

Novell NetWare[®] 6.0

www.novell.com

June 30, 2004

STORAGE MANAGEMENT SERVICES
ADMINISTRATION GUIDE



Novell[®]

Legal Notices

Novell, Inc. makes no representations or warranties with respect to the contents or use of this documentation, and specifically disclaims any express or implied warranties of merchantability or fitness for any particular purpose. Further, Novell, Inc. reserves the right to revise this publication and to make changes to its content, at any time, without obligation to notify any person or entity of such revisions or changes.

Further, Novell, Inc. makes no representations or warranties with respect to any software, and specifically disclaims any express or implied warranties of merchantability or fitness for any particular purpose. Further, Novell, Inc. reserves the right to make changes to any and all parts of Novell software, at any time, without any obligation to notify any person or entity of such changes.

This product may require export authorization from the U.S. Department of Commerce prior to exporting from the U.S. or Canada.

Copyright © 1999-2003 Novell, Inc. All rights reserved. No part of this publication may be reproduced, photocopied, stored on a retrieval system, or transmitted without the express written consent of the publisher.

U.S. Patent No. 5,157,663; 5,349,642; 5,455,932; 5,553,139; 5,553,143; 5,572,528; 5,594,863; 5,608,903; 5,633,931; 5,652,854; 5,671,414; 5,677,851; 5,692,129; 5,701,459; 5,717,912; 5,758,069; 5,758,344; 5,781,724; 5,781,733; 5,784,560; 5,787,439; 5,818,936; 5,828,882; 5,832,274; 5,832,275; 5,832,483; 5,832,487; 5,859,978; 5,870,561; 5,870,739; 5,873,079; 5,878,415; 5,884,304; 5,893,118; 5,903,650; 5,903,720; 5,905,860; 5,910,803; 5,913,025; 5,913,209; 5,915,253; 5,925,108; 5,933,503; 5,933,826; 5,946,002; 5,946,467; 5,956,718; 5,956,745; 5,964,872; 5,974,474; 5,983,223; 5,983,234; 5,987,471; 5,991,810; 6,002,398; 6,014,667; 6,016,499; 6,023,586; 6,029,247; 6,052,724; 6,061,726; 6,061,740; 6,061,743; 6,065,017; 6,081,774; 6,081,814; 6,094,672; 6,098,090; 6,105,062; 6,105,069; 6,105,132; 6,115,039; 6,119,122; 6,144,959; 6,151,688; 6,157,925; 6,167,393; 6,173,289; 6,216,123; 6,219,652; 6,233,859; 6,247,149; 6,269,391; 6,286,010; 6,308,181; 6,314,520; 6,324,670; 6,338,112; 6,345,266; 6,353,898; 6,424,976; 6,466,944; 6,477,583; 6,477,648; 6,484,186; 6,496,865; 6,510,450; 6,516,325; 6,519,610; 6,532,451; 6,532,491; 6,539,381; 6,560,615; 6,567,873; 6,578,035; 6,591,397; 6,609,158; 6,615,350; 6,629,105; 6,629,132; 6,647,408; 6,651,242 & RE37,178. Patents Pending.

Novell, Inc.
1800 South Novell Place
Provo, UT 84606
U.S.A.

www.novell.com

NetWare 6.0 Storage Management Services Administration Guide

June 30, 2004

Online Documentation: To access the online documentation for this and other Novell products, and to get updates, see www.novell.com/documentation.

Novell Trademarks

eDirectory is a trademark of Novell, Inc.

GroupWise is a registered trademark of Novell, Inc. in the United States and other countries.

Internetwork Packet Exchange and IPX are trademarks of Novell, Inc.

NetWare is a registered trademark of Novell, Inc., in the United States and other countries.

NetWare Loadable Module and NLM are trademarks of Novell, Inc.

NetWare SMS is a trademark of Novell, Inc

NetWare Storage Management Services and SMS are trademarks of Novell, Inc.

Novell is a registered trademark of Novell, Inc. in the United States and other countries.

Novell Cluster Services is a trademark of Novell, Inc.

Novell Directory Services and NDS are registered trademarks of Novell, Inc. in the United States and other countries.

Novell Storage Services is a trademark of Novell, Inc.

Sequenced Packet Exchange and SPX are trademarks of Novell, Inc.

SPX/IPX is a trademark of Novell, Inc.

Third-Party Trademarks

All third-party trademarks are the property of their respective owners.

Preface

Novell[®] Storage Management Services[™] (SMS[™]) is a collection of software programs that provides backup and restore services. These services are performed by a collection of components that are independent of operating systems and hardware.

You can backup SMS Targets such as Novell eDirectory[™], binderies, the file system, cluster-enabled pools or an individual workstation's hard disk onto a media that can be stored off-site. If you have a hardware failure, natural catastrophe, corrupted data, or incorrectly deleted or changed data, you can recover a previous version of the data.

Documentation Conventions

In this documentation, a greater-than symbol (>) is used to separate actions within a step and items in a cross-reference path.

Also, a trademark symbol ([®], [™], etc.) denotes a Novell trademark. An asterisk (*) denotes a third-party trademark.

1

Overview

Backup services is an essential piece of the Novell® Storage Management Services™ (SMS™) software, a collection of services that provides backup, restore, and data migration. These services are performed by various components and are independent of operating systems and hardware.

This chapter discusses the following:

- ◆ “Backup Services” on page 7
- ◆ “SMS Components” on page 8
- ◆ “Restore Services” on page 22
- ◆ “Session Files” on page 26

Backup Services

Storage Management Services (SMS) allows you to back up SMS targets such as eDirectory™, the file system, cluster-enabled pools, or hard disks on individual workstations to media such as a tape drive for off-site storage, and gives you a periodic view (daily, weekly, monthly) of your data. Then in case of hardware failure, natural catastrophe, corrupted data, or incorrectly deleted or changed data, you can retrieve a previous version of the data.

Backup services provides information on supported devices, the SMS architecture, memory requirements, types of backup offered, customizable options, schedules, how to back up eDirectory, how to back up the file system (both traditional NetWare® file system and Novell Storage Services™), how to back up cluster-enabled pools, and how to use Target Service Agents (TSA).

The following topics are discussed in this section:

- ◆ “SMS Components” on page 8
- ◆ “Backup Types” on page 11
- ◆ “Customizing Your Backup” on page 12
- ◆ “Keeping a Backup Logbook” on page 15
- ◆ “Preparing to Back Up” on page 15
- ◆ “Backups and eDirectory” on page 16
- ◆ “Backups and the File System” on page 19
- ◆ “Backups and Clusters” on page 20
- ◆ “Session Files” on page 26

SMS Components

- ◆ **Storage Management Engine (SME)** for backup and restore operations. Novell provides the SBCON utility as a basic SME for NetWare.

See “[Storage Management Engine \(SME\)](#)” on page 9 for more information.

- ◆ **Storage Management Data Requester (SMDR)** for passing communication between the backup program (see “[Storage Management Data Requester](#)” on page 9) and the TSA software (see “[Target Service Agents \(TSAs\)](#)” on page 21.)
- ◆ **Storage device interface** is used to pass information between the SME and the **storage device**.
- ◆ **Device drivers** are used to control the behavior of the storage devices.
- ◆ **Target Service Agents (TSAs)** pass requests and commands between the SME and server or eDirectory database, and prepare the data for the SME. TSAFS.NLM must be loaded on the server where the data is to be backed up.


NOTE: The TSA600 has been replaced by TSAFS. From this release onwards, we’ll will not be supporting TSA600.

See “[Target Service Agents \(TSAs\)](#)” on page 21 for more information.

