30	9	4 B	browsable
11/25/94 18:30	9	4 B	browsable
11/26/94 18:30	3	22 MB	browsable
11/27/94 18:30	full	73 MB	browsable
11/28/94 18:30	9	2 KB	browsable
11/29/94 14:07	9	2 KB	browsable

These are the backups you need

The Instances display shows the backup history of /usr. To recover /usr, you need to mark the most recent full backup of /usr and the most recent level 9 to restore /usr to the state it was in before the system crash. In other words, you need to mark the appropriate save set levels for each filesystem you are recovering. As indicated in the previous note, you may recover files you were not expecting and do not need. Just delete these files.

Normally during disaster recovery, you want to force the recover program to overwrite existing files. Using save set recover, this is even more important since the same file may be recovered multiple times with each successive version coming from a later save set. Since each file in each save set is recovered with save set recover, files or directories previously deleted or renamed (with the mv command) between backups are still recovered.

These files or directories need to be manually deleted. You may even run out of space during the recovery if there are too many instances of previously deleted files or directories recovered by save set recover.

The save set recover feature reads each save set in its entirety during recovery. If you have to recover many save sets, you may be better off using the "normal" recover instead of save set recover. However, if you only need the last full backup of a save set, save set recover is a better method to use.

By contrast, if you decide to use the Recover command from the `nwrecover` program, you will recover your filesystems exactly as they were as of the last backup before the crash, and Backup will read the minimum amount of tape to recover the files. This method has these disadvantages:

- You spend more time browsing and marking files for recovery, since Backup needs to find an entry for each file in the index to add all the files in one filesystem.
- You may run out of swap space if you add too many filesystems at once to your list of data to recover.

In summary, we recommend the save set recover feature for faster disaster recovery:

- if you can determine the correct save sets to recover
- if there are only a few save sets to recover for each filesystem
- if you are not recovering the Backup indexes
- if recovering extra files is acceptable (not an issue if you are recovering from full backups only)

Use the normal recovery procedure:

- if you cannot determine which save sets to recover
- if you need to recover files from many backups to restore the filesystem to an acceptable level
- if you are recovering the Backup indexes
- if recovering extra files is not acceptable, and you are recovering from several backups, not just a full one

▼ How to Recover a Secondary Disk

This section provides an example of how to recover a secondary disk using Backup. The example may apply to either a Backup server or a client.

Caution – It is impossible to provide step-by-step instructions on how to recover your system from a disaster, since every site is unique. The examples in this chapter are designed to give you general principles on how to recover a primary or secondary disk, and to help you understand the procedure. They are meant to be examples only, not instructions.

The following example assumes the primary disk is still operational, so the system has an operating system and can run Backup. However, a secondary disk is lost due to a head crash.

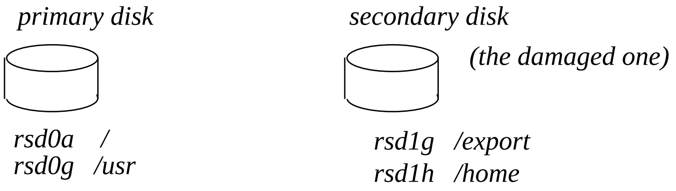

If the disk is damaged, replace it with a new disk of the same type and size as the old one. You will need a disk large enough to hold all the filesystems to be recovered.

1. **Install the replacement disk. Make sure the operating system and kernel recognize the new disk.**

2. **Label and partition the new disk so that you can recover the filesystems.** Use the hard copy of the disk information to remember how large each partition was. (See "File the Disk Information" on page 82)

Note – If you do not have this information, look at the file `/etc/vfstab` for Solaris or `/etc/fstab` for SunOS systems to find out how the disk was partitioned into filesystems. You will have to guess how much space to give each partition. Since you still have the primary disk, the partition information is available.

3. **Make filesystems for each raw partition that you are going to recover and mount the block partition, consulting the hard copy of the df output.** (Backup does not initialize filesystems; it recovers data into existing ones.)

Caution – Using the `newfs` command destroys the disk contents. Be sure this disk is really destroyed before using `newfs`.

For example, for a SunOS system:

```
# newfs /dev/rsd1g
...
# mount /dev/sd1g /export
# newfs /dev/rsd1h
# mount /dev/sd1h /home
```

For example, for a Solaris system:

```
# newfs /dev/rdsk/c0t1d0s5
# mount /dev/dsk/c0t1d0s5
# newfs /dev/rdsk/c0t1d0s7
# mount /dev/dsk/c0t1d0s7
```

After creating and mounting all the filesystems on the replacement disk, use the Backup save set recover feature or the Backup Recover window to recover the files. For more detailed information about save set recover, refer to Chapter 6 of the *Solstice Backup 4.2 Administration Guide*, “Recovering and Cloning Save Sets.”

Recovering a Primary Disk

In this example, a disk with the operating system or the Backup binaries is damaged. You have to first re-install the operating system and boot the system. If necessary, remake and mount all filesystems. Next, re-install Backup from the original software distribution, so you can recover the data.

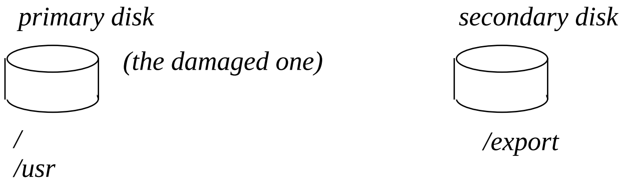

Caution – Any time you have to recover the primary disk (for example, *root*), you should do so in single-user mode from the system console, not from the X window system. Before starting this procedure, make sure all your filesystems are mounted.

After replacing the damaged disk, format it and re-install the operating system, using the original software distribution. Consult the documentation included with the operating system for instructions on how to re-install the system. On a Backup server you should re-install the Backup software using the same paths and directory locations. On Backup clients you only need access to the Backup binaries. You may run Backup from the `cd-rom` directory or NFS-mount the binaries from another system running Backup.

Using the original partition information, make filesystems for each partition you are going to recover and mount them. If a filesystem is already created and mounted, you do not need to do this. For example, if you re-installed `/` and `/usr`, you do not need to re-create them.

Next, re-install Backup from the distribution media. Since you may have several different versions of Backup on the media, use the version number that matches the version of Backup you were running before you lost your disk. For example, if you were running Backup version 4.2, you should install version 4.2 because that is what you were running before your disk crash. The version must be equal to or later than the version used for the backups. Refer to the installation chapters in this guide for detailed instructions.

Note – You do not need to reload the license enablers. If the `/nsr` partition is present, the license enabler is still loaded. If the `/nsr` partition was also lost, recovering `/nsr` (described in the next section) recovers your license enablers.

There is one other way to access the Backup binaries if these were located within one of the damaged filesystems. If there is another system of the same type as the system being recovered on the network which has Backup running, you may NFS-mount the Backup binaries on the damaged system. For example:

```
# mount venus:/usr/etc /mnt
# /mnt/recover -s server -q
recover> add /
recover> force
recover> recover
```

With the operating system and Backup back in place, you are ready to start recovering the remainder of the data lost from the disk.

Use `mmrecov -s` *server_name* or the `recover -s` *server_name* command to recover each filesystem on the disk being recovered.

Caution – If you are recovering `/` (`root`), delete `/boot` for SunOS systems and `/*boot` for Solaris from the save set recover list, **not** the filesystem. If you recover these files, you will not be able to reboot the system. If you recover `/boot` on a SunOS system, for example, you will need to use the `installboot` command. You should always reboot a system after recovering a primary disk. For Solaris, you must also unmark the `/dev` and `/devices` directories from the save set recover list.

▼ How to Recover `/nsr` on a Backup Server

This section addresses the case where the `/nsr` filesystem on a Backup server is lost due to a disk crash. The `/nsr` filesystem contains the indexes which hold the necessary information to recover the Backup clients.

If the server loses its operating system and Backup programs, they must be re-installed first. (See “Recovering a Primary Disk” on page 89)

The next important step is to recover the server indexes from the backup media, using the `mmrecov` command. The `mmrecov` command asks you for the *bootstrap* save set identification number (ssid). If you followed the procedure recommended to prepare for a disk crash, you have a piece of paper with the name of the backup media you need and the *bootstrap* ssid.

In the following example, ssid "1148869870" is the most recent *bootstrap* backup:

```
August 20 03:30 1995 Backup bootstrap information Page 1

  date      time   level        ssid  file  record      volume
8/08/95  7:44:38   full  1148869706     55       0    mars.004
8/09/95  6:12:09      9  1148869754     48       0    mars.005
8/10/95  6:14:23      9  1148869808     63       0    mars.006
8/11/95  6:29:58      9  1148869870     88       0    mars.006
```

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

You may need more than one backup media to recover the server indexes. During the recovery, you can use the `nsrwatch` command or open the Backup Administrator window to watch for pending messages requesting backup media.

With the operating system and Backup in place, recover the indexes from the backup media:

1. **Find the printout of the *bootstrap* save set id information. You need it for the next two steps.**

2. **Retrieve the backup media that contains the most recent backup named *bootstrap* and load it into the server's device.**

3. **Use the mmrecov command to extract the contents of the *bootstrap* backup. For example:**

```
# mmrecov
$pathname/mmrecov: Using mars as server

NOTICE: mmrecov is used to recover the Backup server's on-line file
and media indexes from media (backup tapes or disks) when either of
the server's on-line file or media index has been lost or damaged.
```

```
Note that this command will OVERWRITE the server’s
existing on-line file and media indexes. mmrecov is not used to
recover Backup clients’ on-line indexes; normal recover procedures
may be used for this purpose. See the mmrecov(8) and nsr_crash(8)man
pages for more details.

What is the name of the tape drive you plan on using [/dev/nrst8]?
[Return]
Enter the latest bootstrap save set id []: 1148869870
Enter starting file number (if known) [0]: 88
Enter starting record number (if known) [0]: 0

Please insert the volume on which save set id 1148869870 started into
/dev/nrst8. When you have done this, press <RETURN>: [Return]

Scanning /dev/nrst8 for save set 1148869870; this may take a while...
scanner: scanning 8mm 5GB tape mars.006 on /dev/nrst8
uasm -r /nsr/res/nsr.res
uasm -r /nsr/res/nsrjb.res
uasm -r /nsr/res/
nsrmmdbasm -r /nsr/mm/mmvolume
/nsr/mm/mmvolume: file exists, overwriting
uasm -r /nsr/index/mars/
nsrindexasm -r /nsr/index/mars/db
scanner: ssid 1148869870: scan complete
scanner: ssid 1148869870: 31 KB, 10 files
nsr/index/mars/db: file exists, overwriting
uasm -r /nsr/index/
uasm -r /nsr/mm/
uasm -r /nsr/
uasm -r /
mars: 31 records recovered, 0 discarded.
nsrindexasm: Building indexes for mars...
8mm 5GB tape mars.006 mounted on /dev/nrst8, write protected

The bootstrap entry in the on-line index for mars has been recovered.
The complete index is now being reconstructed from the various
partial indexes which were saved during the normal saves for this
server.
mars# nsrwatch
nsrindexasm: Pursuing index pieces of nsr/index/mars/db from mars.
Recovering 2 files into their original locations
Total estimated disk space needed for recover is 11 MB
Requesting 2 files, this may take a while...

nsrindexasm -r ./db
```

```
merging with existing mars index
mars: 25711 records recovered, 0 discarded.
nsrindexasm -r ./db
merging with existing mars index
Received 2 files from NSR server ‘mars’
mars: 733 records recovered, 0 discarded.
nsrindexasm: Building indexes for mars...
nsrindexasm: Suppressing duplicate entries in mars - 50 duplicates
discarded.

The on-line index for ‘mars’ is now fully recovered.
```

Notice how in the example above, the shell prompt appears during the *bootstrap* recovery. You can use Backup commands such as `nsrwatch` to watch the progress of the server or `nwadmin` to bring up the Backup Administrator window during the recovery of the index. Open a new window (shell tool) to monitor the recovery so that the `mmrecov` output does not display on top of the `nsrwatch` output.

