Solstice Backup™ 5.5 Disaster Recovery Guide

Sun Microsystems, Inc.
901 San Antonio Road
Palo Alto, CA 94303
U.S.A. 650-960-1300

Part No. 806-1133-10
September 1999, Revision A

DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID.

CETTE PUBLICATION EST FOURNIE "EN L'ETAT" ET AUCUNE GARANTIE, EXPRESSE OU IMPLICITE, N'EST ACCORDEE, Y COMPRIS DES GARANTIES CONCERNANT LA VALEUR MARCHANDE, L'APTITUDE DE LA PUBLICATION A REPONDRE A UNE UTILISATION PARTICULIERE, OU LE FAIT QU'ELLE NE SOIT PAS CONTREFAISANTE DE PRODUIT DE TIERS. CE DENI DE GARANTIE NE S'APPLIQUERAIT PAS, DANS LA MESURE OU IL SERAIT TENU JURIDIQUEMENT NUL ET NON AVENU.

Please
Recycle

Adobe PostScript

Contents

Figures

Preface

The *Solstice Backup 5.5 Disaster Recovery Guide* contains information on how to prepare for a disaster, whether you are running a single system or an entire network. Using Backup to back up your data is an excellent way to begin. However, you must also consider how to recover your data and systems if a disk crashes or an entire system is lost. If you back up your data regularly and implement the planning procedures outlined in this guide, you are well prepared to recover from a disaster.

For instructions about configuring and administering the administration program for Backup, refer to the *Solstice Backup 5.5 Administrator's Guide* that pertains to your platform.

About This Guide

This guide is for system administrators who are responsible for performing backups and recovers and for maintaining the safety of the data on the network.

This guide includes the following information:

- Descriptions of different types of disasters
- Platform-specific information that prepares you for a disaster
- Choices to consider while performing a disaster recovery
- Step-by-step instructions for recovering from a disaster for a Solaris backup server.

How This Book Is Organized

This book is organized as follows:

Chapter 1, "Introduction," contains concepts, procedures, and information that help you prepare for recovering data after a disaster.

Chapter 2, "Disaster Recovery," contains information that helps you determine which disaster recovery procedures you should follow for Backup for servers and clients.

Appendix A, "Client Recovery," provides special instructions necessary to perform a disaster recovery on Backup clients.

What Typographic Changes Mean

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

TABLE P-1 Typographic Conventions

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 *User's 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 P-2 Shell Prompts

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	`#`

CHAPTER 1

Introduction

This chapter contains concepts, procedures, and information that help you prepare for recovering data after a disaster. It is important that you develop a plan for recovering from a disaster where valuable data, a disk, or an entire system has been destroyed.

Different Types of Disasters

You can use Backup to recover from different types of disasters on your Backup clients. The degree of loss during a disaster can range from one or more files lost when a disk crashes to an entire computer system. The degree of severity of the disaster determines the procedures you need to perform to recover data to a Backup client.

Typically, the four types of disasters you might experience are as follows:

Case 1: Backup Client Data Corruption

Critical data other than the operating system (OS) or the Backup software is damaged or destroyed.

FIGURE 1-1 on page 2 illustrates a Backup client whose critial data is damaged or destroyed, but the operating system (OS) or the Backup software is not. In this example a Windows NT Backup client has two disks. The second disk containing the critical data has been destroyed by a disk crash. However, the disk containing the operating system and Backup software is still operational. To recover from this type of disaster, use the Backup `recover` program to recover the lost applications and data.

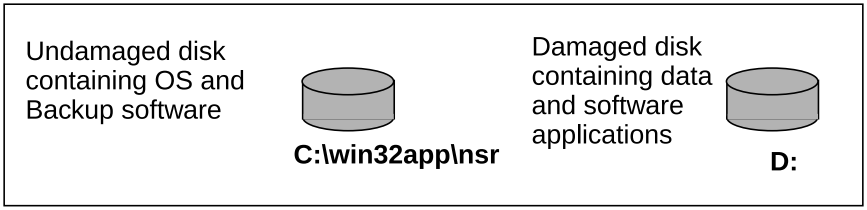
Undamaged disk containing OS and Backup software
C:\win32app\nsr
Damaged disk containing data and software applications
D:

FIGURE 1-1 Critical Data Is Lost on a Secondary Disk

Case 2: Backup Server Data Corruption

FIGURE 1-2 on page 2 illustrates a Backup server whose operating system and Backup software has been damaged or destroyed. This situation can occur on Backup clients and servers.

In this case, a Backup server has several physical disks. A power outage has corrupted the filesystem on Disk 0, which destroyed the operating system and Backup software residing on Disk 0. To recover from this type of disaster, you will need to replace the disk, reinstall the operating system and the Backup software. Then use Backup to recover the lost server configuration and any data that was lost when the filesystem system was corrupted.

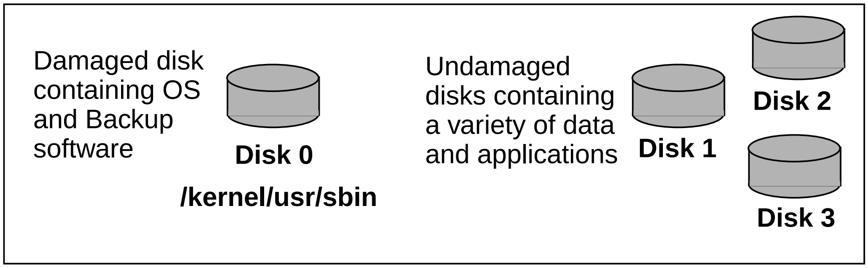
Damaged disk containing OS and Backup software
Disk 0
/kernel/usr/sbin
Undamaged disks containing a variety of data and applications
Disk 1
Disk 2
Disk 3

FIGURE 1-2 Damaged Disk Containing OS and Backup Software

Caution – In a situation where the operating system was destroyed, you must always reinstall the operating system, reinstall Backup, and then use Backup to recover the remainder of your data. You cannot recover data backed up by Backup without reinstalling the operating system and Backup software first.

Case 3: Backup Server and Client Data Corruption

FIGURE 1-3 on page 3 illustrates a directory containing the Backup software, online indexes and resource configuration files that have been damaged or destroyed. In this case, a disaster has corrupted the disk on the Backup server that contains the Backup indexes and configuration files. To recover from a disaster of this type, install Backup and then recover the contents of the bootstrap save set.

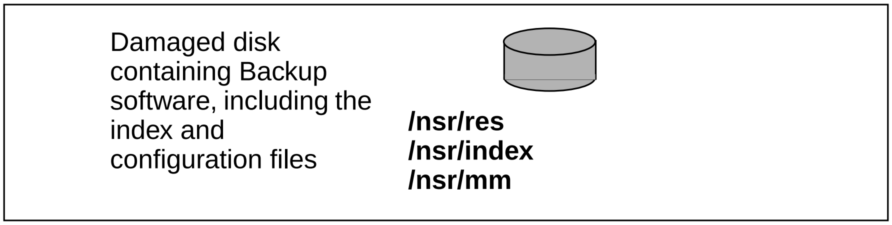
Damaged disk containing Backup software, including the index and configuration files
/nsr/res
/nsr/index
/nsr/mm

FIGURE 1-3 Disk Containing Backup Indexes Is Damaged

Case 4: Backup Server Is Unrecoverable

FIGURE 1-4 on page 3 illustrates a Backup server that has been destroyed. To recover from this type of disaster, you will need to recover all the data to a new system with *the same name*.

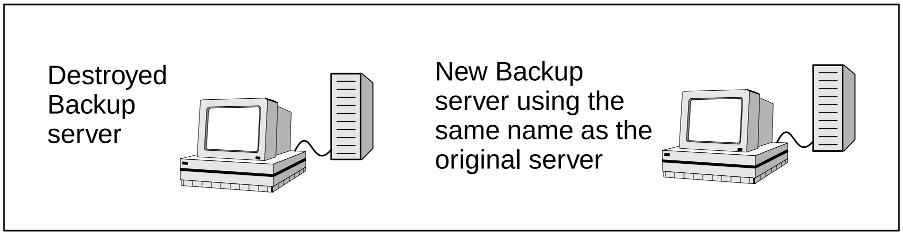
Destroyed Backup server
New Backup server using the same name as the original server

FIGURE 1-4 Backup Server Is Destroyed

Preparing for Disaster

It is important to develop a plan for recovering from a disaster on your Backup server and clients. Not only do you need to back up important data on a daily basis, you need to develop and test a plan for recovering your data should you experience a disk crash or loss of data. The more time and effort you invest in creating and testing your disaster recovery plan, the better prepared you will be should a disaster strike.