- ◆ **TSAProxy** is used to register the workstation with the **host server**. TSAProxy also identifies and keeps track of the stations waiting to be backed up. It receives *I am here* messages from workstations available to be backed up. The TSAProxy keeps the names of these workstations in an internal list and displays the list, allowing you to select a **target** for a backup or restore procedure.

The SBCON process involves two machines:

- ◆ The host server is typically the NetWare server running the backup program. The SMS components will include the backup engine, the SMDR, the Storage Device Driver, the tape device, and the required Target Service Agents.
- ◆ A target server is the NetWare server that contains the data to be backed up or restored. Depending on the data to be backed up, the appropriate Target Service Agent must be loaded.



SBCON can also be loaded and used on one machine.

SBCON uses an application on the host server to communicate with modules on target devices. The application reads the information from the target device and sends it to a storage medium, such as a tape drive.

Supported Storage Devices and Drivers

SBCON supports 0.25-inch, 4mm, and 8mm storage devices. If you are using 4mm tape, use only DDS (Digital Data Storage)-certified, computer-grade tapes.

IMPORTANT: To ensure reliable operations, pretest all media storage devices that are not Novell certified with the appropriate NetWare device driver and SBCON backup and restore utility.

Use the driver files recommended by your hardware manufacturer.

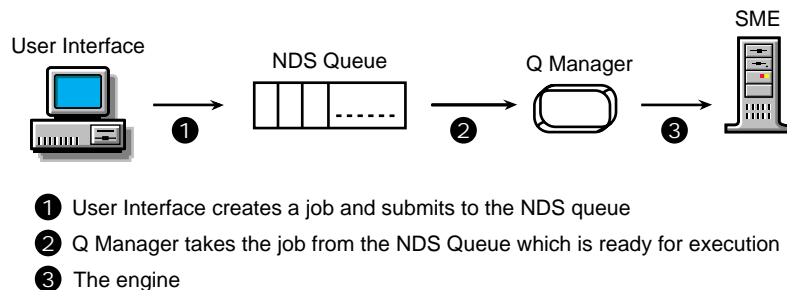
Storage Management Engine (SME)

The Storage Management Engine (SME) is central to the SMS architecture. The SME communicates with the network clients to back up and restore information.

SBCON has three modules:

- ◆ User interface - Creates a job and submits to the NDS Queue.
- ◆ Q Manager (QMAN) - Takes the job from the NDS Queue which is ready for execution.
- ◆ Engine - Completes the job.

Q Manager facilitates multiple job scheduling plus other features. Loading Q Manager automatically loads the backup engine. The user interface can be loaded after you load the Q Manager. See [“Setting Up” on page 29](#) for information on Q Manager.



Storage Management Data Requester

The Storage Management Data Requester (SMDR) is the communication module in the SMS architecture. It provides transparent access to SMS services in an intranet as it allows access to local or remote SMS services. The SMDR APIs are used by SBCON and other third-party applications as well. SMDR uses TCP Port Number 413.

Features of SMDR

The features of the SMDR 6.00 include the following:

Protocol Independence: SMDR 6.00 is protocol independent and does not depend on Sequenced Packet Exchange™ (SPX™) or Internetwork Packet Exchange™ (IPX™) protocols. From NetWare 5.1 onwards, the requester also uses TCP/IP for communicating with other SMDRs. Although

SMDR 6.00 can be configured to support TCP/IP, SPX/IPX™, or TCP/IP and SPX/IPX, both protocols are supported by default. SMDR versions prior to 5.00 use the SPX protocol.

If cluster-enabled pools are to be backed up or restored, use SLP as the discovery mechanism.

The protocols can be specified in the configuration file, SMDR.CFG (see [“Using the SMDR Configuration File” on page 10](#) for more information).

NDS Registration and Name Resolution: SMDR creates an SMS Remote Procedure Call (RPC) object in eDirectory. The default tree of the server (the tree in which the server is present) is used for eDirectory registration.

The SMS RPC object is defined with the following attributes:

- ◆ Common Name - Name of the server on which the SMDR is located
- ◆ Status - Indicates whether the SMDR is active or inactive
- ◆ Protocol - List of protocols for which the SMDR has listeners
- ◆ Version - Version number of the SMDR
- ◆ Services - List of services registered with the SMDR

The SMDR creates an instance of this RPC class at the SMDR Context location in the server's default tree. The SMDR Context is specified in the SMDR.CFG file, which can be edited at any given time.

Multiple SMDRs are grouped together to reduce the search scope in eDirectory. A SMDR Group object defines this search scope. This group represents an instance of a predefined group class in the eDirectory schema. Any number of such groups can exist in eDirectory. The SMDR can become a member of one or more groups by registering its object's (SMS RPC object) context.

When SMDR requires name resolution, it searches all members of the SMDR Group at SMDR Group Context. The SMDR Group Context and SMDR Group are specified in the SMDR.CFG file.

Discovery and Name Resolution Using SAP: SMDR can be configured to use Service Advertising Protocol (SAP) for locating other SMDRs in an IPX environment. Each SMDR advertises the server name where it is loaded using service type 0x23F. But in an IP environment, NDS and Service Location Protocol (SLP) replaces SAP.

Discovery and Name Resolution Using SLP: SMDR can be configured to use Service Location Protocol (SLP) for locating other SMDRs. This enables SMDRs to locate other SMDRs running on servers that belong to different trees. Every SLP enabled SMDR will register itself in the smdr.novell domain when loaded. The SLP enabled SMDRs will query this domain for locating registered SMDRs.

Name Resolution Using HOSTS File: SMDR can also be configured to use the SYS:\ETC\HOSTS file for IP name resolution. It enables SMDR to do name resolution along with SLP and NDS mechanisms. The HOSTS file is automatically installed in the SYS:\ETC directory when you install TCP/IP.

Name resolution using HOSTS file and the discovery and name resolution mechanisms are enabled by default.

Using the SMDR Configuration File

The SMDR.CFG file is a text file located in the SYS:\ETC\SMS\ directory on your server.

You can modify the configuration file from the command prompt by entering:

LOAD SMDR NEW

The SMDR Configuration screen is displayed where you can make the required modifications.

SMDR Configuration Problem

If you try to load SBICON when you have not set the SMDR on a server with the corresponding eDirectory objects and configuration, SBICON will autoloading the SMDR and prompt you for configuration information. (This screen is hidden by the SBICON screen). To rectify this problem, press Alt+Esc to allow the SMDR to complete its setup before loading SBICON.

If NetWare Common Install is used to install SMS (see [“Customizing the NetWare Server as the Backup Server” on page 30](#) for more information), this problem will not occur. If the SMDR is explicitly loaded for the first time, the screen for configuration information will not be hidden.

Memory Requirements

To run SBICON, the host server requires the following:

- A minimum of 3 MB free memory for SBICON
- Memory required to run NetWare 6 (currently 256 MB RAM)
- Memory required to run the clients

If 3 MB of memory is not available, try setting the storage buffers lower than the default and still run SBICON.

Backup Files

Each backup session produces three types of files:

- ◆ **Data files** are copied to the selected storage media.
- ◆ **Log files** are produced by the engine during backup and restore. Log files are placed in a directory on the host server and accessed through the SBICON Main Menu or from a Windows* 95, 98, 2000 or Windows NT* workstation using NWBACK32.
- ◆ **Error files** are produced by the engine while backing up. Error files are placed in a directory on the host server and accessed through the SBICON Main Menu or from a Windows 95, 98, 2000 or Windows NT workstation using NWBACK32.

Both log and error files contain information such as the date, time, and media identification for a session. But the error file also contains a list of any errors that occurred during the backup session, such as files that were not backed up (see [“Log and Error Files” on page 66](#) for more information on these files). The log and error files are labeled with the same description you give the session.