How to Replace the /nsr/res Directory

The `mmrecov` command also recovers the `/nsr/res` directory, which is used by Backup to store configuration information such as the list of Backup clients and registration information. Unlike the indexes, the contents of this directory can not be reliably overwritten while Backup is running. Therefore, the `mmrecov` program recovers the `/nsr/res` directory as `/nsr/res.R`.

To complete the recovery of the `/nsr/res` directory, shut down Backup, move the recovered `/nsr/res` directory into its original location, and then restart Backup.

Complete these steps after `mmrecov` has finished and this final message appears:

```
The on-line index for ‘server’ is now fully recovered.
```

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

```
# nsr_shutdown
```

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

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

3. **Restart the Backup server. When it restarts, the server uses the recovered configuration data.**

```
# cd /
# nsrd
```

4. **Once you have verified that the Backup configuration is correct, you can remove the *\/nsr/res.orig* directory.**

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

To recover other filesystems, see “How to Recover a Secondary Disk” on page 87.

Finding the Bootstrap Save Set ID

If you did not file a hard copy of the *bootstrap* information, you can still find the save set id of the most recent *bootstrap* by using the `scanner -B` command.

For example:

1. **Place the most recent media used for scheduled backups in the server device.**

2. **Use the scanner -B command as *root* to locate the most recent *bootstrap* on the backup media.**

3. **Replace the device name in the appropriate example with the name of your device.**
   For SunOS servers: `# scanner -B /dev/nrst8`

   For Solaris servers: `# scanner -B /dev/rmt/0hbn`

The `scanner -B` command displays information on the latest *bootstrap* save set found on the backup volume, as illustrated below:

```
scanner: scanning 8mm tape mars.006 on /dev/rmt/0hbn
scanner: Bootstrap 1148869870 8/11/95 6:29:58 mars.006, file 88
```

If the *bootstrap* date looks reasonable, run the `mmrecov` command and supply the save set id and file number displayed by the `scanner` command. Otherwise, use another backup volume to try to find a more recent *bootstrap*.

▼ How to Recover to a New Server

This section describes the case where your old Backup server is beyond repair, and you wish to recover Backup to a new server.

Caution – The new Backup server must have the same *hostname* as the old Backup server.

1. **Install the Backup software from the original distribution media on the new server.**

Note – If you have a jukebox, do not start the Backup daemons. Refer to the instructions in this guide for jukebox device driver installation and testing.

2. **Find the printout of the *bootstrap* save set id information from the old server. You need it for the next two steps.**

3. **Retrieve the backup media that contains the most recent backup named *bootstrap*, and load it into the new server device.**

4. **Add the name of the old server as an alias for the new server.**

5. **Use the mmrecov command to extract the contents of the *bootstrap* backup.**
   The Backup daemons should start up on the new server and display the following messages:

```
new_server syslog: Backup Server: (notice) started
new_server syslog: Backup Registration: (notice) invalid auth
codes detected.
new_server syslog:
new_server syslog: The auth codes for the following licenses
enablers are now invalid.
new_server syslog: The cause may be that you moved the Backup
server to a new computer.
new_server syslog: You must re-register these enablers within
15 days to obtain new codes.
new_server syslog:
new_server syslog: License enabler #xxxxxx-xxxxxx-xxxxxx
(Backup TurboPak)
```

Re-register your new Backup server. After moving Backup from one system to another, you have 15 days to register the new server with Sun. Follow the steps in Chapter 4, “Enabling and Registering Backup.”

Sun will send you a Solstice Backup Host Transfer Affidavit which you must complete and return to Sun. Once a signed affidavit is returned to Sun, you will receive a new authorization code, which you must enter into the Auth code field of the Registration window.

Follow these steps after successfully moving your server:

1. **Verify that all the clients are included in the scheduled backups.**

2. **Recover the client filesystems and indexes.**

3. **Use the Recover window to make sure all the client indexes are visible and therefore “recoverable.”**

4. **Back up the indexes on the new server by doing a full backup of the new server as soon as possible.**

▼ How to Use Jukeboxes with Disaster Recovery

The following is a description of how to use jukeboxes during disaster recovery:

1. **Read the disaster recovery procedures listed in the nsr_crash man page. Perform all steps up to the point where you issue the mmrecov command. If only one volume will be needed to recover your Backup file indexes, follow the instructions in nsr_crash.**

2. **Run the jb_config command to add the jukebox.**

3. **Issue the command nsrjb -H. This resets the jukebox for operation. If any volumes are loaded in the media drives, they are moved back to a slot. This operation may take a few minutes to finish.**

4. **Using the instructions in nsr_crash, determine which volume(s) are needed to retrieve the Backup file indexes. Load these volume(s) into the jukebox. Refer to the nsr_crash man pages for additional information.**

5. **Issue the command nsrjb -I. This re-inventories the jukebox. If you want to speed up this process, issue the command with the -S flag and list the slots where you placed the required backup volumes. You must list the slots in order (for example, "nsrjb -I -S *1-3*"). If you want to inventory slots out of order, (for example, 2, 1, and 4,) you must issue the nsrjb -I -S command separately for each slot. All the volumes currently loaded in the jukebox will be marked with an asterisk since there is no media database.**

6. **Load the first volume that mmrecov requests into the first drive in the jukebox. Issue the command:**

   ```
   nsrjb -l -n -S slot -f device_name
   ```

   where *slot* is the slot where the first volume is located and *device_name* is the *pathname* of the first drive.

7. **Run mmrecov. Provide the same device name as in step 6 above, the last save set id, file number, and record number requested by mmrecov.**
   At this point Backup will recover the server file indexes and configuration information.

8. **Issue the nsrjb -u command after the indexes have been recovered.**

9. **Issue the nsr_shutdown command to shut down Backup.**

10. **Rename the** `nsr.res` **and** `nsrjb.res` **files.**

    a. **Restart Backup.**

Summary

The following steps summarize what you need to do if a primary or secondary disk is damaged, destroying the filesystems of a Backup server or client:

1. **Reload and boot the system, if the operating system is lost, using the same *hostname* and disk partitioning, if possible.**

2. **Replace the damaged disk(s), if necessary, and format it, partition it, make new filesystems, and mount filesystems with the same names as those that were damaged.**

3. **Re-install or access the Backup binaries if they were lost. On a Backup server, re-install them from the distribution media. On a Backup client, mount the CD-ROM or temporarily NFS-mount them over the network.**

4. **Use mmrecov to recover the indexes for the Backup server if the** `/nsr` **directory is destroyed on the server.**

5. **Recover the lost filesystems by using the normal recover process or the save set recover feature. Recover the client indexes, using the normal recover process.**

Managing the Backup Environment A

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

For your convenience, Backup is shipped with preconfigured backup schedules, index policies, pools, label templates, directives, and notifications. A brief description of each preconfigured setting is presented in this appendix.

Guidelines for Choosing a Configuration

There are several factors that affect the Backup server configuration that best suits your backup and recover needs. The configuration consists of the hardware and software, including tape drives, jukeboxes, client systems, and network connections.

This section provides a few simple rules that you can use to guide your choices and focuses on backup, since backup requires far more server capacity than recover.

Caution – Please keep in mind that these are *guidelines*; actual performances may vary.

The goal in selecting a configuration is to balance the different hardware and software limitations to achieve the overall data handling capabilities you require. Start by looking at the limits of the major Backup configuration components: tape drives, clients, network connection, jukeboxes, and the Backup server itself.

Tape Drives

Tape drives have a fixed maximum data transfer rate that they can handle. Since Backup automatically spans multiple tapes, the total tape capacity is not as important as the data rate. Refer to your hardware documentation to find out what the data transfer rate is for your drive.

Backup cannot back up faster than the data rate of your tape drive, but multiple tape drives can decrease backup time.

Clients

Different clients can generate data at different rates and, even within a single client, different types of files can generate different data rates. For example, symbolic links require as much processing as large data files, but produce no data. Consequently, the data rate produced by a backup of a single client can vary quite a bit. It is a good idea to run several clients simultaneously to help smooth out fluctuations in the data transfer rate for each client.

Network

Ethernet has an upper limit on bandwidth of about 1 Megabyte (MB) per second, but in practice, most networks can only handle about 500 KB between a set of clients and a single server. Token ring has a lower maximum bandwidth (8 Mbits/s) but a higher utilization so the data transfer rates are approximately the same.

Network	Rate
Ethernet	500 KB/s
Token Ring	500 KB/s
FDDI	5 MB/s

Server

The server must be able to handle the load of network packets, data movement, and tape drives in order to achieve the rates listed above. Most of the work on the server side is in data movement, context switching, and interrupt handling. The performance of all of these functions improves as the CPU speed increases. It takes approximately 20 MIPS to handle 500 KB/s of data, although this tapers off at high CPU speeds because bus bandwidth and other bottlenecks begin to affect the data movement. Approximately 16 MB of memory is required per 500 KB/s of data rate handled by the server.

Server CPU/Memory	Rate
20 MIPS/16 MB	500 KB/s
50 MIPS/32 MB	1000 KB/s
100 MIPS/64 MB	2000 KB/s

Jukeboxes

Jukeboxes provide automatic loading and unloading of tapes or optical disks. This assists the administrator in two different ways. During nightly backups, Backup uses the jukebox to automatically switch to new media as the data is backed up. During recovers, Backup uses the jukebox to load all of the media needed for the recovery without operator intervention. Refer to the jukebox documentation for data transfer rates and maximum capacity.

To determine the capacity requirements of a jukebox for a scheduled unattended backup, simply pick the jukebox with a capacity large enough to handle the largest possible amount of backup data. For example, a full backup of 60 GB requires an ATL (automatic tape loader) or Lago jukebox.

To determine how much disk space you will need for the online indexes for quick recovers, you first need to do a rough calculation of the amount of data backed up in a single schedule period (for example, week, month, or quarter). To do this, use the following guidelines to determine how much data is backed up with different levels of backup.

Level	% of data backed up
Full	100%
Level 1-9	25%
Incremental	10%

For example, a monthly schedule that has 1 full on the first Sunday, a level 5 on other Sundays, and incrementals every other day will look like this:

Level	% of data backed up
1 Full	100%
26 Incremental	260%
4 Level 5	100%
Total	460%

This illustrates that over the course of a month, 460% of the total amount of data will be backed up. For example, a total of 10 GB of client data backed up using this schedule would result in about 46 GB of data on tape per month.

In this example, different jukeboxes would provide the following amount of online data:

Jukebox	Capacity	Months of data
EXB-10e	50 GB	1
ATL/Lago	270 GB	5
EXB-120CHS	580 GB	12

Guidelines for Configuring Server/Client Network

Using the previous data transfer rate and capacity calculations, you may consider the following guidelines for configuring your network of servers and clients:

1. **Assign 4 simultaneous clients per network, by setting the parallelism to 4 (if the clients are PCs, you may assign more).**

2. **Use 1 Exabyte 8500 or Sony 5200 per network, or 2 Exabyte 8200s.**

3. **Assign a Backup server with approximately 20-30 MIPS and 16 MB of memory per network.**

Configuration Examples

This section includes two examples of Backup configuration to illustrate the reasoning behind selecting the components.

Example 1

Site A has approximately 30 GB of data on 2 networks of 50 clients and wants to schedule full backups for all of their data in one night (12 hours). The following equation calculates the required data transfer rate to achieve this goal.