When you create a disaster recovery plan, consider the following strategy for an ideal backup solution.

- Perform full system backups at least once a week.
- Perform a level or incremental backup of the Backup clients at least once a day.
- Perform a backup of the Backup server’s bootstrap at least once a week.
- Maintain accurate records of your network and system configurations and maintain all your original software in a safe location, including the following:
  - Original operating system media and patches
  - Original Backup media
  - Device drivers and media device names
  - File system configuration
  - IP addresses and host names
  - NDS topology
  - Bootstrap information
  - Disk information
  - Software Enabler and authorization codes
- Keep an extra disk on hand as a “hot spare,” in case a disk failure occurs. Format it and verify that it works.

Disaster Recovery Requirements

Before performing any of the disaster recovery procedures outlined in this guide, keep in mind the hardware, operating system and software requirements outlined in “Disaster Recovery” on page 17. Fulfill requirements that are pertinent to your situation.

If you have an autochanger, you can either configure and use the autochanger during the recovery, or use the drive in the autochanger as a stand-alone device.

Important Information

Use the procedures in this section to collect bootstrap and disk configuration information necessary to perform a disaster recovery.

Bootstrap Information

During each scheduled backup of the backup server, Backup creates a special save set named *bootstrap* that is essential to perform a successful disaster recovery. The bootstrap contains the Backup server file index, media database, and configuration files.

Caution – Backup does not save the bootstrap information during a manual backup; Backup only saves it during a scheduled backup.

Backup prints or saves to a file the most recent bootstrap information that includes dates, locations, and save set ID numbers. The following is an example of the bootstrap information generated each time Backup performs a scheduled backup. Make sure you store the bootstrap printout or electronic file in a safe place.

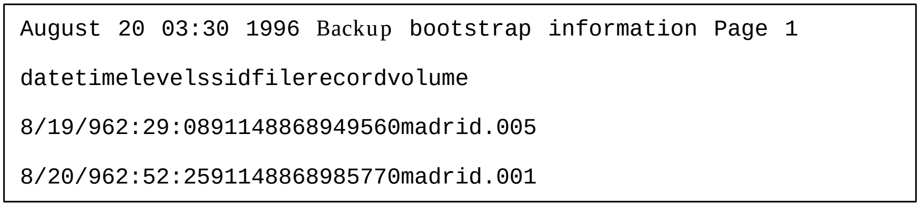

```
August 20 03:30 1996 Backup bootstrap information Page 1

datetimelevelssidfilerecordvolume

8/19/962:29:0891148868949560madrid.005

8/20/962:52:2591148868985770madrid.001
```

FIGURE 1-5 Bootstrap Information for Solaris Systems

The bootstrap displays a listing of the bootstrap save sets backed up for the past month.

For specific instructions about recovering Backup server indexes and configuration files, see “Disaster Recovery” on page 17.

You can also perform scheduled backups of the Backup server indexes by using the `savegrp` command. Using this command also sends the bootstrap information to a printer or electronic file. For example:

```
% savegrp -O -P printer_name -c client_name -s networker_servername
```

- The `-O` option indicates to `savegrp` that only the specified client’s index and the associated Backup server’s bootstrap should be backed up.
- The `-P` *printer_name* option specifies the name of the printer to print the bootstrap information.
- The `-c` *client_name* specifies the host name of the Backup client, whose index is to be backed up.
- The `-s` *networker_servername* specifies the host name of the Backup server, whose boostrap information is to be backed up.

To use the `savegrp` command, you must be `root` on the Backup server.

For information about printing or saving bootstrap data to a file, refer to “Bootstrap Notification” in the *Solstice Backup 5.5 Administrator’s Guide*.

Bootstrap Save Set ID

The most efficient way to recover the bootstrap is to ensure you save the bootstrap information prior to a disaster. However, if you do *not* have the information, you must scan the most recent backup volume to find the save set ID of the most recent bootstrap. Use the `scanner -B` command to locate valid bootstrap data.

Locating the Latest Bootstrap

Use the following steps to locate the save set ID of the most recent bootstrap:

1. **Place the most recent media used for scheduled backups in the Backup server device.**
2. **At the system or command prompt, change to the directory where the Backup binaries and executables are located, typically** `/usr/sbin`**.**
3. **Use the** `scanner -B` **command to locate the most recent bootstrap on the media, for example:**

   For SunOS™ systems:

   ```
   /usr/etc scanner -B /dev/nrst8
   ```

   For Solaris™ systems:

   ```
   /usr/sbin/ scanner -B /dev/rmt/0hbn
   ```

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

   ```
   scanner: scanning 8mm tape madrid.006 on \\.\Tape0
   scanner: Bootstrap 1148869 8/11/96 6:29:58 madrid.006,file 88
   ```

After you locate the bootstrap with the most recent date, run the `mmrecov` command to recover the Backup client indexes and resource configurations. For more information, “Disaster Recovery” on page 17.

Disk Information

As an additional precautionary step to help you recover from loss of critical data, find out how each Backup client’s disk on your network is partitioned and formatted, and print and save this information. If a disk is damaged or destroyed during a disaster, use this disk information to recreate the disk exactly as it was prior to the disk crash.

Caution – When you recreate your disk configuration, you will need to have partitions large enough to hold all the recovered data. Make the partitions at least as large as they were prior to the crash.

Use the `df` command to determine how the Backup server and client disks are partitioned and mounted. The example below shows the output generated by the `df` command. Use the appropriate operating system command to print disk partitioning information:

- For Solaris, use the `df` and `prtvtoc` command

```
# df
/              (/dev/dsk/c0t3d0s0 ):  114600 blocks   45421 files
/usr           (/dev/dsk/c0t3d0s6 ):  223706 blocks  220376 files
/proc            (/proc           ):       0 blocks     427 files
/dev/fd          (fd              ):       0 blocks       0 files
...
#
```

FIGURE 1-6 Example of `df` Command

The following `dkinfo` command examples provide information about how each disk is partitioned for a SunOS 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 `prtvtoc` command example in FIGURE 1-7 on page 8 displays information about the partitions for each disk on a Solaris system. The device name is the “raw” device corresponding to the device name used for the output from the `df` command.

```
# prtvtoc /dev/dsk/c0t0d0s0
* /dev/dsk/c0t0d0s0 partition map
*
* Dimensions:
*     512 bytes/sector
*      80 sectors/track
*      19 tracks/cylinder
*    1520 sectors/cylinder
*    3500 cylinders
*    2733 accessible cylinders
*
* Flags:
*   1: unmountable
*  10: read-only
*
* Unallocated space:
*       First     Sector    Last
*       Sector     Count    Sector
*     4152640      1520   4154159
*
*                          First     Sector    Last
* Partition  Tag  Flags   Sector    Count    Sector  Mount Directory
       0      2    00          0   3128160   3128159   /
       1      3    01    3128160   1024480   4152639
       2      5    00          0   4154160   4154159
#
```

FIGURE 1-7 Output of the `prtvtoc` Command

Disaster Recovery Pre-Configuration

When performing a disaster recovery, some scenarios require additional configuration *prior* to performing the disaster recovery.

If you have more than one media device, you must configure that device prior to beginning the disaster recovery. The default device created by a Backup install is always a 4mm device. If your device is not a 4mm device, you must delete this device from within Backup and create a correct device type.

After reinstalling Backup, and before starting `mmrecov`, edit the client resource for the server and set its browse and retention policies to a decade. This allows all of the server's records to be recovered by `mmrecov`. If you do not do this, all records are recovered, but any records that are more than a month old are then discarded, since the browse policy is one month by default.