Backup Types

SBICON has three types of backup sessions:

- ◆ **Full backup**—Backs up the entire file system of the selected target regardless of whether the data has changed since the last backup, and clears the **Modify bit** after the backup.

- ◆ **Differential backup**—Available only for the file system; backs up only data that has been changed since the last full or incremental backup. When you perform a differential backup, the modify bit is not cleared after the backup. All files modified since the last full backup are included in the backup (unless they have been deleted). Each differential backup uses more media and is slower than an incremental backup because it backs up more files.

IMPORTANT: Do not interchange differential backups and incremental backups. If you do, the differential backup will not contain all changes since the last full backup. Use full backups interspersed with differential backups or full backups interspersed with incremental backups.

- ◆ **Incremental backup**—Available only for the file system; backs up only data that has been changed since the last full or incremental backup (whichever was last). Incremental backup sessions back up only files that have the modify bit set (that is, files that changed since the last full or incremental backup session when the modify bit was cleared).

Customizing Your Backup

All backup types contain advanced options to allow you to customize your backup. These options allow you to

- ◆ Choose subsets of data to back up

You can choose specific subsets of a data set to exclude from or include in the backup session by selecting major resources, such as volumes, files, directories, or path.

See [“Exclude and Include Options” on page 12.](#)

- ◆ Specify how to scan what you are backing up

See [“Scan Data Sets” on page 14.](#)

Exclude and Include Options

Whenever you perform a custom backup or restore, you can use the exclude and include options to select subsets of what you want to back up.

Whether you use exclude or include usually depends on the size of the data you want to back up, compared to the size of the data you do not want to back up.

Exclude

To back up most of the file system structure or eDirectory tree structure while omitting only a small part, use the exclude option to omit the part you do not want to back up. Everything that you do not specifically exclude is included.

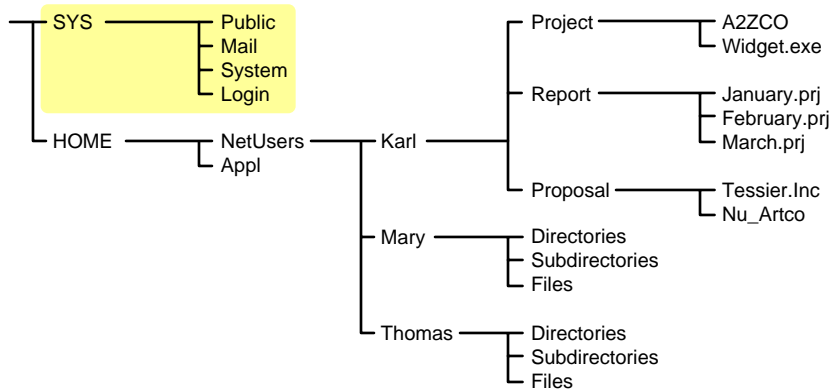
After you exclude part of the structure such as a volume, directory, or container, you cannot include any subdirectories, files, or objects beneath that excluded volume, directory, or container.

Include

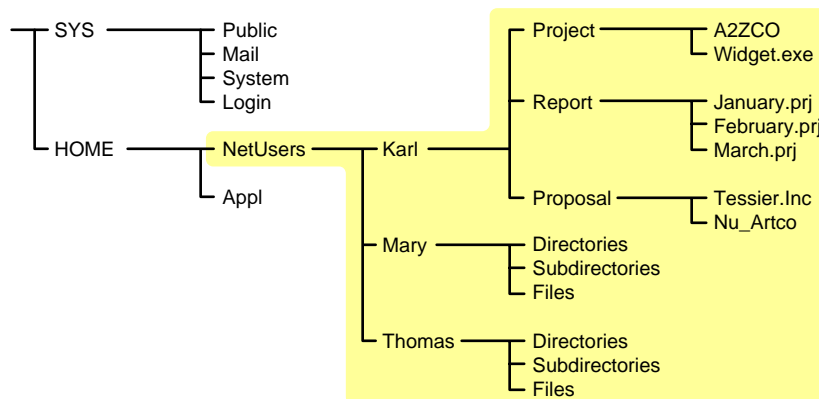
To back up a small part of the file system structure, use the include option to specify the data you want. Everything you do not specifically include is excluded.

When you select only part of the file system structure to include (such as a volume), all directories, subdirectories, and files under that selection are included in the backup by default.

In the figure given below, volume SYS: is selected as an include option. All other areas of the file system structure are excluded from the backup. You can exclude some subdirectories or files beneath your selection if necessary.



The same principle applies when you specify a directory with the include option. The figure below shows that all directories, subdirectories, and files under the NetUsers directory are included in the backup. All other areas of the file system structure are excluded from the backup.



The reverse is true when you select a major TSA resource, a directory, or a file as an exclude option. All other areas of the file system structure are included in the backup.

Combining Include and Exclude Options

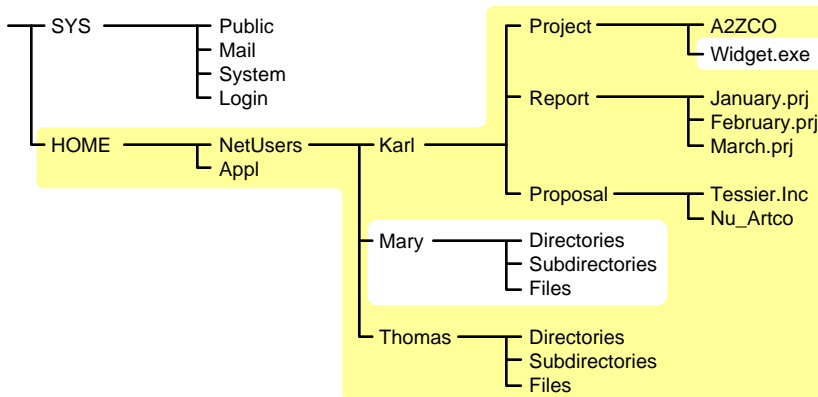
By combining the include and exclude options, you can control what is backed up.

For example, the following command sequence results in volume HOME being included in the backup with the exception of the MARY directory and the WIDGET.EXE file.

Include major TSA resources HOME:

Exclude directories (full path): HOME:NETUSERS/MARY

Exclude path/files HOME:NETUSERS/KARL/PROJECT/WIDGET.EXE



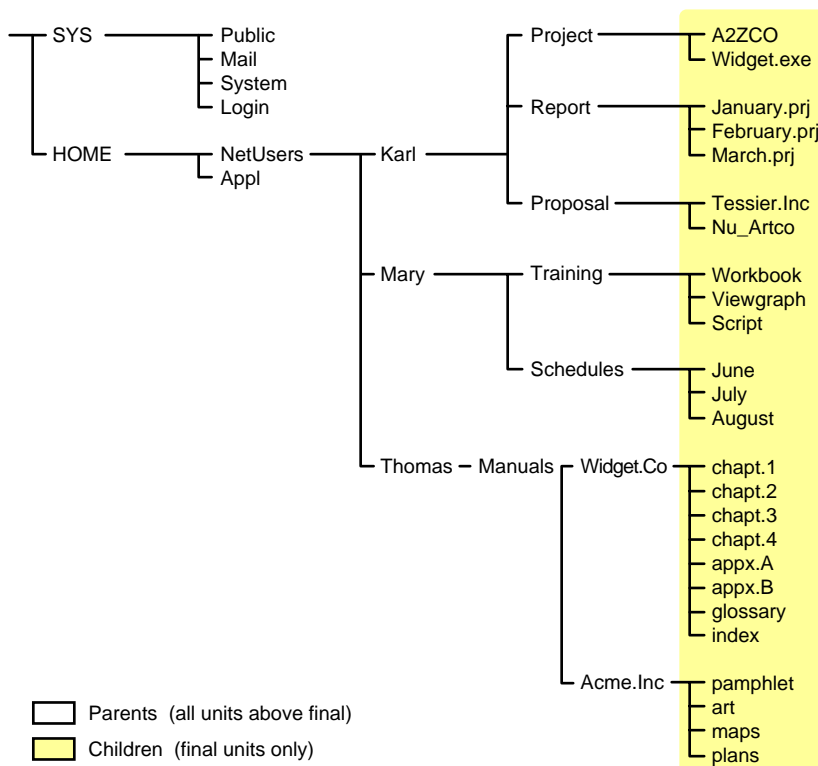
Scan Data Sets

You can specify a different type of **data set** to be scanned.