30000 MB / 12 hours = 2500 MB/hour = 694 KB/sec

To back up the data with a single Backup server, the following configuration is suggested:

- Backup TurboPak installed on a 30 - 40 MIPS CPU system with 24 - 32 MB of memory
- two Exabyte 8500 tape drives connected to the Backup server
- an ATL or Lago jukebox with two Exabyte 8500 tape drives connected to the Backup server
- two network interfaces
- one Backup Jukebox Software Module for the ATL (or Lago)

To back up the data with two Backup servers, the following configuration is suggested:

- Backup TurboPak installed on two 20 MIPS CPU systems with 16 MB of memory
- an EXB-10e jukebox with an Exabyte 8500 tape drive connected to each server
- one network interface for each server
- two Backup Jukebox Software Modules for the EXB-10e

Example 2

Site B has 50 GB of data on a single network with 80 clients and they want to be able to schedule backups in a single night (12 hours). The full backups for the clients must be staggered due to the limit of 500 KB/sec data transfer rate per network. Calculate the backup capacity required to complete the backups in one night:

```
500 KB/sec * 12 hours = 21.6 GB/night
```

By staggering their full backups into three nights instead of one, using three different backup schedules, they reduce the load of the nightly backup data to about 20 GB/night.

Full:50GB/3 = 16.7 GB

Incr:(50 GB - 16.7 GB) * .1 = 3.3 GB

Total:20 GB/12 hrs = 462 KB/s

To back up the data with a single Backup server, the following configuration is suggested:

- Backup TurboPak installed on one 20 MIPS CPU system with 16 MB of memory
- an EXB-10e jukebox with an Exabyte 8500 tape drive connected to the server
- one network interface
- one Backup Jukebox Software module for the EXB-10e

Measuring Performance

If you are interested in measuring the performance of your Backup environment, you must take into consideration the server system and the client system.

The factors to consider for the server system are the speed of the tape drive, the network speed, and the CPU speed. Factors to consider for the client system are filesystem traversing, generation of data, data on multiple disks, and CPU speed.

Server Performance

This section provides examples on how to measure the performance of the server.

How to Measure Tape Drive Speed

Most tapes have step function in data rate. Backup uses 32KB/record. To measure tape speed, follow these steps.

1. **Create a large file (at least 20 MB) with non-zero data. For example:**

```
# cat /vmunix /vmunix /vmunix /vmunix /vmunix > big
```

2. **Use the dd command to write the large file to tape four times and measure the time results:**

```
/bin/time dd if=big of=/dev/nrst8 bs=32k conv=sync
/bin/time dd if=big of=/dev/nrst8 bs=32k conv=sync
/bin/time dd if=big of=/dev/nrst8 bs=32k conv=sync
/bin/time dd if=big of=/dev/nrst8 bs=32k conv=sync
```

3. **Divide the file's size by the average of the last three real times. For example:**

```
-rw-rw-r-- 1 root 20675420 Jan 7 11:04 big
95.2 real 13.0 user 11.9 sys
78.2 real 12.9 user 12.7 sys
78.0 real 12.8 user 12.5 sys
76.8 real 13.0 user 12.4 sys
Rate: 20190 KB / 77.66 sec = 260 KB / sec
```

This number gives you the rate of the tape speed.

How to Measure Network Speed

Backup for UNIX uses TCP and RPC/XDR as network communication protocols. To measure the network speed, follow these steps:

1. **Create a large file (as in the tape speed measurement example) on a fast client.**
2. **Use the rcp command to copy the file from the client to the server and time the result:**

```
# /bin/time rcp big server:/dev/null
```

3. **To find the network speed, divide the number of bytes in the file by the real time. For example:**

```
-rw-rw-r-- 1 root 20675420 Jan 7 11:04 big
38.2 real 0.2 user 30.7 sys
Rate: 20190 KB / 38.2 sec = 529 KB / sec
```

The most important factor affecting network speed is network errors. To determine the input error rate, the output error rate, and the collision rate, use the `netstat -i` command. If the input or output error rate is above 0.5%, or the collision rate is above 5%, network errors are slowing down the network speed.

CPU Speed

The CPU of a server limits the following:

- the total data throughput to tape
- the interrupts per second for network data
- the context switches per second between processes

The best measure is the MIPS rating for the server. A larger MIPS rating means a faster machine.

Memory

The memory on the server limits the amount of data buffered between the Backup `save` command, agent daemon, and media management daemon.

Client Performance

This section provides examples on how to measure the performance of the Backup client.

▼ How to Measure Filesystem Traversing

To measure the filesystem traversing speed, follow these steps:

1. **Time the uasm command with the -bi option. For example:**

```
# /bin/time uasm -bi /usr
13931 records 2667396 header bytes 350849148 data bytes 124.9
real 10.8 user 34.0 sys
```

2. **Divide the number of records by real time for rate per file. For example:**

```
# 13931 records / 124.9 sec = 111.5 files/sec
```

▼ How to Measure the Data Generation Rate

To measure the rate at which a client generates data for a backup, follow these steps:

1. **Time the uasm command with the -si option and redirect the output to** `/dev/null`**. For example:**

```
# /bin/time uasm -si /usr > /dev/null
```

2. **Divide the number of bytes obtained (filesystem traversing) with the uasm -bi command by the real time generated by the uasm -si command. For example:**

```
342626 KB / 1199 = 286 KB / sec
```

Data on Multiple Disks

Backup automatically backs up multiple disks in parallel.

To measure parallel disk speeds, follow these steps:

1. **Use the df or du command to find two directories of approximately the same size located on different disks.**
2. **Run the same uasm speed tests for filesystem traversing and data generation rate as for one disk, but run the tests simultaneously on the two directories.**
3. **Add the data from each test (*files/sec* and *KB/sec*) to obtain a combined rate.**

This rate reflects the performance of Backup backing up data on multiple disks.

CPU Speed

The CPU of a client limits the following:

- the total data throughput to tape
- the interrupts per second for network data
- the context switches per second between processes

The best measure is the MIPS rating for the client. A larger MIPS rating means a faster machine.

Backup Schedules

The Backup server backs up each client system across your network according to a backup schedule. Create schedules in the Schedules window and assign individual clients to schedules in the Clients window. Schedules can be very simple or very sophisticated, depending on the needs of your environment. All clients can share the same schedule, or each client can have its own unique schedule. This section discusses some of the considerations you should keep in mind while determining which schedule best fits your situation.

Backup Levels

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

- Full – backs up all files, regardless of whether or not they have changed since the previous backup. A full backup is equal to a level 0 backup.
- Level 1 through level 9 – back up files that have been modified since a previous full backup or a backup of a lower numbered level. For example, a level 3 backs up all the files that have changed since the previous level 2, level 1, or full backup.
- Incremental – backs up all files that have changed since the previous backup of any level. An incremental backup is equal to a level 10.

You can also skip a backup on a given day. You may want to schedule a “skip” backup on weekends or holidays when no one is available to load backup media. If you have a jukebox, this option becomes less important, since a jukebox automatically loads and unloads backup media.

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

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

Full Backups versus Incremental Backups

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

The situations you should consider are listed below:

- full backups take more time to complete than incremental backups.
- if the full backup does not fit on a single piece of media, someone has to monitor the backup and change the media (unless you have a jukebox).
- full backups cause the online indexes to grow more rapidly than incremental or level backups.

You may decide to schedule a full backup at the beginning of the period and then schedule incremental backups the rest of the period. This schedule minimizes the amount of time that the backups take, the size of the backups, and the size of the Backup indexes.

However, if you need to recover from a disk crash, you may need all the tapes used during the schedule, because the most current version of your files may be scattered across several different tapes. Although Backup asks for each tape that it needs for the recovery by name, loading and unloading them can be time-consuming (unless you have a jukebox, or all the incremental backups fit on one tape).

Using Level Backups

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

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

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

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

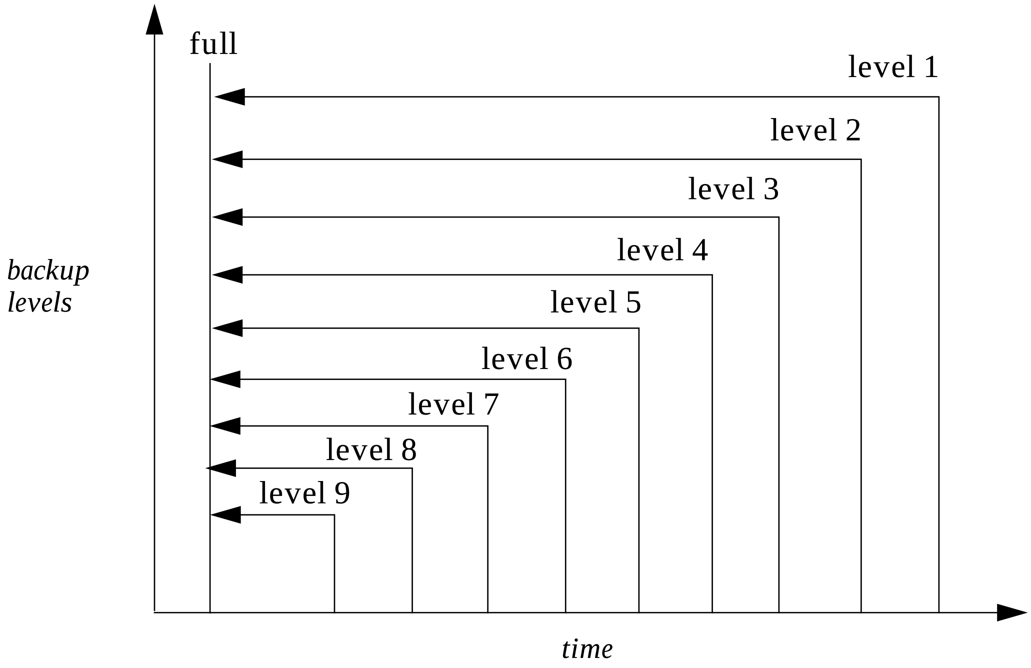

The advantage of this schedule is that to recover from a disk crash, you only need two tapes: the one with the full backup, and the one with the last level backup. The disadvantage is that with each day, there are more changed files to back up, so the backups take longer to complete.

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

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

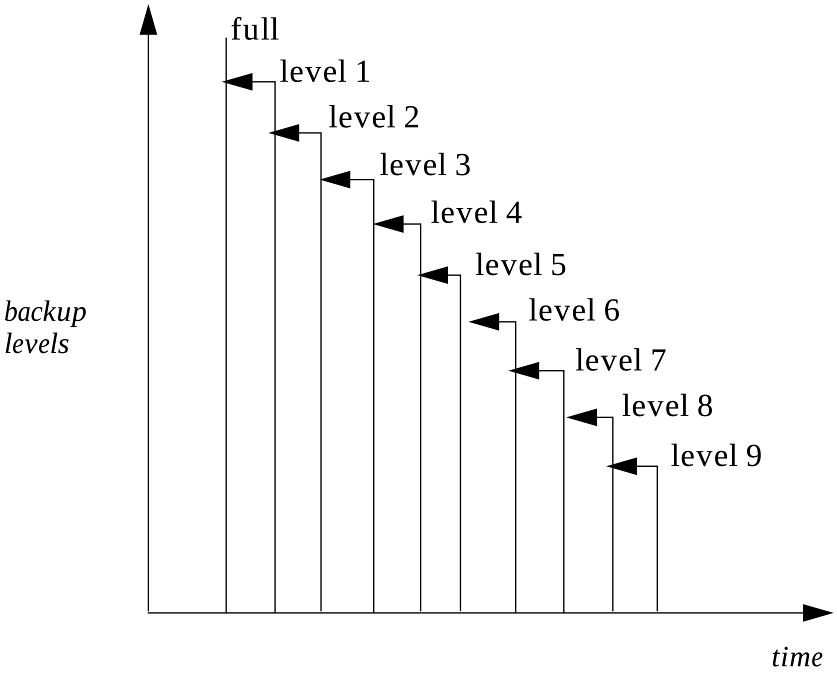

The advantage of this schedule is that each day's backup will be small and will complete in a short period of time. The disadvantage is that recovering from a disk crash will require the full backup tape and all of the level backup tapes up until the day of the disk crash.

Neither of these backup schedules is practical. They simply illustrate how level backups work. The real power of level backups comes into play when you combine multiple levels along with fulls and incrementals.

A Typical Monthly Backup Schedule

Sites with even a few gigabytes of files to back up often choose a monthly schedule based on fulls, incremental, and level backups. The example described in this section performs a full backup on the first day of each month, a level 5 backup on the 10th and 20th of the month, and incremental backups on all other days.