If you relocated the server's file index path to a new location, you must edit the client index path resource *prior* to beginning the disaster recovery. The correct steps are to change the index path to match the path used prior to the disaster. Restart the services to alert Backup that the index path has changed. Finally, run `nsrck -c` from the `/nsr/bin` directory to create an empty file index into which the records are recovered.

Disaster Recovery Choices

You have several options for recovering the operating system. You also have the option of using an autochanger or a stand-alone drive. This section outlines the differences so you can decide which methods best suit your situation.

Restoring the Operating System

You can use one of two methods for recovering the operating system during a disaster recovery: complete or partial. When you perform a complete reinstallation, you install all operating system files and recreate any unique configurations that existed before you lost data or experienced a disk crash. To perform a partial reinstallation, install the minimum number of files and make the minimum number of configurations necessary for creating a fully operational networked system. Then, later, recover the remaining operating system and configuration files using Backup.

The following figure illustrates the steps for recovering from a disk crash where you lost the operating system, Backup software, and server indexes and configuration files. It also outlines the two choices you have for reinstalling the operating system.

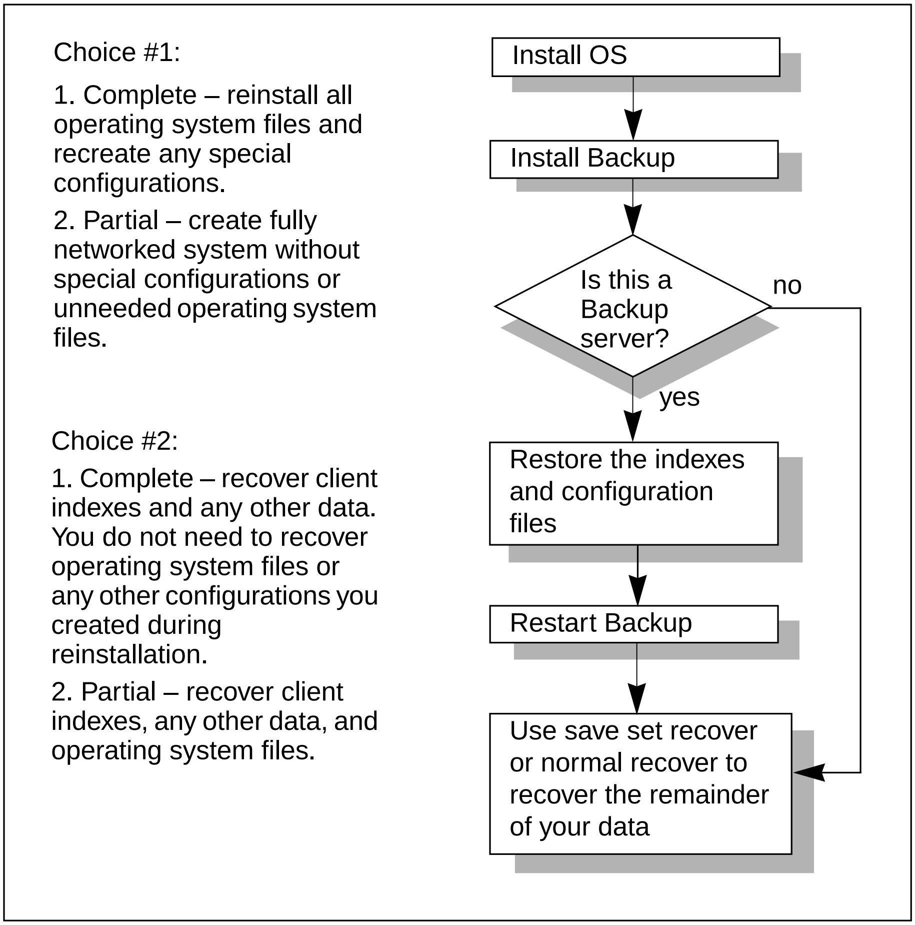

Choice #1:

1. Complete – reinstall all operating system files and recreate any special configurations.

2. Partial – create fully networked system without special configurations or unneeded operating system files.

Choice #2:

1. Complete – recover client indexes and any other data. You do not need to recover operating system files or any other configurations you created during reinstallation.

2. Partial – recover client indexes, any other data, and operating system files.

Install OS
Install Backup
Is this a Backup server?
no
yes
Restore the indexes and configuration files
Restart Backup
Use save set recover or normal recover to recover the remainder of your data

FIGURE 1-8 Recovering the Operating System

Performing a disaster recovery for Backup servers and clients is very similar, except on client systems you do not need to recover the server indexes and configuration files.

Complete Installation

In some cases, it might be faster to perform a complete reinstallation of the operating system, especially if you install the operating system from a CD and have very few unique configurations to recreate. Depending on the speed of your backup device and network, it might take longer to recover the remainder of your files and configurations using Backup during the disaster recovery procedure.

If you use a device with a default configuration that is not directly supported by the operating system, you also need to modify the device configuration files during installation:

- For Solaris systems, you might need to modify the `/kernel/drv/st.conf` file to support a DLT tape drive.
- For SunOS systems, modify the `/usr/sys/scsi/targets/st_conf.legato.h` file.

When you recover the remainder of your data, you can decide whether you want to replace the operating system files you just reinstalled with the operating system files backed up by Backup. If you want to guarantee that you have the same configurations prior to the disaster, replace the files and configurations you created during the installation.

Partial Installation

On the other hand, a partial installation might get your Backup server up and running more quickly, allowing you to continue with the disaster recovery. Later, you can recover the remainder of your operating system files using Backup. You will save time if you have a large number of clients and devices on the network that need to be configured; it will take you time to find IP addresses, hostnames, and recreate configurations.

Furthermore, if you wait to recover the remainder of the operating system files with Backup, the server, clients, and devices will be reconfigured exactly as they were prior to the disaster.

If you choose to do a partial install, you need to perform the following tasks:

- If necessary, select a domain for the system.
- Install the basic operating system files and device driver software.
- Ensure the system communicates properly over the network.

After a full or partial reinstallation of the operating system, you should verify that the tape device is functioning properly:

- Run the `mt` command, for example:

```
# mt -f $TAPE rewind
```

Recovery With Autochangers

This section explains how to use your autochanger during a disaster recovery where you have lost, at a minimum, the Backup server indexes and configuration files. Typically, the configuration files reside in the following directory `nsr/res`.

The configuration files include the `nsrjb.res` file, which contains autochanger configuration information.

This section assumes that you have lost the Backup server indexes and configuration files on the original server, or you are moving Backup and need to recover the existing index and configuration files to the new server.

The programs that recover the indexes and configuration files do not recognize autochangers. Consequently, you need to use the autochanger as if it were a stand-alone drive for this portion of the recovery. Use the autochanger's control panel to mount and unmount the necessary backup volumes.

After recovering the indexes and configuration files, all the original autochanger configuration files are back in place. You can now use the autochanger to recover the remainder of your data.

Caution – If you did not lose the server indexes and they are over 30 days old, you must reenable the server and autochanger to use the autochanger during a disaster recovery.

The rest of this section describes the issues that might influence your choices for using either the autochanger or the drive located inside the autochanger and how to recover the server's indexes and configuration files.

Autochanger Addition and Configuration

If you choose to recover with an autochanger, review the following issues about recovering data prior to restoring the server indexes and configuration files:

- If the autochanger has more than one drive, use the first drive for recovery.
- You cannot use the full functionality of the autochanger while restoring the server indexes and configuration files. `mmrecov` does not support autochangers. This command only supports stand-alone devices.
- The robotic device does not locate, load, and mount volumes automatically. You must use the Backup Mount and Unmount buttons and the autochanger control panel to mount and unmount volumes. If you use the autochanger control panel, Backup does not have a record of where the volumes have been moved. Inventory the autochanger contents after you complete the recovery.
- When you recover the server indexes and configuration files, recover the autochanger configuration files as they existed during the last backup, including the inventory of the autochanger. If you moved backup volumes inside the autochanger during the disaster recovery, the location of the volumes probably no longer matches the recovered inventory contents. After the recover operation, inventory the autochanger.