A data set is a group of data that can be manipulated by SBCON. Each data set in the file system structure can be classified as a parent or a child, and each class includes different types of data items.

Within SBCON, a parent might be a server, eDirectory, a volume, or a directory. A child is a file, which is the lowest level of the directory structure.

The unit below a parent is not necessarily a child; it might be another parent, or the line might end with the parent. The unit above a child must always be a parent.



□ Parents (all units above final)

■ Children (final units only)

Items in a data set for either a parent or child should be items that do not frequently change. You might choose to exclude from the backup session one or more items in the data set of your target.

Overwriting a Parent or Child

SBCON allows you to overwrite all existing parents or children. Children can be overwritten only if the date on the data set on the hard disk is more recent than the date of the data set backup.

Keeping a Backup Logbook

Keep a hard copy log of your backups in case your online log and error files become corrupted. The log should contain the following information:

- ◆ Source backed up (such as the server or workstation name)
- ◆ Full path for **session** log and error files directory
- ◆ Full path for backed-up data
- ◆ Label of the media on which the backup is stored
- ◆ Description of the session
- ◆ Name or initials of the person performing the backup
- ◆ Date and time of the backup session

Planning a Backup Schedule

Before you begin backup procedures, plan a backup schedule based on your needs. Consider such factors as the number of users and frequency of changes to files.

You can perform different types of backups on different schedules:

- ◆ **Daily**—Perform an incremental or differential backup after the close of business. If revisions are heavy and rapid, consider several backup sessions each day.
- ◆ **Weekly**—Perform an incremental or differential backup after the close of business on the last day of the week for three of the four weeks in the month.
- ◆ **Monthly**—Perform a full backup on the last business day of the month (for example, the last Friday).
- ◆ **Major changes**—Perform a full backup before and after you change your configuration, and before and after you upgrade your server to a new version of NetWare.
- ◆ **Application changes**—Perform a custom backup before and after you modify applications.

Preparing to Back Up

Careful planning can help you minimize the impact of data loss. Before you back up, consider the following:

- ◆ What should be backed up
- ◆ Where to back up
- ◆ Backup type
- ◆ Who should back up
- ◆ When to back up
- ◆ How to prepare for a backup
- ◆ How to track backup

Open Files Backup

TSAFS.NLM supports backup of open files on Novell Storage Services (NSS) volumes if the CopyOnWrite feature is enabled.

To enable CopyOnWrite on a single NSS volume, do the following:

- 1** At the server console, enter
`nss /FileCopyOnWrite=volume_name`
 - 2** Dismount and remount the volume.
 - 2a** To dismount the volume, enter
`dismount volume_name`
 - 2b** To remount the volume, enter
`mount volume_name`
- To enable the feature on all NSS volumes, enter
`nss /FileCopyOnWrite=all`

Continue with the next sections to prepare for the backup session.

Determining an Appropriate Backup Type

Each type of backup has a different effect on the backup and restore process. When planning your backup schedule, consider all of the following variables before determining which schedule is right for you.

Media usage and backup speed. This helps increase the speed of the restore.

Restoring after incremental backups. If you have performed full and incremental backups and need to restore data, you must restore the last full backup as well as all subsequent incremental backups.

Restoring after differential backups. If you have performed full and differential backups and need to restore data after an unexpected loss, restore only the last full and differential backup.

Backups and eDirectory

The best way to protect your eDirectory database is to use replicas. Replication, however, is not sufficient protection for a single server network or when all copies of the replicas are destroyed or corrupted. In these instances, if the eDirectory data has been backed up regularly, the eDirectory tree structure can be restored using SMS.

You can back up the entire tree or a selected section of the tree starting with a particular container. You can back up the schema and schema extensions. Trustee assignments are backed up as part of the file system.

You cannot back up partition information. If the eDirectory tree structure becomes corrupted and you restore the eDirectory data, all data is restored to one partition, [Root]. You need to repartition that portion of the tree.

It is important that you keep a written copy of the tree structure and the partitions. You can use the DSMISC.LOG file that is backed up with the file system as part of the server-specific information.

This section discusses the following:

- ◆ Managing “Distributed Database” on page 17, “Server Interdependence” on page 17, “Object ID Numbers” on page 17, “Placeholder (Unknown) Objects” on page 17, and “Schema Backup” on page 17
- ◆ “Backup Software for eDirectory” on page 18
- ◆ “Setting Rights to Back Up Portions of the eDirectory Tree” on page 18
- ◆ “Frequency of Backup” on page 19

Distributed Database

The network of servers that comprise an eDirectory tree structure continually exchange updates and other time-sensitive information. The eDirectory database exists as a set of files that are stored in the SYS: volume and are hidden so they are not accidentally tampered with or deleted.

The eDirectory database files cannot be backed up, as was the case with bindery files in NetWare 3.12 or earlier versions.

Server Interdependence

eDirectory is not server-centric, and neither are its backup and restore processes. Backing up eDirectory, for example, backs up data that is spread out over multiple servers. SMS Directory database backups gather all the necessary eDirectory data.

To handle the necessary links and dependencies between objects, the backup and restore system must be able to navigate the entire eDirectory tree structure.

Object ID Numbers

In NetWare, a random ID number is assigned when an object is created. NetWare uses object ID numbers to keep track of information such as users' trustee rights to directories and files in the file system. These object ID numbers are stored in the directory entry table (DET) of each file and directory and are server-centric.

When NetWare is backed up, SMS-compatible products store the objects' *fully distinguished names* on the backup media, not the objects' ID numbers. If an object with the same distinguished name as on the backup media already exists in the eDirectory tree structure, its object ID is not overwritten during a restore. If an object with the same name does not already exist in the eDirectory tree structure, it is assigned a new object ID when it is restored. This occurs on every server where the object is used.

Schema Backup

The **schema** is backed up automatically with a full eDirectory backup. You can also choose to back up the schema separately using a custom eDirectory backup.

Placeholder (Unknown) Objects

Whenever insufficient information is known about an object, such as when one of its mandatory attributes is missing, eDirectory creates as a placeholder an Unknown object.

During a restore session of the eDirectory database information, Unknown objects are created when restoring an object that has an **access control list** (ACL) or any other attribute that refers to other objects that do not currently exist in the eDirectory tree structure.

This condition is common in a restore, because only one object can be restored at a time. When this condition arises, an Unknown object is created until the real object is restored.

For example, User object User1 has been given property and object rights to User object User2. If User1 and User2 are deleted and only User2 is restored, an object named User1 will be created but it will have a base class of Unknown. This occurs because the access control list of User2 lists User1, which was not restored. The Unknown object is used as a placeholder in the tree. If User1 is later restored, it will replace the Unknown object.

If the restore session does not include the object for which the placeholder was created, the object remains in the eDirectory tree structure as type Unknown. Expect to see Unknown objects after a restore session if all network resources such as servers, volumes, and users are not in place before the restore session starts.

Objects that remain unknown after a restoration is completed are objects for which eDirectory could not resolve the dependencies.

In this case, you can do one of two things:

1. Delete the Unknown objects and re-create the original object.
2. Perform a selective restore to overwrite the Unknown objects.