This monthly backup schedule minimizes the size of daily backups while also making it relatively easy to recover in the event of a disk crash. This schedule offers several advantages.

First, the level 5 backups simplify recovery. Assume that a disaster strikes on the 24th of the month. All the files that you need to recover an entire client system are located on tapes from just five backup sessions:

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

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

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

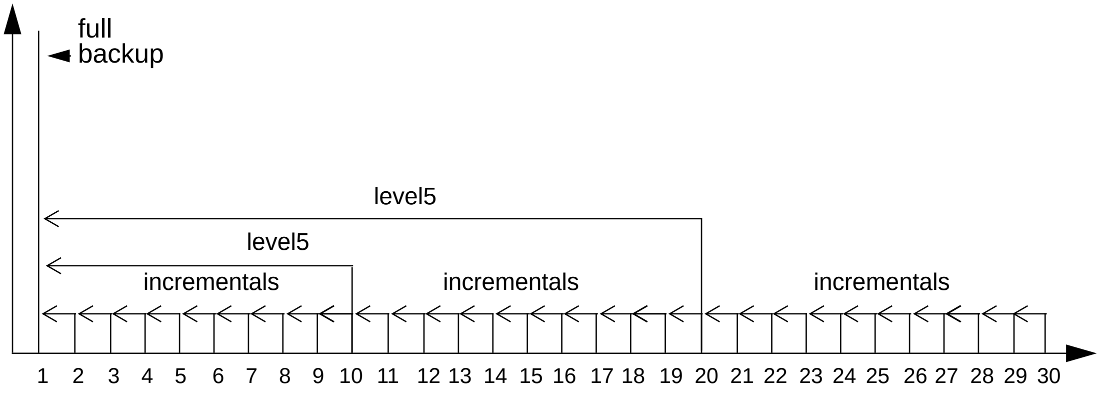

Backups Take Time

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

- Select a backup server with enough CPU power, memory, and bus bandwidth so the backup server is not a bottleneck.
- Leave Backup’s parallelism feature turned on. This feature causes multiple client systems to send their files to the backup server in parallel. This keeps a stream of files ready for the tape drive, so that it does not start and stop. Find a parallelism that keeps the tape drive streaming without overloading the CPU.
- Experiment with compressing files on the client systems to reduce the size of the data sent across the network and written to tape. Compression may speed your backup as long as the client systems are still able to supply files to the backup server fast enough to keep the tape drive streaming.
- Take advantage of Backup’s ability to skip specified files during the backup. For example, you could choose to skip *core* files.
- Add a second backup device and add Backup TurboPak to your backup server. With TurboPak Backup can simultaneously back up to more than one device.

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

Unless you have a jukebox and Backup’s optional Jukebox Software Module, you also have to schedule backups based on someone being available to load and unload media. Many administrators find that an incremental backup of their network fits onto a single 4mm or 8mm tape, but they must schedule multi-tape full and level backups for specific nights or weekends when an operator is on duty to load additional tapes. If an operator will not be available over a holiday weekend then you can set an override in the schedule to skip the backup on that day. You may also want to override the schedule just before a holiday with a full backup – for added peace of mind.

Using Compression During Backup

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

This reduces network traffic, as compressasm typically achieves 2:1 compression (your mileage may vary). There is no harm in using compressasm in conjunction with a compressing tape drive, but the drive will probably not achieve much compression on the data.

If you are deciding between compressasm and a compressing drive solely to increase the amount of data on a tape, obtain the compressing drive – the hardware on the drive will generally compress faster than Backup and place *no* load on the CPU due to the compression.

Follow these guidelines when compressing data during backup:

- Use compressasm to minimize network bandwidth if you have available CPU power.
- Use compressing drives to get more data on a tape.
- Any generic compressing algorithm typically achieves 2:1 compression and tape drives are no exception. Sometimes you get more, sometimes less.
- Compressing already compressed data has no effect (and may even expand the data!)
- Do not use compressasm if you have a compressing drive and no networked clients.

Staggering the Backup Schedules

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

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

Convenience versus Security

Backup is a sophisticated product. You may leave the same backup volume mounted in the server's backup device throughout a week or month, and when it becomes full, replace it with a new labeled backup volume. Backup tracks all the backups, no matter what day of the week or month, or what part of the backup schedule cycle is in effect. The same backup volume may contain full, level [1-9], or incremental backups, and it makes no difference to Backup. For you, the benefits are fewer backup volumes to manage and the ability to recover from a disk crash with a minimum number of backup volumes.

Some sites prefer to segregate the full backups from the level [1-9] and incremental ones. The full backups protect the network from a catastrophic disk loss, and you want to guarantee their integrity. There is always a very small risk that if you leave the backup volume with the full backup sitting in the backup device, something could happen to it.

If a backup volume with incremental backups is ruined, users may lose one day of work. If the backup volume with the full backup is destroyed, users may lose all the work done since the last full backup. Therefore, some administrators prefer to remove the backup volume used for a full backup, put it in a safe place, and mount another backup volume for the following level [1-9] and incremental backups. The trade-off is that you may need a few more backup volumes to recover from a disk crash – the one with the last full, and the other volumes that contain the most recent level [1-9] and incremental backups.

Backup Browse and Retention Policies

Backup maintains online indexes of all the files backed up for each client and an index of the files stored on each piece of media. Backup lets you set policies that automatically control how long the information is retained in these online indexes. This section explains Backup's browse and retention policies and the trade-off between providing faster, easier recovery for your users or conserving disk space.

Browse Policy

One of Backup's popular features is the ability to browse many versions of a file that have been backed up over time and to choose which one to recover. However, each version of a file that Backup tracks takes up space in the client online index (about 220 bytes each). Since disk space is limited, you need to establish a policy of how far back in time you will keep information about backed-up files in the indexes.

The browse policy that you select specifies how long the entries for your files remain in the file indexes. A browse policy can be any number of days, weeks, months, or years. Backup automatically deletes entries older than the browse policy time and frees up disk space. The browse policy you select, like the backup schedule, can be different for each client.

How Browse Policies Work

To recover a complete directory or file system, you often need to recover some files from incremental and level backups as well as from a full. The incremental backup is dependent on the level backups and, in turn, on the full. Backup will not delete the entries from any backups on which other backups depend. As a result, you may find that entries are deleted later than you expect.

In the illustration shown below the browse policy is set to one week, which happens to equal one complete backup cycle.

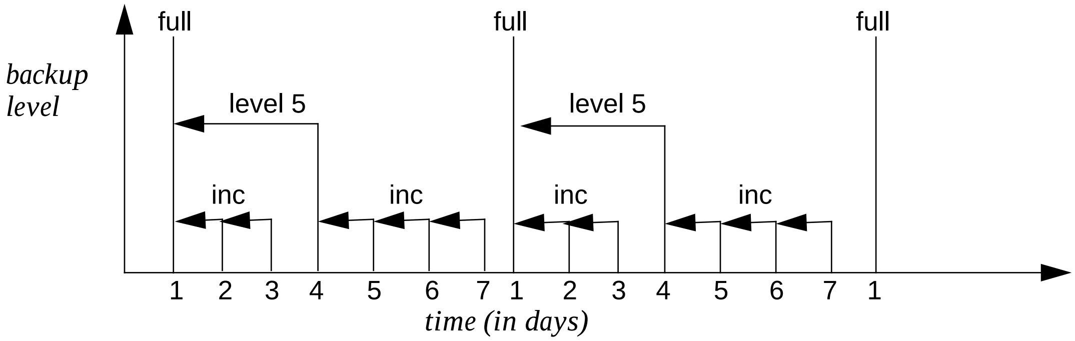

Backup will not remove the first full backup from the online file index until all the incremental and level 5 backups that depend on it have expired. As a result, the full backup actually stays in the online index for a period of time equal to the browse policy plus one full backup cycle.

The first full backup will not be removed from the online index in exactly one week, however, because there are incrementals and a level 5 backup, which have not yet expired, that depend on the full. Each incremental backup will be removed from the online index one week from the time it was completed. The level 5 backup will be removed one week after the last incremental that depends on it is removed, and then the full backup will be removed at that same time.

The rule to remember is that a full backup actually remains in the online index for a period of time equal to the browse policy plus one complete backup cycle. A backup cycle is measured from one full backup to the next full backup. Also note that the browse policy is set for an entire client (or filesystem, if the filesystems are separately scheduled). Consequently, whatever policy you have for keeping full backups online and browsable in the file index you must also use for all incremental and level backups. With Backup you manage backup cycles (the period from one full backup to the next); you do not independently manage different levels of backups.

Reclaiming Disk Space

Backup automatically reclaims disk space that is freed up when entries are deleted from the online file indexes. However, the space is not returned immediately to your system. Backup takes some time, processing power, and swap space in order to reclaim this space and to have this constantly taking place on your backup server is inefficient. Instead Backup first reuses this space to store information about new files that are backed up. When the file index for a client reaches a point where less than 50% of its space is being used by files that have not reached the end of their browse period, then Backup automatically invokes a process that returns the space to your system.

You may also reclaim disk space at any time by using the Reclaim space button in the Indexes window.

Recovering Files Removed from the Index

You can recover files whose entries have been removed from the online index because they have passed the Browse policy period as long as the files are still stored on a backup volume. However, the recover process is not as convenient as when the entries are still in the online index.

If you do not want to rebuild the index, the save sets you need are still in the media index. Since you know which save set contains the file you want, you can use the save set recover feature to recover the entire save set or selected directories and files. The save set recover feature is most useful for recovering from full backups and is limited to *root* and users belonging to the group *operator*.

If you want to rebuild the file index so that you can browse for the file you lost, here is the basic procedure to follow.

1. **Use the Volume window to find the name of the backup volume that contains the save set.**
2. **Use the mminfo command to determine the save set id. Use this syntax:**

```
mminfo -v -s server -c client -N saveset volume_name
```

3. **Rebuild the file index entries for the save set using the scanner -i -s *save_set_id#* command at the system prompt. Enter the save set id number determined above for save set id#. Rebuilding the file index using the scanner command may take some time.**
4. **Use the Backup Recover window to identify the needed file(s) and initiate the recovery.**

Recovery is considerably easier if the file information is still in Backup's online index. That is why you want to set a browse policy long enough to cover most recover requests.

Media Retention Policy

Your need to conserve disk space may lead you to establish a short browse period. Backup's media retention policies complement the browse policy by letting you specify a longer period of time during which files can still be recovered, although with more difficulty. Backup uses the retention policy to automatically recycle backup volumes.

Backup maintains a file index for each client system and a much smaller media index that tracks which save sets are stored on each backup volume. When Backup removes entries that are older than the specified browse time from a file index, it leaves the corresponding save set information in the media index. The retention policy controls how long this information is kept and, as a result, how long a backup volume is kept before it can be overwritten with new backups.

As with the backup schedule and browse policy, you set the retention policy for each Backup client. Different clients can have different policies. The retention period can be any number of days, weeks, months, or years as long as the retention period is equal to or longer than the browse policy.

A Backup backup volume can contain save sets for many different clients over many days. As the retention period is reached for each save set, information about that save set is removed from the media index. When the retention period for every save set on a backup volume is reached, Backup marks the volume "recyclable." This volume can then be reused for backups. At the time that the volume is actually reused, the old files are overwritten and can no longer be recovered.

Backup's browse and retention policies combine to give you a hierarchy of recovery capability while keeping the disk space needed for the online indexes to a minimum. Recovering a file is quick and easy using the Backup Recover window until the browse policy time is reached and the file information is removed from the file index. Then you can use save set recover or the more tedious process described to recover your files until the retention policy time is reached and the backup volume is recycled.

Setting Policies When Using a Jukebox

Backup’s Jukebox Software Module automates your backup and recover activity. The capacity of the jukebox, the backup schedule you select, and the browse and retention policies you use determine whether you can walk away from backups for a week, a month, or even longer.