How to Recover With an Autochanger

To perform a disaster recovery with an autochanger, follow these steps:

1. **If necessary, reinstall the operating system and Backup software. During installation, use the same path name for the indexes that you previously used and backed up.**

2. **Run the** `jbconfig` **command to add and configure the autochanger.**

3. **Issue the** `nsrjb -vHE` **command.**

   This command resets the autochanger for operation, ejects backup volumes, reintializes the element status, and checks each slot for a volume. If the `-E` option is not supported for your autochanger, use the `sjiielm` program (for example, `/etc/LGTOusesi/sjiielm`) to initialize element status.

   If a volume is loaded in the drive, it is removed and placed into a slot. This operation might take a few minutes to complete.

   If you receive an error, typically the robotic device is having trouble finding a slot for a volume it has removed from the drive. Try moving some backup volumes around to make room for the volume, or, if possible, remove the volume from the robotic arm and manually place it in a slot.

4. **Locate your bootstrap data, either an electronic file or hard copy. With this information, determine which volumes are necessary for recovering the server indexes and configuration files.**

5. **Enter the** `nsrjb - I` **command to inventory the contents of the autochanger, and to help you determine whether the volumes required for recovering the bootstrap are inside the autochanger. The volume currently loaded in the drive probably contains the most current bootstrap.**

   If you want to speed up this process, issue the command with the `-S` flag and list only the slots where you think the required backup volumes reside. This saves you from having to inventory the entire autochanger contents. 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 1, 3, and 6), you must issue the `nsrjb -I -S` command separately for each slot. All the volumes currently loaded in the autochanger are marked with an asterisk because you have not yet recovered the media index.

6. **Load the appropriate volume by entering the following command:**

   ```
   # nsrjb -l -S slot -f device-name
   ```

   where *slot* is the slot where the first volume is located and *device-name* is the path name of the first drive. You can also use the Backup Mount button.

7. **Enter the `mmrecov` command. If the bootstrap spans across more than one volume, Backup prompts you to load another backup volume.**

8. **Enter the `nsrjb -u` command to unmount the volume after the indexes have been recovered. You can also use the Backup Unmount button.**

```
# nsrjb -u -S slot -f device-name
```

9. **Shut down Backup.**

10. **Rename the `/nsr/res` directory to `/nsr/res.orig`.**

11. **Rename the `/nsr/res.R` directory to `/nsr/res`.**

    When you recover and rename the `/nsr/res` files, you replace the configuration files you created when you reinstalled and configured the autochanger. This step ensures that you have all the configurations that existed on the last backup prior to the disaster.

12. **Restart Backup.**

    After the server indexes and configuration files are recovered, you have a fully functioning autochanger. Inventory the contents of your autochanger, especially if you manually moved volumes as part of the disaster recovery.

Recovery With a Stand-Alone Drive

If you choose to recover with a drive in the autochanger, review the following issues about recovering data prior to restoring the indexes and configuration files:

- If the autochanger has more than one drive, use the first drive for recovery.
- You must manually mount the backup volumes required for recovering the server indexes and configurations files.
- If you remove backup volumes from the autochanger cartridge used for recovering the Backup indexes and configuration files, put them back in the same slots when you finish.

How to Recover with a Stand-Alone Drive

To perform a disaster recovery using only the drive inside the autochanger for a Backup server, follow these steps:

1. **If necessary, reinstall the operating system and Backup software. If you need to reinstall the Backup software, use the same pathname for the indexes that you previously used and backed up.**

2. **Locate your bootstrap data, either an electronic file or hard copy. With this information, determine which volumes are necessary for recovering the server indexes and configuration files.**

3. **Manually mount the appropriate volume into the drive.**

4. **Enter the** `mmrecov` **command.**

5. **Shut down Backup.**

6. **Rename the original** `/nsr/res` **directory to** `/nsr/res.orig`**.**

7. **Rename the recovered** `/nsr/res.R` **directory to** `/nsr/res`**.**

8. **Restart Backup.**

9. **Issue the** `nsrjb -vHE` **command.**

   This command resets the autochanger for operation, ejects backup volumes, reintializes the element status, and checks each slot for a volume. If a volume is loaded in the drive, it is removed and placed into a slot. This operation might take a few minutes to complete.

10. **Inventory the autochanger contents by using the** `nsrjb - I` **command or use the Inventory command in the administrator program.**

    After you recover the server indexes and configuration files, you should have a fully functioning autochanger.

CHAPTER 2

Disaster Recovery

Use the information in this chapter to determine which disaster recovery procedures you should follow for Backup servers and clients. However, it is important that you read and follow the procedures in Chapter 1 before you need to recover from a disaster. Chapter 1 explains how to prepare for recovering from a disaster and defines basic terms, procedures, and concepts used throughout this guide.

This chapter includes procedures for recovering from the following kinds of disasters:

- Loss of a drive or partition that contains critical data other than the operating system or Backup software
- Loss of the operating system
- Loss of a drive or partition that contains Backup software, which typically includes the Backup indexes and configuration files
- Loss of an entire Backup server to the extent that you need to recover to a new system

It is difficult to provide step-by-step instructions for performing a disaster recovery for a specific situation, because every situation is unique. The examples in this chapter are designed to give you *general principles* to recover from a disaster and to help you understand the procedures.

Disaster Recovery Requirements

While performing any of the disaster recovery procedures included in this chapter, keep in mind the requirements listed in this section. Fulfill the requirements pertinent to the disaster recovery procedure you are following.

Hardware Requirements

If hardware becomes damaged or destroyed, adhere to the following guidelines to install and configure your new system hardware correctly:

- Ensure that the replacement disk is as large or larger than the original disk.
- When replacing the hardware, try to use the same controller, driver, and SCSI ID used prior to the disaster.
- Re-create the same size or larger disk partitions on the new disk/system.
- Format the disk partitions using the same formats used by the original disk.

Operating System Requirements

If the operating system is damaged or destroyed, adhere to the following guidelines when you reinstall Solaris:

- Reinstall the same version of Solaris.
- Use the same computer name, TCP/IP host name, and DNS domain name.
- Reinstall any operating system patches that existed before the disaster.
- Reinstall the device and SCSI drivers.
- Make sure all network protocols are working properly.
- After reinstalling Solaris, reboot your system, and log on as root. Make sure no error messages occur when you start up the system and that all the devices are recognized by the operating system.

Requirements for Reinstalling Backup

Adhere to the following guidelines to ensure successful reinstallation of Backup.

- Reinstall the same version of the Backup software.
- Reinstall Backup where it originally resided.
- Reinstall any patches that were installed prior to the disaster.
- For Backup servers, you must run additional procedures to retrieve the Backup server's indexes and configuration files. See “Recovering Backup Indexes and Configuration Files” on page 25 for information.
- For Backup clients or storage nodes see “How to Recover Backup Clients and Storage Nodes” on page 24.

Critical Data Recovery

The following example assumes the disk containing the operating system and Backup software is still operational, but another disk containing critical data has been lost. The example applies to both Backup servers and clients.

If the disk is damaged beyond repair, replace it with a new disk the same size or larger than the original disk. You need a disk large enough to hold all the data you plan to recover.

▼ How to Recover Critical Data

To recover the critical data, follow these steps:

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

2. **Use the saved disk partition information to re-create the disk partitions with the same structure as the original disk. See "Disk Information" on page 7.**

   If you did not save the disk information, it should still be available on the primary disk, which is operational in this case. Look at :

   - `/etc/vfstab` for a Solaris™ system
   - `/etc/fstab` for a SunOS™ system

3. **Use the output from the** `disk information` **command to make a filesystem for each raw partition you plan to recover, then mount the block partition. (Backup does not initialize or create filesystems; it recovers data into existing filesystems.)**

4. **Use the appropriate command to format the replacement disk.**

   For SunOS and Solaris systems, use `newfs` or `mkfs`.

Caution – Make sure the disk is no longer needed, because you will completely destroy the disk contents when you use `newfs` or `mkfs`.

- Run `newfs` on a SunOS system:

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

- Run `newfs` on a Solaris system:

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

5. **After creating and mounting all filesystems on the replacement disk, use the Save Set Recover feature in the `nwadmin` program or the normal recovery procedure in the `nwrecover` program to recover the files.**

   To determine which recovery method is best suited to your situation, refer to the *Solstice Backup 5.5 Administrator's Guide*.