Backup Software for eDirectory

In order to back up the eDirectory database, the TSANDS.NLM software must be loaded on one server in the eDirectory tree structure—preferably the server containing a replica of the largest partition.

For large or complex networks, you can improve performance by loading the TSANDS.NLM software for a particular partition. This minimizes network traffic during the backup process and improves performance when the backup program performs name resolution across the eDirectory tree structure.

The version of TSANDS.NLM that ships with NetWare allows selective backup and restoration of an eDirectory tree structure.

TIP: Not all third-party backup applications support this selective backup and restoration. Check with the application vendor for details on product features.

In SBCON, you can begin the backup of the eDirectory database from any server in the eDirectory tree structure. The backup process continues from that point downward to the end of that portion of the tree. If the selected container is [Root], the entire eDirectory tree structure is processed.

This allows you to back up the entire eDirectory tree structure or subsets such as a single branch, a single container, or even a single **leaf object**. Also, a scan option allows backup of only those objects for which the backup user has the Supervisor right.

When you back up eDirectory, we recommend that you back up the eDirectory tree structure in one session whenever possible. Although partial eDirectory backups and restores are possible, numerous precautions and additional issues must be noted. See [“Partial eDirectory Restores” on page 55](#) for more information.

Setting Rights to Back Up Portions of the eDirectory Tree

The network administrator can assign backup administrators with limited rights to the eDirectory tree structure.

For example, suppose in your company you have three Organizational Units that need to be backed up (East, West, Mid). You could create three User objects—BackAdmin1, BackAdmin2, and BackAdmin3—and give them rights to the Organizational Unit that they are responsible to back up.

You then create a TSANDS.CFG file that lists the fully distinguished name of the contexts where the backup administrators' rights begin. It would look similar to the following:

```
.OU=East.O=Acme  
.OU=West.O=Acme  
.OU=Mid.O=Acme
```

Backup administrators have rights to back up the eDirectory tree structure beginning only at the context listed, and the rights continue until the tree stops or the rights are filtered out. Backup administrators should use a custom eDirectory backup to back up the portions of the tree for which they have rights.

The network administrator assigns the Supervisor right to the backup administrators for the section of the eDirectory tree structure that they are responsible to back up. The network administrator then needs to create a TSANDS.CFG file that lists the fully distinguished names of the containers where each of the backup administrators' rights begin. The TSANDS.CFG file should be saved in the SYS:SYSTEM\TSA directory of the server.

Frequency of Backup

In general, the eDirectory database should be backed up on a weekly basis. The frequency of this backup depends on how often changes and updates are made to the eDirectory tree structure. For a tree that changes often, you might want to perform an eDirectory backup every time you do a full backup of all servers on the network.

IMPORTANT: Always back up eDirectory prior to major tree modifications.

To get a full backup, the entire eDirectory tree structure needs to be functioning, meaning that all partitions are synchronizing normally. An eDirectory tree cannot be backed up entirely if any replicas of any partition are offline.

Backups and the File System

Back up your volumes so that in the case of hardware failure, natural catastrophe, or accidental change or deletion of files, you can restore the file system to a previous state and not lose the data.

To back up file system data, an appropriate SMS TSA must be loaded on each server for which a file system backup is to be created (see [“Loading the Target Service Agents” on page 31](#)). To back up file system information, make sure your backup application can handle the NetWare file system name spaces, extended attributes, trustee rights, compression, etc.

Once the device drivers for your backup hardware and the SMS TSA software is loaded, you can run the backup program of your choice (see [“Loading SBCON” on page 32](#)).

Trustee Assignments

Trustee assignments are stored as part of the file system as an ID. They are backed up by default when the file system is backed up with the SMS TSA software. If a User object is deleted and then re-created or restored, its object ID changes. This is why the SMS TSA module uses fully distinguished names for objects to back up the trustee rights from the file system. If a User object

is deleted and re-created with a new ID, the user's trustee assignments in the file system can be restored.

As long as an object with the same name on the backup media exists in the eDirectory tree structure when the file system is restored, the TSA can interact with eDirectory to rebuild the **directory entry table** (DET) to reflect new object ID numbers.

For additional information about object ID and trustee issues, see “**Restoring eDirectory**” on **page 48**.

Server-Specific Information

Server-specific information such as the replica information, ID information, name spaces loaded, and system configuration is stored on the volume SYS:. This information is backed up as part of the file system as a single resource. This resource includes the following five files:

- ◆ **SERVDATA.NDS** contains server-specific eDirectory data.
- ◆ **DSMISC.LOG** contains a replica list and replica types on the server at backup.
- ◆ **STARTUP.NCF** contains a disk driver, name spaces, and SET parameters.
- ◆ **AUTOEXEC.NCF** contains load modules and the NetWare operating system configuration.
- ◆ **VOLSINFO.TXT** contains volumes on the server, name spaces loaded, compression, and migration information.

You can also choose to back up this information individually. The information is not restored unless you specifically choose to restore it. It does not need to be restored unless you have lost the SYS: volume. In that case, you must replace the hardware and restore this information. For more information, see “**Restoring the Entire eDirectory Tree Structure**” on **page 50**.

Backups and Clusters

Novell Cluster Services™ allows you to configure up to 32 NetWare servers into high-availability cluster, where resources can be dynamically switched or moved to any server in the cluster. Consolidation of applications and operations on a cluster has benefits such as lower costs, scalability, and increased availability. See the [Novell Cluster Services documentation \(http://www.novell.com/documentation\)](http://www.novell.com/documentation) for more information.

For a cluster to work as a high-availability system, the file system, the applications, and services that run on the cluster should be cluster-enabled. SBCON supports backup and restore of cluster-enabled pools. In addition, the backup session can be automatically recovered in case of a failover or failback condition.

NOTE: Backup and restore of cluster-enabled pools is not supported in NetWare versions earlier than NetWare 6.

SBCON supports automatic recovery of backup sessions if failover or failback occurs. The backup engine reconnects to the Target Service Agent and resumes the backup from where it had terminated. Various cluster options like Enable Auto-Recovery and Retry Interval are provided by SBCON. You can use the default values or reset the values while submitting the backup job. The engine begins the reconnection attempts after waiting for a configurable time, and then retries at regular intervals until the connection is re-established or the number of retries has expired.

After the connection is re-established, the internal structures of the TSA are built and the recovered session is continued from where it had terminated. User intervention is not required during the recovery period. You can view the status in the Session Report screen.

Consider the following before preparing for backup and restore of cluster-enabled pools:

- ❑ Use TSAFSTSAFS.NLM to backup or restore cluster-enabled pools. (See “[Target Service Agents \(TSAs\)](#)” on page 21 for details.)
- ❑ If SBCON is loaded on a cluster node and a backup session for a cluster-enabled pool is submitted, the session cannot be recovered if a failover or failback condition occurs on the host server.
- ❑ SBCON can backup and restore only one cluster-enabled pool per session. This is because the engine is connected to the cluster pool instead of the server.
- ❑ If a cluster server is chosen as the target, only noncluster volumes will be backed up. You have to back up each pool individually.
- ❑ You might abort an ongoing backup job with the intention of resubmitting the same job later. In such situations, the job will not restart from where it was terminated; it will restart from the beginning.

For more information, see “[Backing Up Cluster-enabled Pools from the Server](#)” on page 40.

Target Service Agents (TSAs)

In SMS, a target is any machine on the network that requires backup. Examples of targets include SQL database engines, eDirectory databases, workstations, and NetWare servers.