Jukebox Capacity

A jukebox is most useful if it has at least enough capacity to complete one entire backup cycle without intervention. This allows backups to run while you are out ill, on vacation, or busy with a user emergency and helps minimize the time that you spend on backup (particularly if the backup server and jukebox are located some distance away). At the end of the cycle, you can move the used backup volumes offsite and load fresh tapes into the jukebox.

A jukebox with the capacity for one entire backup cycle also speeds file recovery. If a user accidentally deletes a file, there is at least one version (more if the user has recently edited the file) in the jukebox. With Backup, the user can quickly identify the lost file and initiate the recovery. The jukebox will load the needed tape and Backup will complete the recovery without your help. Depending on the speed of the jukebox and the device used, the file should be recovered very quickly.

You need to design a schedule that fits the capacity of your jukebox. Start with your ideal schedule and then consider these suggestions to reduce the size of your complete backup cycle:

- use more incremental backups and fewer level 1-9 backups
- back up systems with less critical files less often – perhaps only once a week
- use Backup’s directives to skip files during the backup, for example, *core* files
- shorten the length of the backup cycle

Although your jukebox may only have enough capacity for one backup cycle, you can still set the browse and retention policies for a longer period. If a user tries to recover a file stored on a volume that is not in the jukebox, Backup will prompt you to load that volume. You can use the Location field in the Volumes window to keep track of volumes. Users can refer to this information when deciding which version of a file to recover and choose the one stored on a tape that is located in the jukebox.

Choosing the Jukebox Capacity

With just enough capacity in the jukebox for a single backup cycle, you must reload tapes at the end of each cycle. With more capacity you can set the schedule and the browse and recover policies so that the jukebox runs unattended for a long period of time. The jukebox will automatically recycle tapes containing save sets that have passed their browse and retention times to continue backups virtually indefinitely.

Suppose you established a backup schedule for your network of systems that takes one week to complete (for example, you schedule a full backup once a week) and which consumes a total of 12 GB of tape during the week. Assume that you are using a 50 GB EXB-10e jukebox. Each of the following combinations of browse and retention times will allow the jukebox to operate without intervention for an extended period of time:

- browse policy = 1 week, retention policy = 1 week
- browse policy = 1 week, retention policy = 2 weeks
- browse policy = 2 weeks, retention policy = 2 weeks

Each of these sets of policies has its advantages. With a browse and retention policy of just one week, your online indexes will be kept small. With a browse and retention policy of two weeks, your indexes will be larger but your users will have more versions to select from when they need to recover a file. A browse policy of one week and a retention policy of two weeks keeps your indexes small and allows you to recover older files, although with a great deal more effort than if those files were still browsable in the index.

If you set the browse policy to four weeks, 4 * 12 GB = 48 GB will fit in the jukebox. First, a full backup actually remains in the online index for a period of time equal to the browse policy plus one complete backup cycle. Thus with a browse policy of four weeks, essentially five weeks of backups would need to fit into the jukebox.

Second, since Backup cannot recycle a tape until all the save sets on that tape have expired there is often some amount of “unavailable” tape in the jukebox.

Now suppose that one year later the number of files that you have has grown so that the one week backup cycle needs 18 GB of tape capacity. A browse policy of one week and a retention policy of one week still allow the jukebox to run unattended on an on-going basis.

If you want to keep files online in the jukebox for a longer period of time, then you can use the methods listed earlier to reduce the size of the backup cycle. As an alternative, you can stretch out the backup cycle. For example, you can perform full backups every other week rather than every week. This should not greatly increase the size of a backup cycle and gives you more versions of files online in the jukebox.

Choosing a Jukebox

In the ideal situation, you first design the best backup schedule and set of policies for your environment and then determine the jukebox size that you need to purchase.

Assume that you have a network of systems with a total of 25 GB of files to back up and that you have selected a schedule that includes a full backup at the beginning of each month, a level 5 on the 10th and 20th of each month, and incrementals on all other days. The calculation below illustrates that one complete backup cycle will be about 64 GB in size.

Level	**Size**	**Frequency**	**Total**
Full	25 GB	* 1 time /month	25 GB
Level	52.5 GB	* 2 times/ month	5 GB
Incremental	1.25 GB	*27 times/month	34 GB
			64 GB

Note – These size percentages are based upon experience with Backup over the past 3 years

To determine the size jukebox that you need, start by estimating the size of a complete backup cycle. Now assume that you have decided on a browse policy of 2 months for all the client systems and a retention policy of 6 months. These policies let your users quickly recover any file and any version of a file that they had during the past 2 months. And with some effort you can recover for them files that they had any time during the past 6 months. So you will need 6 months * 64 GB = 384 GB of capacity.

In practice you need a little extra jukebox capacity since there will be a small number of "unavailable" volumes as Backup must wait to recycle a tape until after all the save sets on that tape have expired.

Finally, remember to plan for growth in the number of your files. While sites differ in the rate at which their files are growing, a rule of thumb is that you should purchase a jukebox with about 50% more capacity than your current requirement.

Preconfigured Selections

Backup provides preconfigured settings for you to use so you can immediately start backing up your systems. This section offers an explanation of the different preconfigured settings.

Preconfigured Backup Schedules

For your convenience, Backup is shipped with several preconfigured backup schedules. If these schedules fit your backup requirements, you can use them "out of the box," or you can create new ones to accommodate your site-specific needs.

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

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

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

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

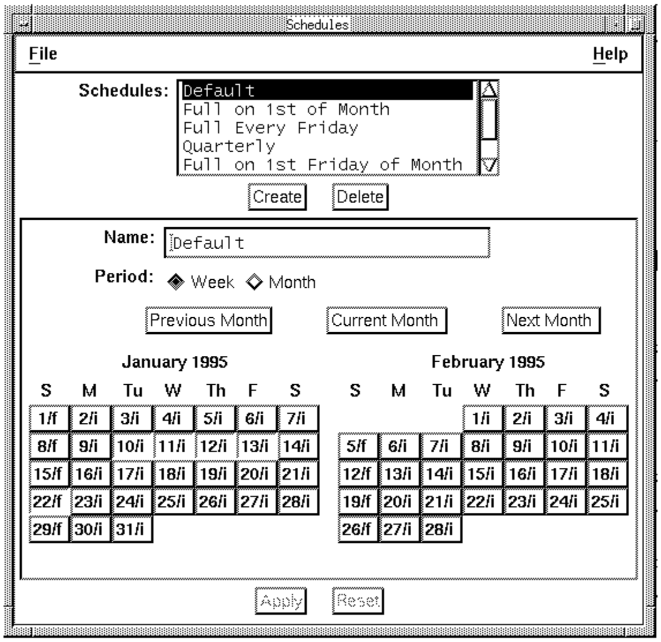

Schedules: Default
Full on 1st of Month
Full Every Friday
Quarterly
Full on 1st Friday of Month

Name: Default

Period: ◆ Week ◇ Month

January 1995

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

February 1995

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

- Default – this is the only schedule you may not delete. It is a weekly schedule and completes a full backup every Sunday, followed by incremental backups all other days of the week.
  This schedule is convenient if you want to premount the backup volume Friday night before you go home for the weekend. On Monday mornings, check your messages from Backup to make sure the backup completed. If you want to separate the full backups from the incrementals, remove the backup volume with the full backup and mount another one for the incremental backups.

- Full Every Friday -— this weekly schedule completes a full backup every Friday, followed by incremental backups the other days of the week.
  This schedule is identical to the Default schedule, except that instead of completing a full backup on Sundays, the full backup takes place on Fridays. Depending upon how much data changes on the network, the daily incremental backups might all fit onto one backup volume. In that case, if you had to recover from a disk crash, you would need only two backup volumes – the one with the last full backup, and the one with the incremental backups.

- Full on 1st Friday of Month – this monthly schedule completes a full backup on the first Friday of the month, (not the first calendar day of the month. Incremental backups take place on all the other days.
  The advantage of this schedule is that you complete a full backup only once a month. If you use this schedule, it would be a good idea to store the backup volume with the full backup in a safe place, and use other backup volumes for the incremental backups. It would also be a good idea to change backup volumes every few days for the incremental backups. If you allow all the incremental backups to be stored on one backup volume, and it is destroyed near the end of the month, you may not be able to fully recover from a disk crash.

  Whenever you create a monthly schedule for a full backup on a *weekday* instead of a *calendar* day (like Friday, in this example), you must set the overrides in each month. (Notice the "f*" in the first Friday of each month.) This is because the first weekday (Monday through Friday) in a month may fall on any calendar day from 1 to 7.

Note – The Overrides you select for individual days do not carry over from one year to the next. Preconfigured schedules, however, do maintain the overrides for years into the future.

- Full on 1st of Month -– this monthly schedule completes a full backup on the first calendar day of the month. On the other days of the month, an incremental backup takes place. This schedule has the same advantages and disadvantages as the "Full on 1st Friday of Month" schedule. This schedule is easier to create because you do not have to manually set any overrides.
- Quarterly – the quarterly schedule completes a full backup on the first day of the quarter. A level 5 backup takes place on the first day of the other months in the quarter. Every seven days, a level 7 backup takes place. The other days of the month, an incremental backup takes place.
  This schedule is convenient because a full backup takes place only once a quarter. On the first day of the month, a level 5 backs up everything that has changed since the first day of the quarter. Every seven days, the level 7 backup protects all the data that has changed since the first day of the month. The daily changes are protected by incremental backups.

  If you use this schedule, it is a good idea to segregate the backup volumes and store them in a safe place. Use one volume for the full backup, one for the level 5 backups, one for the level 7 backups, and another one for the incremental backups. If you have a disk crash or a disaster, you only risk losing a few days' work (the backup(s) on the mounted volume). If you change the backup volume every day for the incremental backups, you only risk losing one day's work; however, you must use more tapes to recover from a disaster.

When you create a quarterly schedule like this one, use the Month period to set the level backups, then set each quarterly full backup on the calendar with an override.

To recover from a disk crash, you would need the backup volume with the full backup, the latest level 5, the latest level 7, and the incremental backups for the week.

Preconfigured Policies

Backup is shipped with five preconfigured policies: Decade, Month, Quarter, Week, and Year. Use these policies to choose the length of time to retain the entries in both the file index and media index. Remember, the *retention* policy you select affects the size of the *media* index and controls the length of time Backup tracks the backup volumes and the data on each volume.

The *browse* policy affects the size of the *file* index and the length of time Backup retains entries for every file backed-up and visible in the Recover window. You must always choose a retention policy which is greater than or equal to the browse policy.

For example, if you choose Quarter for the retention policy for a client, and Month as the browse policy, the client will be able to browse all the file entries for backed-up files dating back a month. Each month the oldest entries for the client's files will automatically be removed from the server's file index. However, the backup volumes which contain the files are still tracked by Backup in the media index.

The Policies window and the five preconfigured policies are shown below.