Operating System Recovery

This example assumes a disk with the operating systems has been damaged or completely destroyed. You need to replace the damaged disk and reinstall the operating system. If the disk was not completely destroyed and either the operating system or Backup is still operational, use only the steps in this section that apply to your situation. Instructions in this section apply to Backup servers and clients, unless otherwise specified.

Caution – When you recover the operating system, you must do so in single-user mode from the system console, not from the X Window System.

▼ How to Recover the Operating System and Software

To prepare for recovering the operating system for either a server or client, follow these steps:

1. **Replace the damaged disk if necessary. Make sure the replacement disk is as large or larger than the original disk.**

2. **Use the saved disk partition information to re-create the disk partitions with the same structure as the original. See "Disk Information" on page 7.**

3. **Use the output from the `disk information` command to make a filesystem for each raw partition that you plan to recover, then mount the block partition. (Backup does not initialize or create filesystems; it recovers data into existing filesystems.)**

4. **Use the appropriate command to format the replacement disk:**
   - For SunOS and Solaris systems, use `newfs` or `mkfs`.

5. **Reinstall the operating system in the same location where it originally resided, using the original software and accompanying documentation. Use the same system name, TCP/IP hostname, and DNS Domain name used prior to losing the operating system.**

   You can choose to fully configure the operating system now, or you can install the minimum number of files and make the minimum number of configurations required to create an operational networked system. See "Restoring the Operating System" on page 9 for more information.

6. **Install and configure the SCSI controller and tape device drivers.**

7. **Reinstall the Backup software, using the original software and accompanying documentation. On a Backup client, you only need access to the Backup binaries. You can run Backup from the `nsr_extract` directory or NFS-mount the binaries from another system running Backup. Refer to the** *Solstice Backup Installation Guide and Release Notes* **for installation instructions. Reinstall any Backup patches you had installed prior to the disaster.**

   You might have several different releases of Backup software. The release must be equal to or later than the release used for the backups.

   Backup servers only: When you reinstall the Backup server software, Backup automatically locates indexes and configuration files if they are not corrupted.

   Backup clients only: The client system is now ready to recover its data from the Backup server.

You can also use the following method to access the Backup binaries for recovery. If you have another system running Backup that is similar to the system being recovered on the network, you can 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
```

8. **Install and configure the SCSI controller and tape device drivers.**

9. **Reboot the system, and log on as root.**

▼ How to Recover the Operating System

First create and mount all filesystems on the replacement disk. Then, to recover the Solaris operating environment, use the Save Set recover feature in the `nwadmin` program or the normal recovery procedure in the `nwrecover` program.

Backup Software Recovery

The following example for recovering the Backup binaries assumes a disk containing the Backup software has been damaged or completely destroyed. This example also assumes that the Solaris operating environment is installed and operating properly.

The set of instructions you need to follow in this section depends upon which system you lost (server, client, or storage node) and the extent of the damage. Refer to the following list of disaster recovery scenarios to determine which set of instructions apply to your situation.

If You Are Recovering ...	See The Following Sections ...
Backup clients and storage nodes	• "How to Prepare for Recovering Backup Software" on page 23 • "How to Recover Backup Clients and Storage Nodes" on page 24

If You Are Recovering ...	See The Following Sections ...
A Backup server that lost its indexes and configuration files	• "How to Prepare for Recovering Backup Software" on page 23 • "Recovering Backup Indexes and Configuration Files" on page 25 • "How to Rename the Backup Server Configuration Files" on page 30 • "How to Complete the Recovery of the Backup Server Data" on page 31
A Backup server from clone volumes	• "How to Prepare for Recovering Backup Software" on page 23 • "Recovering Data from Clone Volumes" on page 29 • "How to Rename the Backup Server Configuration Files" on page 30 • "How to Complete the Recovery of the Backup Server Data" on page 31
Backup to a new server	• "Recovery to a New Server" on page 32

▼ How to Prepare for Recovering Backup Software

Before you can restore Backup configuration files or indexes, you must reinstall the Backup software from the original media on the damaged system.

To reinstall the Backup software, follow these steps:

1. **Replace the damaged disk if necessary. Make sure the replacement disk is as large or larger than the original disk.**

2. **Use the saved disk partition information to re-create the disk partitions with the same structure as the original disk. See "Disk Information" on page 7.**

3. **Use the output from the disk information command to make a filesystem for each raw partition that you plan to recover, then mount the block partition. (Backup does not initialize or create filesystems; it recovers data into existing filesystems.)**

4. **Use the appropriate command to format the replacement disk.**
   - For SunOS and Solaris systems, use `newfs` or `mkfs`.

5. **Reinstall the Backup software, using the original software and accompanying documentation. On a Backup client, you only need access to the Backup binaries. You can run Backup from the** *nsr_extract* **directory or NFS-mount the binaries from another system running Backup. Refer to the appropriate** *Solstice Backup Installation Guide and Release Notes* **for detailed instructions. Reinstall any Backup patches you had installed prior to the disaster.**

Backup servers only: You do not need to reload the license enablers if the /nsr/res directory (configuration files) still exists. If the /nsr/res directory was destroyed, the license enablers are recovered when you recover the configuration files.

6. **If you had a link to another disk that contains the Backup indexes and configuration files (/nsr) or any other Backup directories located on another disk, re-create it now.**

   Backup servers only: If you back up to an autochanger and want to use it during the remainder of the disaster recovery, add and configure the autochanger with the jbconfig command after installing Backup. See “Recovery With Autochangers” on page 11 for more information.

   You can also use the following method to access the Backup binaries for recovery. If you have another system running Backup that is similar to the system being recovered on the network, you can NFS-mount the Backup binaries to the damaged system.

   For example:

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

   If this system is a server, continue with the disaster recovery by restoring the Backup indexes and configuration files. See “Recovering Backup Indexes and Configuration Files” on page 25 for instructions.

   If this system is a Backup client or storage node, see “How to Recover Backup Clients and Storage Nodes” on page 24 for instructions.

▼ How to Recover Backup Clients and Storage Nodes

To recover clients and storage nodes, you only need to reinstall the Backup client and storage node software, and then recover the configuration files using the nwrecover program.

Similar to Backup servers, Backup clients and storage nodes, each have a /nsr directory that contains special configurations created during the initial installation. During the disaster recovery procedure, you will recover the /nsr directory, which restores the clients and storage nodes to their status prior to the disaster.

To recover a Backup client or storage node, follow these steps:

1. **Log in as root.**
2. **Start the `nwrecover` program.**
3. **Click the Recover speedbar button to open the Recover window. Backup displays the system's directory structure in the Recover window.**
4. **Select and mark the Backup directory for recovery.**
5. **Click the Start speedbar button to begin the recovery.**
6. **Restart `nsrexecd`.**

   The Backup client or storage node should be restored to the status it had prior to the disk crash.

Recovering Backup Indexes and Configuration Files

These steps only apply to Backup servers, because only servers store and maintain the indexes and configuration files. Use the `mmrecov` command to recover the Backup indexes and configuration files that reside in the `/nsr` directory.

If the operating system and the Backup software were also destroyed, you must reinstall them prior to recovering the `/nsr` directory contents. See “Operating System Recovery” on page 20 and “Backup Software Recovery” on page 22.

When you use the `mmrecov` command to recover the `/nsr` directory, you recover the contents of three important directories:

- `/nsr/mm` (media manager) directory – contains the Backup media index that tracks all the Backup volumes and their save sets.
- `/nsr/index/server-name` directory – contains the server file index that contains a list of all the server files that were backed up prior to the disaster. The server file index includes information about the client indexes, for example, where they are located and how to recover them. Once the recovery of the server file index and media index is complete, use the `nwrecover` program to recover the client file indexes.
- `/nsr/res` directory – contains special Backup configuration files. The `nsr.res` file includes the list of clients that belong to the server, customized client configurations or selections, and device and registration information. The `nsrjb.res` file includes the location of the backup volumes in the jukebox and label template information. Unlike the indexes, the contents of this directory cannot be reliably overwritten while Backup is running. Therefore, `mmrecov` recovers the `/nsr/res` directory as `/nsr/res.R`, which you rename later.

Using the `mmrecov` Command

If the Backup indexes and configuration files that reside in the `/nsr/res` directory have been destroyed, you will need to use the `mmrecov` command to recover them. Information in this section only applies to Backup servers.

Use the `mmrecov` command to recover from the following situations:

- Corrupted indexes
- Incomplete index conversions
- Disk crash or other disaster

The `mmrecov` command prompts you for the bootstrap save set identification number (`ssid` or save set ID). If you followed the recommended procedures to prepare for loss of critical data, you have a copy of the bootstrap file (either hardcopy or an electronic file) with the name of the backup media you require and the bootstrap ssid.

Caution – Never run the `mmrecov` command from root (`/`); you can use any other directory.

In the following example, `ssid` 17851237 is the most recent bootstrap backup:

```
Jun 17  22:21 1997 madrid’s Backup bootstrap information
datetimelevelssidfilerecordvolume
6/14/9223:46:13full17826163480madrid.1
6/15/9222:45:15917836325870madrid.2
6/16/9222:50:349178465051340madrid.2
6/17/9222:20:25917851237520madrid.3
```

If you do not have this information, you can still recover the indexes by finding the ssid using the scanner `-B` command. See “Bootstrap Save Set ID” on page 6.

▼ How to Recover Backup Indexes and Configuration Files From Media

When the operating system and Backup software are in place, recover the Backup server’s online file and media indexes and configuration files from the backup media:

1. **Find the bootstrap information (required for the next two steps).**

2. **Mount the backup media that contains the most recent backup named bootstrap in a storage device.**

3. **Use the** `mmrecov` **command to extract the contents of the bootstrap:**

```
# mmrecov
```

The following output appears:

```
mmrecov: Using madrid.spain.com as server
```

Caution – The `mmrecov` command will overwrite the server's existing online file and media indexes. `mmrecov` is not used to recover Backup clients' online indexes; normal recover procedures may be used for this purpose. See the `mmrecov(8)` and `nsr_crash(8)` man pages for more details.

4. **Enter the latest bootstrap save set ID.**

5. **Enter the file number for the first file in the save set; if unknown enter zero (numeral, 0).**

6. **Enter the first record number; if unknown enter zero (numeral, 0).**

A message will appear, indicating the name of the volume that contains the associated save set ID, for example:

```
Please insert the volume on which save set id 17851237 started into
/disk1/file.tape.
```

7. **Once you have mounted the appropriate volume, the following messages appear:**

```
Scanning /disk1/file.tape for save set 17851237; this may take a
while...
scanner: scanning file disk file.tape on /disk1/file.tape scanner:
ssid 17851237: scan complete
scanner: ssid 17851237: 28 KB, 11 file(s)
/nsr/res/nsr.res
/nsr/res/nsr.res: file exists,
/nsr/res/nsrjb.res
/nsr/res/nsrjb.res: file exists,
/nsr/res/nsrla.res
/nsr/res/nsrla.res: file exists,
/nsr/res/
/nsr/mm/
/nsr/index/madrid.spain.com/
/nsr/index/
/nsr/
/
nsrmmdbasm -r /nsr/mm/mmvolume/
nsrindexasm -r /nsr/index/madrid.spain.com/db/
/disk1/file.tape: mount operation in progress
madrid.spain.com: 7 records recovered, 0 discarded.
/disk1/file.tape: mounted file disk file.tape
The bootstrap entry in the online index for madrid.spain.com has
been recovered. The complete index is now being reconstructed from
the various partial indexes which were saved during the normal save
for this server.
If your resource files were lost, they are now recovered in the
'res.R' directory.  Copy or move them to the 'res' directory, after
the index has been reconstructed and you have shut down the
daemons.  Then restart the daemons.
Otherwise, just restart the daemons after the index has been
reconstructed.
nsrindexasm: Pursuing index pieces of /nsr/index/madrid.spain.com/
db from madrid.spain.com.
Recovering files into their original locations.
nsrindexasm -r ./madrid.spain.com/db/
merging with existing madrid.spain.com index
madrid.spain.com: 753 records recovered, 0 discarded.
Received 1 matching file(s) from NSR server `madrid.spain.com'
Recover completion time: Wed Jan 28 08:37:38 1998
The index for `madrid.spain.com' is now fully recovered.
```

You can use Backup commands such as nsrwatch or nwadmin to watch the progress of the server during the recovery of the indexes and configuration files. Open a new window (shell tool) to monitor the recovery so that the mmrecov output is not displayed on top of the nsrwatch output.

```
Server: madrid.spain.com Wed Jan 28 08:53:54 1998
Up since: Wed Jan 28 08:35:15 1998
Version: NetWorker 5.5
Saves: 0 session(s)  Recovers: 1 session(s), 131 KB total
Device            type     volume
Disk1/file.tape    file     file.tape    reading, done
Messages:
Wed 08:35:11 server notice: started
Wed 08:35:22 index notice: completed checking 1 client(s)
Wed 08:36:44 /disk1/file.tape mount operation in progress
Wed 08:36:48 /disk1/file.tape mounted file disk file.tape
Wed 08:37:36 /disk1/file.tape mounted file disk file.tape
Wed 08:37:36 madrid.spain.com:/nsr/index/madrid.spain.com (1/28/
98)
starting read from file.tape of 131 KB
Wed 08:37:37 madrid.spain.com:/nsr/index/madrid.spain.com (1/28/
98) done
```

Recovering Data from Clone Volumes

For recovery from clone volumes, use the mmrecov command, as described in “Using the mmrecov Command” on page 26.

Select the bootstrap save set ID that includes the information associated with the cloned save set. The most recent bootstrap is the last save set listed in the bootstrap output.

In the following example, the ssid of the most recent bootstrap is 17851237. The clone of the bootstrap save set resides on madrid_c.3. The value for the file location is 6, and the value for the record location is 0.

```
Jun 17  22:21 1996  madrid’s Backup bootstrap information Page 1
datetimelevelssidfilerecordvolume
6/14/9623:46:13full17826163480madrid.1
6/14/9623:46:13full17826163120madridc.1
6/15/9622:45:15917836325870madrid.2
6/15/9622:45:15917836325240madridc.2
6/17/9622:20:25917851237520madrid.3
6/17/9622:20:259178512376 0madridc.3
```

After `mmrecov` recovers the bootstrap save set, it continues recovering the remainder of the server's client index to complete the recovery. The cloned bootstrap contains information about the original and cloned volumes.

Caution – To recover data from clone volumes, make sure that all the required clone volumes are mounted in attached devices at the time you run `mmrecov`. If some of the clone volumes are not online, `mmrecov` attempts to recover the server's client index from the original volume, not the clone volume.

Based on the preceding example of bootstrap output, the `madrid_c.1` and `madrid_c.3` volumes both need to be online. If the `madrid_c.3` volume is the only one online, `mmrecov` also requests `madrid.1`.

▼ How to Rename the Backup Server Configuration Files

Unlike the `/nsr/index` directory, the `/nsr/res` directory that contains the configuration files cannot be reliably overwritten while Backup is running. Therefore, `mmrecov` recovers the `/nsr/res` directory as `/nsr/res.R`.

When the `mmrecov` program is complete, it displays the following message:

```
The index for `server_name' is now fully recovered.
```

To complete the recovery of the configuration files, shut down Backup, rename the recovered `/nsr/res.R` directory to `/nsr/res`, and then restart Backup.

To complete the recovery of the Backup configuration files, follow these steps:

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

```
# nsr_shutdown
```

2. **Save the original** `/res` **directory as** `/res.orig`, **and rename the recovered file (**`res.R`**) to** `res`**.**

```
# mv res res.orig
# mv res.R res
```

3. **Restart Backup. When it restarts, the server uses the recovered configuration data in the recovered `/res` directory.**