Through specific TSAs, SBCON allows you to back up the information that exists on the following targets. These TSAs are listed in general as follows:

Table 1 Target Services and Their Corresponding Target Service Agents

| Target Service | Corresponding Target Service Agent |
|---|------------------------------------|
| NetWare 6 | TSAFS |
| eDirectory | TSANDS |
| Windows 95 and 98 workstation | W95TSA |
| Windows 2000 and Windows NT workstation | Windows NT TSA |
| GroupWise® data | GW TSA |

A Target Service Agent (TSA) is a software module that understands how to scan, read, and write the target data. The primary functions of an SMS TSA are to prepare the target data for backup or restoration, and to communicate with the storage management engine (SME). For example, an SMS TSA for a NetWare server understands name spaces, file and directory attributes, security privileges, etc., for the data on that server.

The TSA packages data from the target and presents it to the SME in a generic format. This allows one SME to interact with many types of TSAs.

NetWare 6 provides TSAFS.NLM, with the following features:

- ◆ Supports backup and restore of traditional NetWare file system and NSS file system

- ◆ Supports backup and restore of cluster-enabled pools (see “Backups and Clusters” on page 20)
- ◆ Recovers the backup session in failover and failback conditions

Restore Services

Restore services provides information on how SMS restores data from eDirectory, volume SYS: and other volumes, single servers, multiple servers, and the eDirectory tree structure.

Use SBCON to retrieve and reinstate data you have backed up to storage media. You restore data if it has been lost or corrupted since a backup was made.

A restore session restores data from a backup. The restore session produces the requested data, which is retrieved from the storage media and restored to the location you specify. If an error occurs during the restore session, an error message is appended to the error file on the host server.

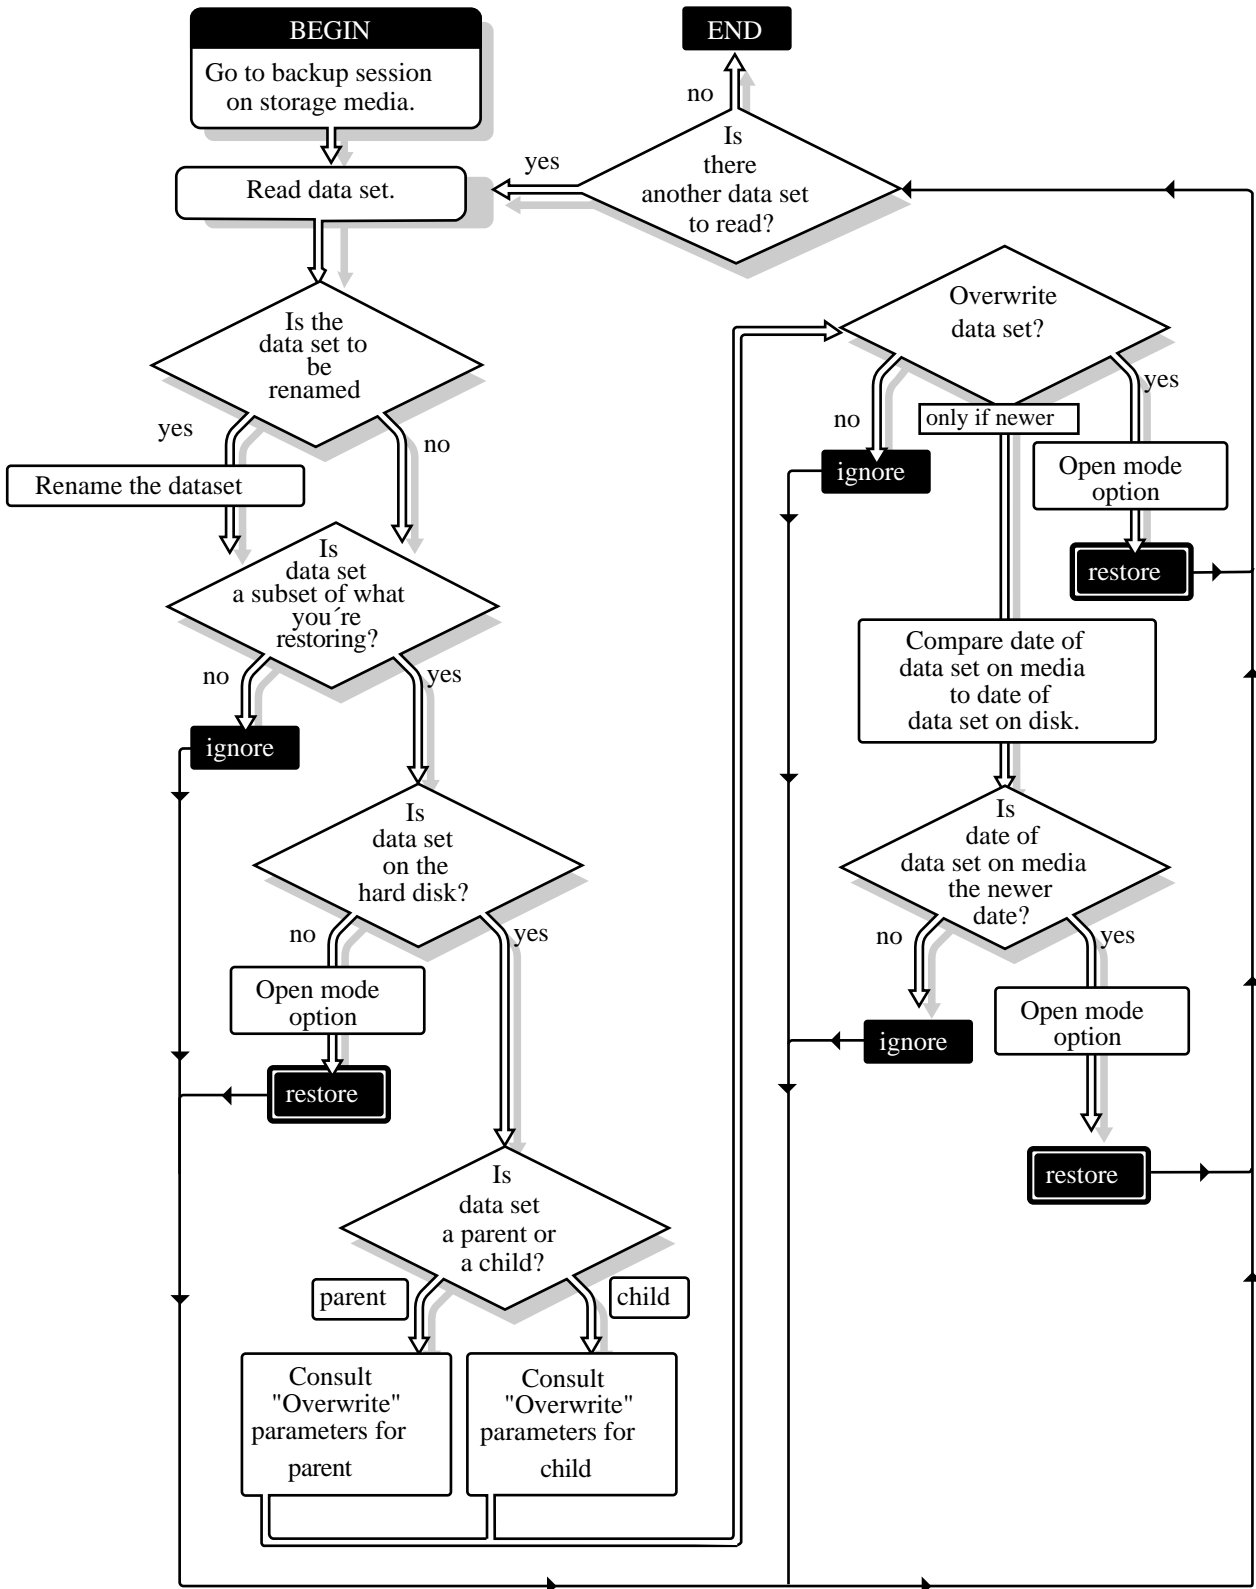
The following topics are discussed in this section:

- ◆ “How SMS Restores Data” on page 22
- ◆ “Restore Options” on page 24
- ◆ “Restoring eDirectory and the File System” on page 24
- ◆ “Restoring Cluster-enabled Pools” on page 26
- ◆ “Session Files” on page 26

How SMS Restores Data

During a restore session, SBCON reads the backup storage media, and the Target Service Agent (TSA) compares the media data set to the existing hard disk data set. The Target Service Agent evaluates each data set according to the following criteria:

- ◆ Is this **data set** a subset of what is being restored?
- ◆ Is this data set found on the hard disk?
- ◆ Which parts of the data set are subject to restoring?
- ◆ Is this data set a parent or a child, and is the Overwrite parameter set to Yes or No?
- ◆ If the parameters for a child are set to Overwrite Only if Newer, does the backup copy have a more recent date than the existing copy?