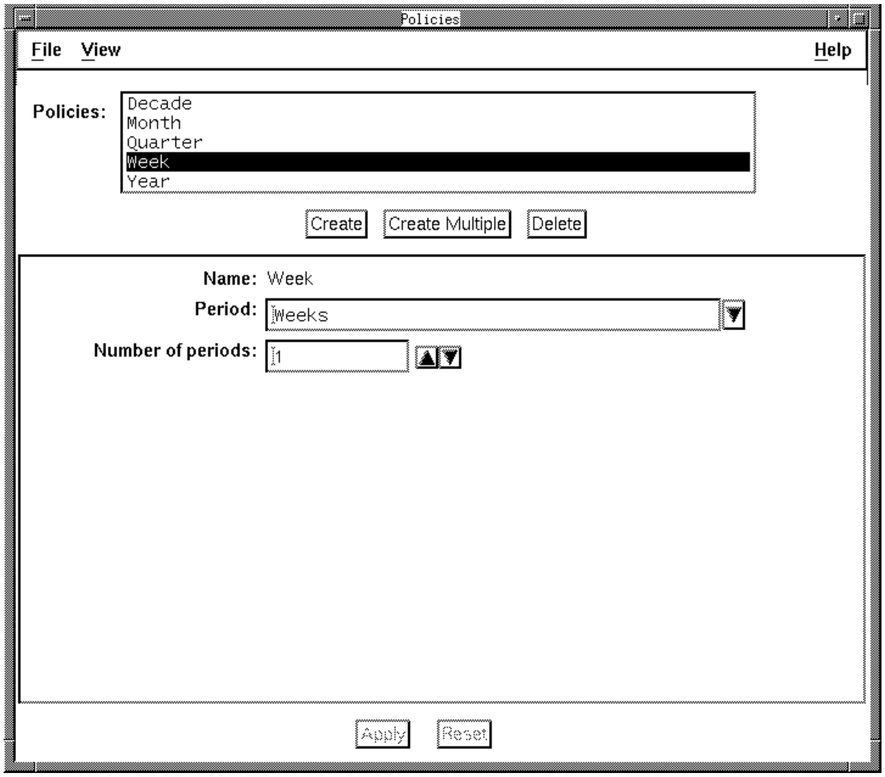

Policies: Decade
Month
Quarter
Week
Year

Name: Week

Period: Weeks

Number of periods: 1

Refer to Chapter 3, “Configuring and Monitoring Clients,” of the *Solstice Backup 4.2 Administration Guide* for an illustration of how browse and retention policies work.

- Week – this policy maintains the file index entries or the media index entries for one week after the last full backup. If you use this browse policy, the users will only be able to view and mark files for recovery which go back in time for a week. It is a useful browse policy when you have a limited amount of disk space and users do not expect to be able to recover versions of their data which are older than one week.

  As a retention policy, Week means that your backup volumes will turn over quickly, and Backup will recycle through the tapes at a faster rate. Use this policy if you schedule weekly full backups and only need to keep backup data for one backup cycle plus a week.

- Month – This browse policy allows users to view and recover versions of files dating back at least a month. The Recover window will display versions for files backed-up for one full month plus a number of weeks. As a retention policy, Backup will maintain and track the backup volumes for one full backup cycle plus a month.
- Quarter – Use this policy if you need to keep backed-up data longer than a month. With this browse policy, the client can view and recover files for at least three months into the past. The retention policy tracks the backup volumes for at least three months plus one full backup cycle.
- Year – If you need to keep backed-up data online for several months, use the Year policy. For example, if your company requires ready access to information going back in time for at least three quarters, this is a good browse and retention policy. Realize, however, that Backup requires more disk space to maintain all the information online.
- Decade - This policy retains the entries in the server's indexes for ten years. It is useful for organizations which are required to retrieve individual files for very long periods of time.
  Your Backup server will require lots of disk space for the online indexes if you choose Decade for your browse policy. Depending upon how much data you are backing up, ten years of file index entries could take up gigabytes of disk space.

  It would make more sense to use Decade as the retention policy and use Quarter or Year as the browse policy. Backup can then track the backup volumes and the data on each one. You would always be able to retrieve data from an old backup volume using the save set recover feature if you needed to do so. Backup would still require disk space to maintain the media index, but it would be a much smaller amount of space using the Quarter or Year browse policies.

Preconfigured Pools

Backup is shipped with preconfigured pools and matching label templates. Each preconfigured volume pool has a set of unique preselected choices. If you do not choose a pool for your backups, they are automatically assigned to the preconfigured Default pool and are labeled using the Default label template.

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

The preconfigured volume pools have matching label templates. The Two Sided label template is for labeling optical media and is the only template that does not have a matching volume pool.

You can use the Default, Default Clone, Archive, and Archive Clone pools without making any additional selections in the Pools window. To use the other preconfigured pools, you must first complete the selections and choose Yes from the Enabled choices. A pool must be enabled in order for Backup to sort data to that pool.

The preconfigured pools are described below.

- Archive – for archiving client data only. This pool cannot be modified or deleted. The preconfigured settings are: Enabled – Yes, Label template – Archive, Pool type – Archive, Store Index entries – Yes. The are no selections for you to make for this pool.
- Archive Clone – for cloning archive data only. This pool cannot be modified or deleted. The preconfigured settings are: Enabled – Yes, Label template – Archive Clone, Pool type – Archive Clone, Store Index entries – No. There are no selections for you to make for this pool.
- Default – automatically used if you do not choose a pool. If you decide not to use the pools feature, Backup automatically places all of your backup volumes in this pool. The Default pool cannot be deleted or modified. The preconfigured settings are: Enabled – Yes, Label template – Default, Pool type – Default, Store Index entries – Yes. There are no selections for you to make for this pool.
- Default Clone – automatically used if you do not choose a pool for cloned data. If you decide not to use the pools feature, Backup automatically places all of your cloned backup volumes in this pool. The Default Clone pool cannot be deleted or modified. The preconfigured settings are: Enabled – Yes, Label template – Default Clone, Pool type – Backup Clone, Store Index entries – No. There are no selections for you to make for this pool.

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

- Full – use this pool for full backups only. This pool separates all of your full backups from the incremental and level backups. Using the Full pool provides you with the ability to easily track and separate your full backups from the incremental and level backups. Typically, you use this pool in conjunction with the NonFull pool. The preconfigured settings are: Enabled – No, Label template – Full, Levels – full, Pool type – Backup, Store Index entries – Yes.

- NonFull – use for any backups other than full backups. This pool includes all incremental and level backups. Use the NonFull pool to easily keep your incremental and level backups separate from the fulls. Typically you use this pool in conjunction with the Full pool. The preconfigured settings are: Enabled – No, Label template – NonFull, Levels – all level and incremental backups, Pool type – Backup, Store Index entries – Yes.
- Offsite – for volumes being stored offsite. The Offsite pool allows you to easily create a set of volumes to be stored offsite. If your onsite backup volumes are destroyed, you can still recover your valuable data with the volumes you have stored offsite. If you are also using the Full pool, you must disable it while you are sending data to the Offsite pool to ensure that all of the full backups will go only to the Offsite pool. The preconfigured settings are: Enabled – No, Label template – Offsite, Levels – full, Pool type – Backup, Store Index entries – No.

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

Preconfigured Label Templates

The preconfigured label templates shipped with Backup are: Archive, Archive Clone, Default, Default Clone, Full, NonFull, Offsite, and Two Sided. These are provided so that you can easily start labeling your backup volumes. There are also preconfigured volume pools with corresponding names (except Two Sided). The preconfigured volume pools automatically use the preconfigured label template with the same name.

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

The Archive label template is used only for clients that need to archive data. It has three fields each separated with a period. The first field contains the Backup server name, the second field is “archive,” and the third field contains a number.

For example:

```
server.archive.number

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

The Archive Clone template has three fields separated by periods. Backup uses this template for backup volumes belonging to the Archive Clone pool. The first field contains the name of the Backup server and the letter “c.” The second field is “archive,” and the third field contains a number. For example:

```
moon_c.archive.001
```

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

For example:

```
server.number

space.675
space.800
atlas.054
```

The Default Clone template has two fields separated by a period. Backup uses this template for backup volumes belonging to the Default Clone pool. The first field contains the server name and the letter “c.” The second field contains a number.

For example:

```
moon_c.002
```

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

For example:

```
label name.number

Full.076
NonFull.003
Offsite.120
```

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

```
server.number.side

phoenix.001.a
phoenix.001.b
```

Preconfigured Directives

Backup is shipped with preconfigured directives. Each directive covers a set of the most important and most useful backup instructions.

The preconfigured directives are listed below.

Unix standard directives – use for most of your backups and when you do not need one of the other specialized directives. This selection:

- applies the directive “+skip: core” to the *root* directory (/), thus skipping the backup of all *core* files
- contains a *swapasm* directive to back up the relevant information about all NFS-based and local swap files, but not the data in them
- contains a *mailasm* directive to ensure that your mail files are backed up, yet not marked as read, and *logasm* for directories containing log files.

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

DOS standard directives – use to back up your DOS clients

NetWare standard directives – use to back up your NetWare clients

NT standard directives – use to back up your NT clients.

NT with compression directives – use to compress the data on an NT client during backup.

Index directives – use to back up the online file index. This option is usually only used by the `savegrp` command.

Preconfigured Notifications

Backup is shipped with several notifications: Bootstrap, Cleaning cartridge expired, Cleaning cartridge required, Device cleaned, Device cleaning required, Index size, Log default, Migration attention, Migration completion, Registration, Savegroup completion, Tape mount request 1, Tape mount request 2, and Tape mount request 3.

Caution – Most of these notices alert you regarding important Backup events. For example, if a group of clients did not complete a nightly backup, Backup sends you a savegroup completion notice by electronic mail.

Registration

The registration notification sends a message to *root* notifying you that your Backup products are not properly registered. You will receive the registration notification once a day or each time you start Backup. The notification message includes related information about each of the Backup products that are not registered correctly.

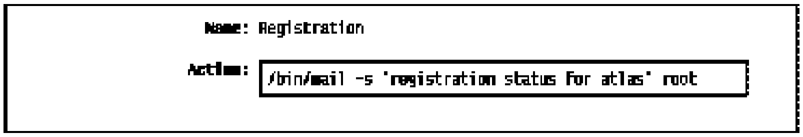

Name: Registration

Action: /bin/mail -s 'registration status for atlas' root

Log Default

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

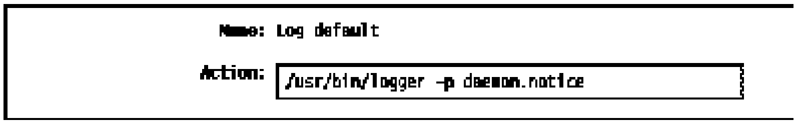

Name: Log default

Action: /usr/bin/logger -p daemon.notice

Index Size

Backup checks the size of its online indexes and sends a notification when it looks as if the indexes may run out of disk space. Backup automatically sends the electronic mail message to *root*.

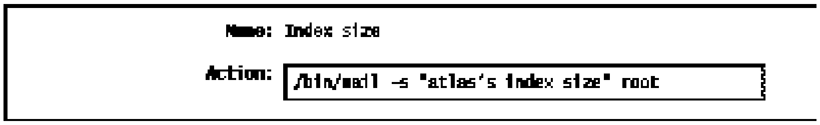

Name: Index size

Action: /bin/mail -s "atlas's index size" root

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

Savegroup Completion

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

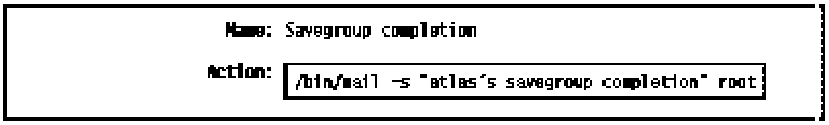

Name: Savegroup completion

Action: /bin/mail -s "atlas's savegroup completion" root

Backup Media Request Notices

When Backup needs backup media mounted for a backup, or a specific backup volume mounted to fill a recovery request, it displays a media request message in the Backup Administrator window. If no one fills the request, Backup sends another request after fifteen minutes. Backup sends a third request after another thirty-seven minutes, if no one fills the request.

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

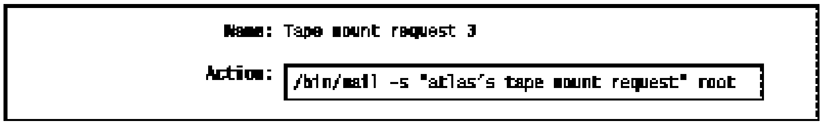

Name: Tape mount request 3

Action: /bin/mail -s "atlas's tape mount request" root

Summary

There are no rules for configuring Backup. The challenge is to understand how to best take advantage of the power and flexibility that Backup offers for your specific environment. You should start using Backup with the preconfigured schedules and policies and then undertake small experiments. As your network of systems grows larger, as there are more and more files to back up, and as users see the advantages of Backup's fast file recovery, you will need to continue making adjustments. Fortunately, Backup was designed to change and adapt as your needs expand.

Troubleshooting B

This appendix contains troubleshooting information that addresses common questions concerning operating and configuring Backup.

Checking the Backup Daemons

If you have trouble starting Backup, the daemons may not be running properly. To check the daemons, enter the following command for a SunOS system:

```
# ps ax | grep nsr
```

For other systems supported by Backup, including Solaris, use the following command:

```
# ps -ef | grep nsr
```

You should see a response similar to the following, showing these five daemons running.