```
# nsrexecd
# nsrd
```

4. **After you verify the Backup configurations are correct, you can remove the `res.orig` directory.**

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

▼ How to Complete the Recovery of the Backup Server Data

After you recover the server's indexes and configuration files, you can recover the remainder of the server data, including the client indexes, by using the `nwrecover` program.

Caution – Do not recover the kernel from the save set recover list. If you recover the kernel, you will not be able to reboot your system. If you recover `/boot` on a SunOS system, for example, you must use the `installboot` command to boot your system. You must always reboot a system after recovering a primary disk. For Solaris and DYNIX/ptx systems, you must also unmark the `/dev` directory from the save set recover list.

To recover the remainder of the Backup data, follow these steps:

1. **Log in as root.**

2. **Open the `nwrecover` program.**

3. **Select and mark the Backup directory for recovery.**

4. **Deselect the following directories and files before you recover the remainder of the server data:**

   - `/nsr/index/server-name` file – recovered when you ran the `mmrecov` command.

- /nsr/res and the /nsr/mm directories – recovered when you ran the mmrecov command. If you recover the /nsr/res directory and you used the autochanger to perform the disaster recovery, you will lose any special configurations you created when you added and configured the autochanger for recovery.

5. **Click the Start speedbar button to begin the recovery.**

6. **Restart nsrexecd and nsrd.**

Caution – You cannot boot from recovered versions of /etc/init or /bin/sh. Preserve the original init file and sh file by moving them to an alternate location before you recover the root filesystem (/). After you recover the root directory, overwrite the recovered versions of the init file and the sh file with the originals. It is also a good practice to preserve the original kernel (/vmunix for SunOS or /kernel/unix for Solaris) by copying it to an alternate location when you recover the root filesystem.

After you recover the server data, inventory the autochanger so Backup knows which slots contain which volumes.

The Backup server should now be restored to the same status it had prior to the disk crash.

Recovery to a New Server

This section describes a case where the original Backup server is damaged beyond repair, and Backup must be moved to a new server. This procedure assumes that you are not updating the operating system or the Backup software.

Caution – Do not make major changes to the operating system or Backup software at the same time you move to a new server.

If you want to make changes to the operating system or the Backup software, we strongly suggest that you configure the new server exactly like the original, using the same version of the operating system and Backup software. After configuring the new server, make sure the system is operational, perform a couple of successful backups, and then update or upgrade the operating system or the Backup software, one at a time.

To move Backup to a new server, use the same steps for recovering the operating system and Backup software, including the indexes and configuration files. Follow the instructions in these sections:

- "Operating System Recovery" on page 20
- "How to Prepare for Recovering Backup Software" on page 23
- "Recovering Backup Indexes and Configuration Files" on page 25
- "How to Rename the Backup Server Configuration Files" on page 30
- "How to Complete the Recovery of the Backup Server Data" on page 31

You should be aware of the following requirements when configuring and registering Backup:

- Use the same `hostname` for the new Backup server. You must use the same hostname, because the server indexes were created under the original Backup server name.
- Make sure the original server name is listed as an alias for the server in the Client window of the `nwadmin` program.
- If the new server has a different host ID, you need to reregister the Backup software.

After you move the Backup server to another system, you must recover the resource database (`nsr.res` file) to ensure that you carry over the same resource and attribute settings to the new Backup server.

If the new server has a different host ID, you have 15 days to reregister the software with Sun. Refer to the "Enabling and Registering Backup" section of the *Solstice Backup 5.5 Installation and Release Guide*.

Sun will send you a Sun Backup *Host Transfer Affidavit*, which you must complete and return. After Sun receives the signed affidavit, you are sent a new authorization code to enter into the Auth Code field of the Registration window.

After you successfully move your server, check the following:

- Verify that the server and all the clients are included in a scheduled backup.
- Schedule a full backup or use the `savegrp -O` command to back up the server and all clients as soon as possible. (Manual backups do not back up the server or client indexes.)
- Use the `nwrecover` program Recover window to make sure all the client indexes are browsable and, therefore, recoverable.

Recovering From an Expired or Bad Enabler

If you get "server disabled" error messages after you enter the permanent enabler and authorization codes, you may have an expired or bad enabler code on your system. This problem happens most often during upgrades, but also happens if the `nsr.res` file is corrupted.

The following solutions are available:

▼ How to Upgrade or Downgrade

1. **To upgrade, enter** `nsrcap -v -u` **and follow the prompts.**
2. **If upgrading does not work, downgrade. Enter** `nsrcap -v -d` **to remove the latest enabler.**
3. **If upgrading and downgrading do not work, use** `nsradmin` **to edit the resource files.**

▼ How to Edit the Resource Files

If you edit the resource files, make only the changes listed below, following the instructions exactly. The resource file is very sensitive to corruption. Corrupted resource files can cause Backup to hang, dump core, or in severe cases, cause the entire system to hang or dump core.

To edit the resource files:

1. **Shut down Backup. Enter** `nsr_shutdown -a` **and make sure that the shutdown completes sucessfully.**
2. **Make backup files of** `nsr.res` **and** `nsrjb.res` **(for example,** `nsr.res.bak` **and** `nsrjb.res.bak`**).**

3. **Start** `nsradmin` **from the command line, specifying the resource file with an** `-f` **option. For example:**

```
# nsradmin -f /nsr/res/nsr.res
```

4. **Once** `nsradmin` **has started, select visual, by typing** `visual` **or** `v`**. Visual mode requires that your** `TERM` **variable be set correctly and a termcap file has been created for your terminal type.**

5. **Select the license resource types: enter** `s`**, tab to the NSR license entry, then press [Return].**

6. **Delete all license entries: enter** `d`**, and then confirm by entering** `y` **when prompted.**

7. **Exit** `nsradmin` **by entering** `q` **twice, once to exit visual mode, then again to exit** `nsradmin`**.**

8. **Start Backup. Enter** `nsrexecd`**, then** `nsrd`**.**

9. **Enter** `nsrcap -v -c`**, as you would normally when entering enabler codes for the first time. After you enter the base enabler, enter the rest of the enablers and authorization codes through** `nwadmin`**.**

APPENDIX **A**

Client Recovery

This appendix provides special instructions necessary to perform a disaster recovery on Backup clients.

Windows 95 Client Recovery

To recover a Windows 95 client, follow these steps:

1. **Perform a full backup of your Windows 95 system.**
2. **Close down all applications.**
3. **Save the Windows 95 Registry to a floppy disk and to your hard drive**:
   a. **Open a DOS window and start** `Regedit`**.**

   b. **In the Registry Menu, provide a filename, select Export Range All, and save to floppy and hard disk.**
4. **Make a Windows 95 boot disk and copy the following files to this boot disk**:
   - `config.sys` (loading `himem.sys` and CD drive)
   - `autoexec.bat` (loading CD drive)
   - `himem.sys` and `format.exe`
   - Files needed for your CD drive and NIC card
5. **Print out your TCP/IP and NIC information.**
6. **To restore the system, install your new hard drive and follow these steps**:
   a. **Partition the new drive and format it (with** `/s` **option) using the** `format.exe` **program on the bootable disk you created.**

b. **Copy the** `config.sys`, `autoexec.bat`, `himem.sys`**, and CD drivers to your** `C:` **drive.**

c. **Reboot the system.**

d. **Change to the CD and install Windows 95 (using** `/setup`**).**

e. **Set up TCP/IP, using the information you printed out in Step 8.**

f. **Install Backup.**

g. **Start Backup on the Windows 95 system and perform a recover:**

- Mark all files to be recovered.

- Select overwrite/suppress.

- Shut down the system after the recovery ends.

- Restart the system.

7. **Open a DOS window and start** `Regedit`**.**
8. **Import the Registry from the floppy.**
9. **Shut down Windows 95 and start it up again.**

Your Windows 95 system should now be successfully recovered.

Windows NT Client Recovery

To recover a Windows NT client, follow these steps:

1. **Install the same version of Windows NT into the same directory as it was before.**
2. **Install the machine as a Workstation or Server, not as a Domain Controller.**
3. **Install the machine into a Workgroup, not a Domain.**
4. **Recreate all of the filesystems that were previously on the machine. The filesystems must be:**
   - The same filesystem type as before (NTFS, FAT, or HPFS).
   - At least the same size as before, to hold all of the backed up data.
5. **Reconfigure the networking as it was configured before, with the same host name and IP address.**
6. **Reboot the machine after reinstalling the Windows NT operating system.**

7. **If the machine had a Windows NT Service Pack installed, reinstall the same Service Pack, as well as any necessary patches for that Service Pack.**

Note – If you use Service Pack 4 for Windows NT 3.51, a patch is required to recover correctly. Before you recover the data from the system, apply Microsoft's patch for Service Pack 4. The patch is documented in Microsoft's "`Q149857 MoveFile_Delay_Until_Reboot`" and is available from:

```
ftp.microsoft.com/bussys/winnt/winnt-public/fixes/usa/nt351/
hotfixes-postSP4/SMSS-fix
```

Reboot the machine after reinstalling the Service Pack.

Recover Windows NT Data With Backup User

- Use the Backup User program to mark the Registry and all data for recovery.
- Do not recover any Service Pack patches that you applied during the reinstallation of Windows NT.
- Do not recover the Backup directory (default `C:\win32app\nsr`).
- For Backup for Windows NT version 4.3 and earlier, do not recover the user profiles:
  - On Windows NT 3.51, the user profiles are part of the `REGISTRY: saveset`. A user profile entry contains the first five letters of a username and some numbers.
  - On Windows NT 4.0, the user profiles are the `ntuser.dat` or `NTUSER.DAT` files located in `%SystemRoot%\Profiles` for each user's directory. The default path is `C:\winnt\profiles`. Unmark the `ntuser.dat` or `NTUSER.DAT` file in every subdirectory located in the Profiles directory.

Caution – Select "Overwrite" for files and Registry keys that already exist.

If you are running Backup for Windows NT version 4.4 or later, reboot the machine. The machine should now be restored to its previous state.

For NetWorker for Windows NT version 4.3 and earlier, follow these steps:

1. **Reboot the system.**

2. **Log on as Administrator.**

3. **Create a new user and assign the user to the Administrators group.**

4. **Log out, then log back on as this new user.**

5. **Recover the user profiles.**

Caution – On Windows NT 3.51 with Backup for Windows NT version 4.3 and earlier, there may be duplicate entries for some user profiles. One entry will have the first five letters of the user name and a three digit number while the other will begin identically, with the same five letters and three numbers, but will be followed by a Windows NT Security Identifier (SID). Only recover the profiles followed by a SID.

---

6. **Log on as Administrator and delete the user created in Step 3.**

7. **Reboot the system.**

   The Windows NT system should now be restored to its previous state.

Glossary

This glossary contains terms and definitions found in this guide. Most of the terms are specific to Backup products.

Administrators group Members of this Windows NT user group have all the rights and capabilities of users in other groups, plus the capability to create and manage all the users and groups in the domain. Only members of the Administrators group can modify Windows NT OS files, maintain the built-in groups, and grant additional rights to groups.

annotation A comment that you associate with an archive save set, to help identify that data later on. Annotations are stored in the media index for ease of searching and are limited to 1024 characters.

Application Specific Module (ASM) 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.

archive The process by which Backup backs up directories or files to an archive volume and then optionally deletes them to free disk space.

archive clone pool A volume pool composed exclusively of archive clone save sets.

archive pool A volume pool composed exclusively of archive save sets.

archive volume A tape or other storage medium used for Backup archives, as opposed to a backup volume.

ASM See *Application Specific Module.*

autochanger A mechanism that uses a robotic arm to move media among various components in a device, including slots, media drives, media access ports, and transports. Autochangers automate media loading and mounting functions during backups and recovers.

browse policy The policy that determines how long entries for your files remain in the online file index.

Backup Operators group Members of this Windows NT group have the capability to log on to a domain from a workstation or a server, back it up, and restore the data. Backup Operators can also shut down servers or workstations.

backup volume Backup media, such as magnetic tape or optical disk.

bootstrap Information that includes the server index, media index, and configuration files needed for recovering Backup after a disk crash.

client A machine that accesses the Backup server to back up or recover files. Clients can be workstations, PCs, or fileservers.

clone The process by which Backup makes an exact copy of saved data (save sets). Backup can clone individual save sets or the entire contents of a backup volume.

clone volume A duplicated volume. Backup can track four types of volumes: backup, archive, backup clone, and archive clone. Save sets of different types can not be intermixed on one volume.

cluster data Data shared by cluster servers that resides on a public disk.

cluster database Database that resides on a public disk that is shared by cluster servers.

cluster server Server that belongs to a cluster. Cluster servers typically provide services that include sharing data and providing failover services to other cluster servers in the group. Cluster servers must have both shared and public disks.

command line The shell prompt, where you enter commands.

compressasm A Backup directive used for compressing and decompressing files.

device The backup device (tape drive, optical drive, or autochanger) connected to the Backup server; it is used for backing up and recovering client files.

directive An instruction directing Backup to take special actions on a given set of files.

enabler codes Special codes provided by Sun that enable you to run your Backup software product.

file index A database of information maintained by Backup that tracks every file or filesystem backed up.

fileserver A machine with disks that provides services to other machines on the network.

filesystem 1. A file tree which is on a specific disk partition or other mount point. 2. The entire set of all 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.

grooming The process of removing files after a successful archive.

group A client or group of clients that starts backing up their files at a designated time.

heterogeneous A type of network with systems of different platforms that interact meaningfully across the network.

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.

carousel A tray or tape cartridge that holds multiple backup volumes.

level [1-9] A backup level that backs up files that have changed since the last backup of any lower level.

machine Any computer, including file servers or compute servers, diskfull workstations, or diskless workstations.

media Magnetic tape or optical disks used to back up files.

media database A database of information maintained by Backup that tracks every backup volume.

media manager The Backup component that tracks save sets to backup volumes.

Backup client A machine that can access the backup and recover services from a Backup server.

Backup server The machine on a network running the Backup software, containing the online indexes, and providing backup and recover services to the clients on the same network.

notice A response to a Backup event.

nsrhost The logical *hostname* of the machine that is the Backup server.

online indexes The databases located on the server that contain all the information pertaining to the client backups and backup volumes.

operator The person who monitors the server status, loads backup volumes into the server devices, and otherwise executes day-to-day tasks using Backup.

override A backup level that takes place instead of the scheduled one.

pathname Instructions for accessing a file. An *absolute pathname* tells you how to find a file beginning at the root directory and working down the directory tree. A *relative pathname* tells you how to find the file starting where you are now.

preconfigured Existing selections or configurations for different Backup features.

private disk Local disk not shared by the other servers in a cluster.

public disk See *shared disk.*

recover The Backup command used to browse the server index and recover files from a backup volume to a client's disk.

recycle A volume whose data has passed both its browse and retention policies and is available for relabeling.

Registry A database of configuration information central to Windows NT operations. The overall effect centralizes all Windows NT settings and provides security and control over system, security, and user account settings.

retention policy A policy that determines how long entries are retained in the media index and, thus, are recoverable.

retrieval The process of locating and copying back files and directories that Backup has archived.

save The Backup command that backs up client files to backup volumes and makes data entries in the online index.

save set A set of files or a filesystem backed up to backup media using Backup.

save set identification (save set ID or ssid) An internal identification number assigned to a save set by Backup.

scanner A Backup command used to read a backup volume when the online indexes are no longer available.

server A machine on a network running the Backup software, that contains the online indexes and provides backup and recover services to the clients on a network.

shared disk Hard disk shared by the servers in a cluster. The shared disk typically stores the quorum database.

shell prompt A cue for input in a shell window where you enter a command.

skip (s) A backup level in which files are skipped and not backed up.

stand-alone device A backup device that contains a single drive for backing up data. Stand-alone devices cannot store or automatically load backup volumes.

Storage Management Services (SMS) An operating system interface that provides functionality specific to storage management products.

system administrator A person typically responsible for installing, configuring, and maintaining Backup.

user People who use Backup from their workstations to back up and recover files.

volume Backup media, such as magnetic tape or optical disk.

volume ID The internal identification assigned to a backup volume by Backup.

volume name The name you assign to a backup volume when it is labeled.

volume pool A feature that enables you to sort backup data to selected volumes. A volume pool contains a collection of backup volumes to which specific data has been backed up.

Index