Restore Options

For a custom restore session, you can specify exactly which data to restore. Several options work together to allow you maximum flexibility in your restore session. These options allow you to do the following:

- ◆ Choose subsets of data to restore
- ◆ Open mode options
- ◆ Overwrite an existing parent (such as a container) or child (such as an object)

Subsets of Data to Restore

You can choose specific subsets of a backup session to include in or exclude from the restore session by selecting major resources (such as volumes, server-specific info, or containers) or minor resources (such as directories, paths, files, or objects).

For more information about including and excluding, see [“Customizing Your Backup” on page 12](#).

Open Mode Options

Open mode options allows you to customize data for restore. File system data can either be included or excluded for the session. The speed of the restore depends on the options you set.

Overwriting Existing Parents or Children

Be careful when you perform a selective restore and choose whether to overwrite existing parents or children, especially eDirectory objects. Objects such as groups and users have references to other objects in the eDirectory tree structure that will be affected by a selective restore.

For example, suppose a part of the eDirectory tree structure gets corrupted and several users are deleted from the tree. There is a group that contains those users, but once the users are gone, the group purges the membership list to remove those users; the group, however, continues to exist in the eDirectory tree structure.

If you perform a selective restore and choose not to overwrite existing objects, the group membership list remains empty even if you restore the users. You need to either add the users manually to the group membership list or restore the original group.

Restoring eDirectory and the File System

The only way to ensure that your eDirectory database can be fully restored is through partition replication, with replicas of the entire database on multiple servers.

On a single-server network, you need to rely more heavily on backing up the data because you do not have replicas to restore information.

If part of the eDirectory tree structure, including partitions and replicas, exist when the eDirectory database information is restored, those partitions and replicas will be restored also, and you will not need to repartition the tree.

In case of corrupted data, follow these general steps:

1. Delete the corrupted eDirectory data.

2. Allow time for the deletion to propagate throughout the network. The allotted time depends on the size of the data to be backed up, the size of your network, the number of servers you have, and the number of containers and users you have.
3. Restore the eDirectory data.

A replica containing the object does not have to be on the server. The eDirectory database creates **external reference** when necessary.

NOTE: An external reference is a pointer to an eDirectory object not found locally on the server; it is used to authenticate and reference objects that are not local to the server.

There are specific recovery procedures for the following scenarios:

- ◆ “Loss of a Volume Other Than SYS:” on page 25
- ◆ “Loss of the SYS: Volume or an Entire Server” on page 25
- ◆ “Loss of the Entire eDirectory Tree” on page 26

Loss of a Volume Other Than SYS:

Loss of a volume other than SYS: does not affect eDirectory. The only requirement is to restore the file system data and trustee rights.

Loss of the SYS: Volume or an Entire Server

A hard disk failure involving the SYS: volume affects the entire server and halts all NetWare operating system activities. Because the eDirectory files are stored on volume SYS:, losing SYS: is equivalent to removing NetWare and eDirectory from the file server. You must reinstall NetWare and eDirectory before you restore your data.

The procedures for this scenario are divided into two cases:

- ◆ Loss of the only server in a single-server network.
- ◆ Loss of a single server in a multiple-server network.

Single-Server Network

In a single-server network, server failure brings all network operations to a halt. The same situation exists if the failure affects only the hard disks containing the SYS: volume.

Since there are no replicas in a single-server network, you cannot recover any eDirectory information from a replica.

Use the following general steps to restore the server:

1. Repair or replace the failed hardware.
2. Reinstall NetWare.
3. Restore eDirectory from an SMS backup. (See “Restoring Data” on page 43.)
4. Restore the file system. (See “Restoring Data” on page 43.)

Multiple-Server Network

In a multiple-server environment, it is possible for one server to go down and for the rest of the servers in its replica list to remain intact. The same situation exists if the hard disks containing volume SYS: on one server gets damaged, causing the failure of the entire server.

Use the following general steps to restore the server:

1. Repair or replace the failed hardware.
2. Restore SERVDATA.NDS (the server-specific eDirectory information) for the failed server to another server on the network. (See [“Restoring Data” on page 43.](#))
3. Reinstall NetWare, including restoring the SERVDATA.NDS file (located in SYS:) to the original server.
4. Restore eDirectory. (See [“Restoring Data” on page 43.](#))
5. Restore the file system. (See [“Restoring Data” on page 43](#) and [“Custom File System Restore” on page 52.](#))
6. Restore any replicas that were removed from the server.

Loss of the Entire eDirectory Tree

If all servers on a network are destroyed because of a disaster, you must perform a complete restore of NetWare, eDirectory, and file system data.

IMPORTANT: We recommend that you document your eDirectory tree structure; the location of Server objects, partitions, and replicas; and record bindery context settings and other relevant information.

Use the following general steps to restore the eDirectory tree structure:

1. Repair or replace the failed hardware.
2. Reinstall NetWare on the first server.
3. Install NetWare on remaining servers to create a skeleton of the tree.
4. Restore eDirectory. (See [“Restoring eDirectory” on page 48.](#))
5. Restore the file system to all servers. (See [“Custom File System Restore” on page 52.](#))
6. Re-establish partition boundaries and distribute replicas.

Restoring Cluster-enabled Pools

Restore of cluster-enabled pools is similar to a normal restore session (see [“Restoring Data” on page 43.](#))

NOTE: Auto recovery of the restore session on failover or failback is currently not supported.

Session Files

Whenever you back up or restore files/directories, a log and corresponding error files are created for that particular session on the server on which you are executing the job. The .LOG file is the session file for that particular job.

The need to create a session file arises when you are restoring the data on a server other than the one you backed up from. The server should possess the session ID of all the jobs that are present on the tape so that the SME can restore it. These session IDs are created from the .LOG files. See [“Creating Session Files” on page 69.](#)

Log and Error Files

The files of the backup session are stored in a default directory (such as SYS:SYSTEM\TSA\LOG). If you prefer, you can create your own directory for the log and error files as long as it resides on the host server.

The files of the restore session are stored in the directory (SYS:SYSTEM\TSA\RESTORE). You cannot modify the location for the files.

The files can be accessed through the SBCON Main Menu or from a Windows 95, 98, 2000 or Windows NT workstation using NWBACK32.

The error file is labeled with the same description that you give the restore session (such as Friday's Full Backup) and is accessed through the Log/Error File Administration option of the SBCON Main menu.

TIP: You might want to create individual log directories for the different types of backup or restore session targets or different organizational units. For example, you could create one directory for workstation backups and another for server backups.

SBCON keeps a list of all the log and error files. These lists show

- ◆ The description you enter for the session.
- ◆ The date and time the session was begun or, in the case of a delayed backup session, the time the session was scheduled.
- ◆ The name of the target the data was backed up from.