```
111 ?  IW    0:10 /usr/etc/nsrexecd -s localhost
116 ?  S   176:15 /usr/etc/nsrd
158 ?  IW    2:48 /usr/etc/nsrmmdbd
159 ?  S    23:45 /usr/etc/nsrindexd
160 ?  IW<  16:07 /usr/etc/nsrmmd -n 1
```

If you discover that you need to start the Backup daemons enter these commands:

```
# cd /
# /etc/rc2.d/S95networker start
```

Command Not Found Message

If you are trying to run any of theBackup commands, such as `nwadmin`, `nwrecover`, `nwbackup`, etc., and you recieve the `command not found` error message, make sure that `/usr/sbin/nsr` and `/usr/bin/nsr` are in your PATH enviromnent variable.

Displaying Backup

If you enter the `nwadmin` command and the Backup Administrator window does not appear, the DISPLAY variable on your system may not be set correctly.

To set the DISPLAY variable correctly, follow these steps.

1. **Enter one of the following at the system prompt:**
   For a C shell:

   ```
   # setenv DISPLAY hostname:0.0
   ```

   For a Korn shell or a Bourne shell:

   ```
   # DISPLAY=hostname:0.0
   # export DISPLAY
   ```

   where hostname is the machine where the user initially logged on.

2. **Enter one of the following at the system prompt:**

   ```
   # xhost machine_name
   ```

   or

   ```
   # xhost +
   ```

   where machine_name is the machine that you are currently logged on to or the machine to which you will log on.

3. **Restart** `nwadmin`**.**

Renaming a Client

Backup maintains an index for every client it backs up. If you change the name of the client, the index for that client is no longer associated with the client, and the client will not be able to recover any files backed up under its old name.

To change the name of a Backup client, you must first delete the old client name, then add the new client name, and rename the directory which contains the corresponding index.

Follow these steps:

1. **Create a new client, using the new name. Use the same configuration choices for the new client that you used originally for the old client.**

2. **Delete the old client.**

3. **As *root* on the Backup server, shut down the Backup daemons, using the nsr_shutdown command.**

4. **Change to the directory containing the client index directory. For example:**

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

5. **Delete the new client index directory (which is empty). For example:**

```
# rmdir new_client_name
```

Note – If you do not remove the new client index directory, the old client index directory will be copied into the new client directory, as a subdirectory, with the old client name.

6. **As *root* on the Backup server, use the mv command to rename the client index directory. For example:**

```
# mv old_client_name new_client_name
```

7. **Restart the Backup daemons, using the the following command:**

```
# /etc/rc2.d/S95networker start
```

The daemon `nsrmmdbd` will rename all the instances of the save set under the old client name to the new client name.

Recover Access Issues

System administrators control client recover access by configuring the client. The Recover access list in the Clients window displays which machines can recover client files.

The following users have the ability to recover any files on any client:

root
operator
a member of the *operator* group

Other users can only recover files for which they have read permission, relative to the file mode and ownership at the time the file was backed up.

Files recovered by a user other than *root*, *operator*, or the *operator* group will be owned by that user.

Previewing a Backup

Every time you add a new client to Backup, it is a good idea to check if Backup can successfully back up the files for the new client. Use the Preview button in the Group Control window to see a "preview" of a group backup without actually backing up any files. You can also use the `savegrp -p` command at the system prompt to see a preview.

For example:

```
# savegrp -p group_name
```

shows you a backup preview of the clients assigned to the backup group *group_name*. If Backup cannot access a client in the backup group, you will see an error message. If this occurs, check the following items:

- Make sure `nsrexecd` is running on the client machine and that it lists the *hostname* of the server in the command line. To make sure that `nsrexecd` is running, use the UNIX command `ps` on the client. See "Installing the Client Software" in the appropriate chapter for your platform for more information on `nsrexecd`.

- Using nsrexecd is the best method for backing up clients over the network. If you choose not to use nsrexecd and the clients cannot find the Backup binaries, add the location of the Backup binaries to the Executable path field in the Client window for each client. In other words, if the default PATH setup for *root* or Remote user does not include the appropriate path to the Backup binaries, add them to the client configuration. Display these hidden attributes using the Details command in the View menu.

Halting a Network Backup

To stop running a network-wide backup using the Backup X Windows interface, use the Stop button in the Group Control window.

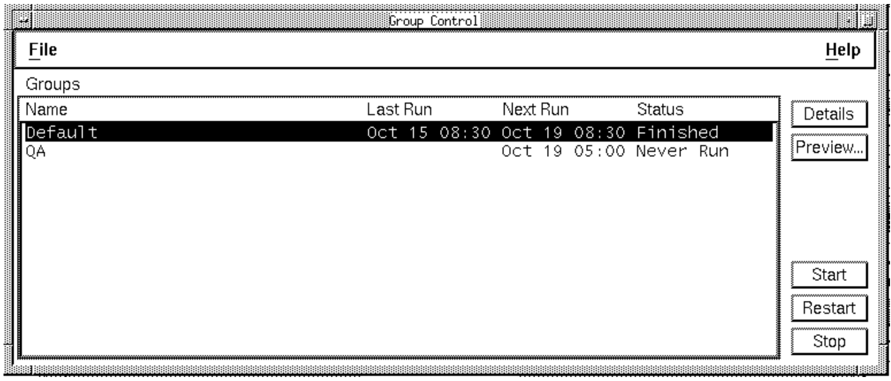

Name	Last Run	Next Run	Status
Default	Oct 15 08:30	Oct 19 08:30	Finished
QA		Oct 19 05:00	Never Run

The Stop button

The next network-wide backup will start as scheduled in the Start time field of the Groups window, or you may restart the backup by clicking the Restart button in the Group Control window.

Backup Media Capacity

Occasionally you will find that Backup marks backup volumes as “full” when they are not really full. (The Volumes window and the output from the mminfo -m command display the details of the backup volumes.)

Backup marks magnetic tape as "full" when it reaches the end of the tape or when there is a bad spot on the tape. For example, a backup tape that is reported as only "13% used" and marked as "full" has a bad spot on 13% of the length at the beginning of the tape. This tape can still be used for recoveries, but may not be used for any more backups.

If you see this "bad spot" behavior on many of the backup volumes, it may indicate the device needs cleaning or maintenance.

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

Determining Jukebox Capacity

To find out how much space is available in the jukebox, you may use either the Jukebox Mounting window or the `nsrjb` command. The Jukebox Mounting window displays all the media in the jukebox and the percentage used of each tape.

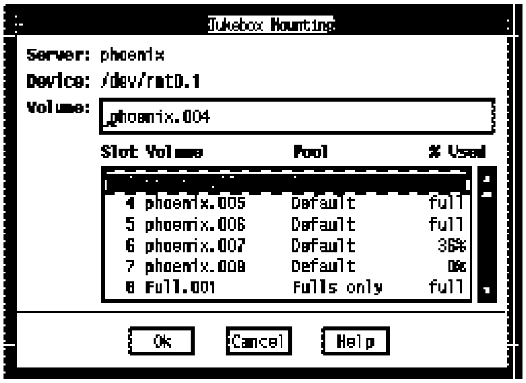

Server: phoenix
Device: /dev/rmt0.1
Volume: phoenix.004

Slot	Volume	Pool	% Used
4	phoenix.005	Default	full
5	phoenix.006	Default	full
6	phoenix.007	Default	36%
7	phoenix.008	Default	0%
8	Full.001	Fulls only	full

If you prefer to use the `nsrjb` command, follow the steps below:

1. Open a shell.

2. **Enter the** `nsrjb -v` **command at the system prompt:**

```
# nsrjb -v
```

Backup displays information about the backup volumes in the jukebox that looks similar to this:

```
Jukebox arc-db:
slot  volume      used   pool  mode
1:    moon.010                 Default
2:    moon.011    full         Default
3:    moon.012                 Default
4:    moon.013    full         Default
4 volumes, 2 less than 80% full.
2305 MB total capacity, 2200 MB remaining (5% full)

drive 1 (/dev/rmt/0hbn) slot 3: moon.012
```

Notice the information about the registered volumes, total capacity, and remaining capacity. This information tells you how much space is still available in the jukebox.

Savegroup Completion Messages

In the Notifications window, you configured Backup to mail the event notification about your savegroups. The Notifications window is preconfigured to mail the savegroup completion messages to *root*. Following are descriptions of error messages that may appear in the savegroup completion mail. Possible solutions are included.

Binding to Server Errors

Backup is designed to follow the client/server model. In a client/server model, servers provide services to the client through the Remote Procedure Call (RPC). These services live inside of long-lived UNIX processes, known as daemons.

For clients to find these services, the services must be registered with a registration service. When daemons start up they register themselves with the registration service. In UNIX the *portmapper* provides the registration service.

Backup servers provide a backup and recover service: they receive data from clients, store the data on backup media, and retrieve it on demand. If the Backup daemons are not running and a service is requested by `nwbackup`, `nwrecover`, or `mminfo`, for example, the following messages may appear in your savegroup completion mail:

```
“Server not available”
“RPC error, remote program is not registered”
```

These messages indicate the Backup daemons *nsrd*, *nsrindexd*, *nsrmmd*, *nsrmmdbd* may not be running.

To restart the *nsr* daemons, enter `nsrd` at the system prompt:

```
# nsrd
```

Saving Remote Filesystems

You may receive the following error message in your Savegroup completion notification when backing up a remote filesystem:

```
All: host hostname cannot request command execution
```

This means the `nsrexecd` on the client was not configured to allow the server *hostname* to back up its files.

You may also see this message:

```
All: sh: permission denied
```

This means `nsrexecd` is not running at all on the clients.

Make sure nsrexecd is running on the client machine and that it lists the server's *hostname* in the command line. To make sure that nsrexecd is running, use the UNIX command ps on the client. See “Installing the Client Software” in the appropriate chapter for more information on nsrexecd.

Using nsrexecd is the best method for backing up clients over the network. If you choose not to use nsrexecd, and the clients cannot find the Backup binaries, add the location of the Backup binaries to the Executable path field in the Client window for each client. In other words, if the default PATH setup for root or Remote user does not include the appropriate path to the Backup binaries, add the path to the client configuration. Display these hidden attributes using the Details command in the View menu.

File Changed During Backup

Backup backs up the image that is in the filesystem at the time it comes across the file. Backup will notify you that the file was changed during the backup in the Backup Status window and the savegroup completion mail. You can back up the file manually after it is through being used, or wait until the next incremental backup.

Cannot Print Bootstrap Information

If your bootstraps are not being printed, you may need to enter the printer name as a hidden attribute using the following steps:

1. **Open the Groups window and select Details from the View menu.**
2. **Enter the name of the printer you are using to print the bootstrap in the Printer field.**
3. **Click Apply to save your changes.**

Copy Violation

If you installed Backup on more than one server using the same Backup enabler code, you will receive the following messages in your savegroup completion mail:

```
--- Unsuccessful Save Sets ---
* quattro:/var save: error, copy violation - servers ‘quattro’ and
‘spim’ have the same software enabler code, ‘12345’ (13)
```

```
* quattro:/var save: cannot start a save for /var with NSR server
'quattro'
* quattro:index save: error, copy violation - servers 'quattro' and
'spim' have the same software enabler code, '12345'
* quattro:index save: cannot start a save for /usr/nsr/index/quattro
with NSR server 'quattro'
* quattro:index save: cannot start a save for bootstrap with NSR
server 'quattro'
* quattro:index /usr/etc/saveindex: bootstrap save of server's index
and volume databases failed
```

To complete a backup, you must kill the Backup daemons on both servers, de-install Backup from the extra server(s), and then restart the Backup daemons on one server.

1. **To kill the Backup daemons, log in to the Backup servers as *root* and enter the following command on all servers that have Backup installed:**

   ```
   spim# nsr_shutdown
   quattro# nsr_shutdown
   ```

2. **Then de-install Backup on the server(s) that you will not be using as a Backup server with the following command:**

   ```
   spim# pkgrm SUNWsbus1 SUNWsbuc
   ```

3. **Finally, restart the Backup daemons on one server with the following commands:**

   ```
   quattro# /etc/rc2.d/S95networker start
   ```

Maximum Filename Length

Backup supports a maximum filename size of 1024 characters. This is the same as the UNIX *svid* limitation.

Savegroup Completion Warning Messages

Occasionally the savegroup completion message includes one or more messages. These messages contain information that help the administrator understand why Backup performs certain tasks.

Below is one of the messages you might see:

```
quattro:/usr no cycles found in media db; doing full save
```

In this example, the filesystem, `/usr`, on the client `quattro` has no full saves listed in the media database. Therefore, despite the backup level pre-selected for that client's schedule, Backup will perform a full backup. This feature is important because it allows you to perform disaster recoveries for that client.

This message may also appear if the server and client clocks are not synchronized. To avoid this, make sure the Backup server and client:

- are in the same time zone
- have their clocks synchronized

The following savegroup message may also appear:

```
Backup_server:index Saving server index because server is not in
an active group
```

If your server belongs to a group that is not enabled Backup will, to avoid a long recovery process, save the server *bootstrap* information along with this group. As soon as possible you should enable the group to which your Backup server belongs.

Xview Errors

The following error message may appear when the nwadmin & command is executed:

```
Xlib: connection to “client:0.0” refused by server
Xlib: Client is not authorized to connect to Server
XView error: Cannot open display on window server: client:0.0
(Server package)
```

This indicates the client is not authorized to display Backup.

To correct this situation do the following at the client machine:

```
client% xhost <Networker_server>
```

Remotely log in to the Backup server and run the following command at the server prompt:

```
Networker_server% setenv DISPLAY client:0.0
```

For shells other than *csh* use the following commands:

```
# DISPLAY=client:0.0
# export DISPLAY
```

Moving Indexes

Because the index databases are holey files, cp will create a file that consumes more disk space than the original file. To move indexes execute the following command in the /nsr/index directory:

```
# uasm -s -i <client index directory name>|(cd target_dir; uasm
-r)
```

Recovering Files from an Interrupted Backup

You will not be able to recover files from a backup terminated by killing the Backup daemons because the media index was not updated before the daemons died. Consequently, Backup will not know on which volume the requested file is located.

Determining the Backup Server

If you start Backup from a remotely mounted directory, you may receive the following message:

```
Using server server_name as server for client_name.
```

Backup looks for the system that is the fileserver of a remotely mounted directory and uses the Backup server assigned to that system as the backup server. To bypass this message, start Backup from a local filesystem.

Using nsrexecd

The `nsrexecd` daemon runs on Backup client machines. This daemon provides a secure and restrictive way for Backup to start automatic backups on clients. The `nsrexecd` daemon allows you to restrict access to a select set of Backup servers.

When you run the `nsr_ize -i -c` or `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.

Security is increased by the use of a challenge/response scheme to ensure that only the Backup server is initiating connections, and not another program.

The file modified for each client type is shown in the table below. If you ever need to reconfigure `nsrexecd`, for example, to allow a different Backup server to back up the client, edit the appropriate file on the client, make the changes to the `nsrexecd` startup command (refer to the `nsrexecd` man page for a description of the command line configuration options), and then restart `nsrexecd`.

Make sure you enter the `nsrexecd` command in exactly the same way as it is listed in the boot-time file, complete with all command-line options listed in the boot-time file. Alternatively, you can use `nsr_ize -c -u` or `pkgrm SUNWsbuc` to de-install the client software, entering **no** whenever it asks you questions. Then, use `nsr_ize -c -i` or `pkgadd SUNWsbuc` and follow the instructions as if you were performing an NFS client install.

Operating System Type	**Boot-time file**
SunOS 4.1.x	*/etc/rc.local*
Solaris 2.x	*/ etc/rc2.d/S95networker*

Command Summary C

This appendix contains a list of the most commonly-used Backup commands entered at the system prompt and a table listing the Backup maintenance commands.

If you are not using the X Window System, or are using Backup from an ASCII terminal, the following commands are the ones you may use most often:

Use:	**To**:
mminfo	display information about the backup volumes and save sets
nsradmin	configure networker resources
nsrclone -c -N	perform a super-full backup
nsrck -F	compact the index size after purging or deleting backup volumes
nsrexecd	NetWorker client execution daemon
nsrhsmck	checks that stubs and file names are consistent
nsrinfo *client*	lists contents of client file index
nsrjb -l *volume_name*	load and mount a backup volume in the jukebox
nsrls	display information about the server online indexes
nsrmig	migrate files meeting defined criteria or rules

nsrmm -d *volume_name*	delete the backup volume named *volume_name*
nsrmm -d -P *volume_name*	purge the file index entries from both the file index and the media index for the backup volume named *volume_name*
nsrmm -l *volume_name*	label a backup volume with the name *volume_name*
nsrmm -m	mount the backup volume
nsrmm -m -l -R	recycle and mount backup volume
nsrmm -u	unmount the backup volume
nsrpmig	pre-migrates files meeting defined criteria or rules
nsr_shutdown -d	shut down the NetWorker daemons and processes dependent on those daemons
nsrwatch	display the character-based Backup status monitor
savegrp *group_name*	start the backup of the clients in the group *group_name*
savegrp -p *group_name*	see a preview of the backup of the group named *group_name*

The table below lists the Backup maintenance commands:

Command	**Manual Page**	**Description**
jb_config	**jb_config**(8)	jukebox resource configuration program
jbexercise	**jbexercise**(8)	jukebox diagnostic program
mmrecov	**mmrecov**(8)	command to recover the online indexes of a Backup server
nsrmm	**nsrmm**(8)	Backup media interface command
mminfo	**mminfo**(8)	Backup media and save set information reporting command

Command	Manual Page	Description
mmlocate	**mmlocate**(8)	accesses and manages the backup volume location information contained in the media database
mmpool	**mmpool**(8)	NetWorker media pool reporting command
nwadmin	**nwadmin**(8)	window-based display of Backup server status
(man pages only)	**nsr**(8)	guide to using Backup
nsradmin	**nsradmin**(8)	Backup character-based program for system administration
nsrck	**nsrck**(8)	Backup check and repair program for the server indexes
nsr_crash	**nsr_crash**(8)	how to use Backup to recover from crashes
nsrclone	**nsrclone**(8)	save set cloning command
nsrd	**nsrd**(8)	Backup server daemon
nsrexecd	**nsrexecd**(8)	NetWorker client execution daemon
nsrhsmck	**nsrhsmck**(8)	checks and corrects inconsistencies in hsm migrated files
nsrim	**nsrim**(8)	NetWorker index management program, usually invoked from the **savegrp** command
nsrindexasm	**nsrindexasm**(8)	module for saving and recovering Backup indexes
nsrindexd	**nsrindexd**(8)	Backup file index daemon
nsrinfo	**nsrinfo**(8)	file index reporting command
nsr_ize	**nsr_ize**(8)	SunOS 4.1.x installation and removal command
nsrjb	**nsrjb**(8)	jukebox control command
nsrls	**nsrls**(8)	lists the statistics of Backup index files
nsrmig	**nsrmig**(8)	migrates files for long term storage
nsrmmd	**nsrmmd**(8)	Backup media management daemon

Command	Manual Page	Description
nsrmmdbasm	**nsrmmdbasm**(8)	module for saving and recovering Backup media databases
nsrmmdbd	**nsrmmdbd**(8)	Backup media index daemon
nsrpmig	**nsrpmig**(8)	pre-migrates files for long term storage
nsrretrieve	**nsrretrieve**(8)	retrieves archived save sets
nsr_shutdown	nsr_shutdown(8)	stops the Backup processes
nsrwatch	**nsrwatch**(8)	displays the Backup server status from an ASCII terminal
recover	**recover**(8)	command to browse the online indexes and recover files from the Backup server
save	**save**(8)	command to save files to the Backup server
savefs	**savefs**(8)	command to save filesystems to the Backup server (see also **savegrp**)
savegrp	**savegrp**(8)	command to save the files of a group of Backup clients
scanner	**scanner**(8)	command to read the contents of a volume to recover from Backup server crashes
tapeexercise	**tapeexercise**(8)	command to exercise a tape drive in order to uncover problems
uasm	**uasm**(8)	Backup module for saving and recovering generic UNIX files

Note – To print the NetWorker man pages using a PostScript printer that has a *troff* filter, enter the following command:

```
% troff -t -man ‘nsr_man -l‘|lpr -t -P<printer_name>
```

This command may vary, depending on your system (for example, `lp` versus `lpr`), and the version of PostScript software you have.

Glossary

This glossary contains terms and definitions found in this manual. Most of the terms are specific to the Backup product.

1-9

These values appear for each calendar day in the Schedules window. Each number represents a backup level.

ASM

Application Specific Module. A program, that when used in a directive, specifies the way that a set of files or directories is to be backed up and recovered.

attribute

A piece of information that describes a Backup resource. It has a name and a list of values.

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.

Backup resources

Components of Backup software configuration information, described by a list of attributes and values.

client

A machine that accesses the Backup server to back up and recover files. Clients may be workstations, PCs, or fileservers with gigabytes of data.

cloning

The process by which NetWorker creates an identical copy of the data it backed up. Cloning applies to save sets.

command line

The shell prompt, where you enter commands.

compressasm

The Backup directive used for compressing and decompressing files.

daemon

A long-lived program that implements a service. For example, *nsrd* is a daemon that implements the Backup backup and recover service.

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.

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.

head (h)

Represents the beginning of a save set that spans two backup volumes.

heterogeneous

Heterogeneous networks are networks with systems of different platforms that interact meaningfully across the network.

holey

A directive used to efficiently back up files that do not have all of their data blocks allocated.

incremental (i)

A backup level in which only files that have changed since the last backup are backed up.

interoperability

The ability of software and hardware on multiple machines from multiple vendors to communicate meaningfully. (From *Internetworking with TCP/IP* by Douglas Comer.)

jukebox

A device which has the ability to randomly move media among various components located in the device including slots, media drives, media access ports, and transports. Jukeboxes automate the media loading, labeling, and mounting functions during backups.

level [1-9]

A backup level that backs up files that have changed since the last backup of any lower level.

machine

Any computer, including file or compute servers, diskfull workstations, or diskless workstations.

mailasm

The directive that adheres to spool mail file-locking conventions and resets a file's access time back to its pre-saved values, so users can still tell if new mail arrived after Backup backed up their mail.

media

Magnetic tape or optical disk 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.

media pool

The collection of backup volumes recognized and managed by Backup.

notice

A Backup event.

nsrhost

The logical *hostname* of the machine that is the default 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 device, and otherwise executes day-to-day tasks using Backup.

override

A backup level that takes place instead of the scheduled one.

preview

A look at what a Backup command will do without actually executing the command.

print

Send data to a printer.

recover

The Backup feature used to browse the server index and recover lost data from backup media to a client disk.

recycle

Make a NetWorker backup volume available for relabeling and new data because the data on the volume has passed both its browse and retention policies.

resources

See *Backup resources*

save

The Backup feature that backs up client files to backup volumes and maps the backups to the backup media.

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.

save set recover

A process by which Backup recovers an entire save set instead of individual files in the save set.

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.

skip

The directive to skip files during a backup. Useful for skipping files that do not require a backup.

super-user

A UNIX user with *root* privileges.

system administrator

The person normally responsible for installing, configuring, and maintaining Backup.

tail (t)

Represents the end of a save set that spans two backup volumes.

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.

Index