Log File

The log file is created on the host server the first time a particular set of data is backed up. This file contains the following:

- ◆ The session date and time and the description you entered.
- ◆ The target from which the data was backed up for a backup session.
- ◆ The target that was backed up and the location on the server where the data was restored to during a restore session.
- ◆ Media set identification information.
- ◆ The area of the file system structure that was backed up or restored (volume name, directory name, etc.).
- ◆ The names of files that were backed up or restored.
- ◆ The numerical location of the data on the storage media.

Error File

The error file is created on the host server the first time a particular set of data is backed up. It contains a list of any errors that occurred during a backup or restore session.

This file contains the following:

- ◆ The session date and time and the description you entered.
- ◆ The target from which the data was backed up for a backup session.
- ◆ The target that was backed up and the location on the server where the data was restored to during a restore session.

- ◆ Media set identification information.
- ◆ The area of the file system structure that was backed up or restored.
- ◆ The total number of parents and children that were backed up or restored.
- ◆ The names of files that were not backed up or restored, along with any error messages or information.
- ◆ Skipped data sets (any file that is open when a session begins, is not backed up or restored, and is listed as a skipped data set).

2

Setting Up

In addition to being able to back up data from a server with SBCON, you can use a Windows* 95, 98, 2000 or Windows NT* workstation to perform your backups using the graphical interface.

This chapter discusses the following:

- ♦ How to customize the NetWare[®] server as the backup server. See “[Customizing the NetWare Server as the Backup Server](#)” on page 30.
- ♦ How to load the backup software. See “[Loading Controller and Storage Device Drivers on the Server](#)” on page 31, “[Loading the Target Service Agents](#)” on page 31, and “[Loading SBCON](#)” on page 32.
- ♦ How to unload the backup software. See “[Unloading SBCON](#)” on page 33.
- ♦ How to back up your data. See “[Backing Up from the Server](#)” on page 35, “[Customizing Backup](#)” on page 37, or “[Backing Up from a Workstation](#)” on page 39.
- ♦ How to restore your data. See “[Customizing Restore](#)” on page 46, or “[Restoring eDirectory](#)” on page 48, and “[Restoring a Volume Other Than SYS:.](#)” on page 48.

SBCON Guidelines

Before you set up and begin backups, follow these guidelines to make backups easier and to avoid problems later.

- ♦ **Log in.** You will have limited success backing up and restoring if you log in without the Supervisor right.
For security reasons, many SBCON options are limited to the network administrator.
- ♦ **Reserve disk space for temporary files.** Make sure you have disk space available (1 to 2 MB) on the target server's volume SYS: to accommodate log and error files.
SBCON creates temporary files on the target server during backup. If you have linked UNIX* files or files with extended attributes, the temporary files might be larger than 1 MB.
- ♦ **Do not mount or dismount volumes during a backup or restore session.** The data might be corrupted or an **association** might occur at the host server.
- ♦ **Use the correct name space and name space formats.** If you do not use the correct name space and name space formats when entering paths and filenames, files can't be restored.
Name spaces: DOS, Macintosh, NFS, and Long
Name space formats:
 - ♦ For Macintosh, use *Volume::directory:directory:filename*
 - ♦ For all others, use *Volume:/directory/directory/filename*

- ◆ **Exit SBCON before unloading drivers.** If you unload a manually loaded driver (such as `aha1740.dsk` or `asptran.dsk`) before exiting SBCON, you might cause the host server to abend.
- ◆ **Use original case for non-DOS names.** Non-DOS pathnames and filenames are case-sensitive. NetWare recognizes DOS pathnames and filenames in uppercase only. If you're not sure of the original case, refer to your log file.

Before Loading the Backup Software

Before you load the backup software on the server or on a Windows machine, you need to configure a server to execute your backups.

Determine the following:

- ◆ The tree to which the server belongs which must be made the current tree
- ◆ The full eDirectory™ context of the server

Before Running the Backup Software

Before backing up software, in general you must:

- ◆ Load the controller and storage device drivers on the server. See [“Loading Controller and Storage Device Drivers on the Server” on page 31](#).
- ◆ Load the TSAs on a server, workstations, or both. See [“Loading the Target Service Agents” on page 31](#).
- ◆ Load the backup engine on the server or the workstation. See [“Loading SBCON” on page 32](#) or [“Loading NWBACK32” on page 33](#).

Customizing the NetWare Server as the Backup Server

SMS allows you to configure a NetWare® server as a backup server while installing NetWare 6. It also allows you to create NDS® Backup Queue and SMDR Group objects.

To customize the settings while running the Install program, do the following:

- 1** In the Final List of Products to be Installed dialog, click Customize.
- 2** Select Storage Management Services from the list of NetWare Services > click Configure.
The SMS Setup dialog box with default values for NDS Backup Queue and SMDR Group objects displays.
- 3** Enter the name of the eDirectory tree where the SMS objects will be created.
- 4** Enter the name of the backup server and the context.
- 5** Accept the default for the SMDR Group context where the SMDR Group object is to be created, or enter a new context.
The SMDR Group object is where all NetWare and workstation SMDRs are registered for name resolution purposes.
- 6** Accept the default name of the Backup Queue where all backup jobs will be submitted, or enter a new name and context.

Loading Controller and Storage Device Drivers on the Server

Controller and device drivers control the mechanical operation, such as read, write, forward, back, and stop of various storage devices and media.

Ensure that you have installed the storage device controller and storage device according to the hardware manufacturer's instructions.

When NetWare 6 is installed, the device drivers are placed in the `startup.ncf` file automatically.

Whenever you want to add new drivers, you can either place them in `STARTUP.NCF` or use the following steps:

- 1 At the server console prompt, enter:

```
LOAD PATH CONTROLLER_DEVICE_DRIVER_NAME(S)
LOAD PATH STORAGE_DEVICE_DRIVER_NAME(S)
```

For example, to load the drivers from drive C:, enter:

```
LOAD SCSI154X.HAM PORT=24
LOAD SCSIHD.CAM
LOAD SCSI2TP.CDM
```

- 2 In `NWCONFIG`, select `Driver Options > Configure Disk and Storage Device Drivers > Discover and Load Additional Drivers [or] Select an Additional Driver`.
- 3 (Conditional) If you loaded `HAM` drivers, or if you added an external device, enter the following at the server console prompt:

```
LIST DEVICES
```

- 4 To register the storage device with the system, enter the following at the server console prompt:

```
SCAN FOR NEW DEVICES
```

NOTE: If you load the drivers from the `STARTUP.NCF` file, you do not need to use the `SCAN FOR NEW DEVICES` command.

Loading the Target Service Agents

Target Service Agents are loaded on servers and workstations. If you want to back up a server, load the TSA on the server. Ensure that the controllers and device drivers are loaded.

Table 2 Server Console Commands for Backup and Restore of Various Target Services

| To Back Up or Restore | At This Console | Enter This Command |
|-----------------------|--|--------------------|
| NetWare 6 | Host server | LOAD TSAFS |
| eDirectory | NetWare 4 or later server (preferably a server with a copy of the largest partition) | LOAD TSANDS |

| To Back Up or Restore | At This Console | Enter This Command |
|-------------------------------|--------------------|--|
| DOS FAT partitions on NetWare | Host server | LOAD DOSFAT.NSS and TSAFS After you load the DOSFAT.NSS module, any DOS FAT partitions are dynamically made available as logical volumes. TSAFS considers the logical volume as the resource for backup and restore. |
| Windows 95 and 98 | Target workstation | W95TSA.EXE (installed with the Novell® Client™). See the online help for more information. |
| Windows 2000 and NT | Target workstation | TSAPREFS.EXE, TSAMAIN.EXE (installed with the Novell Client). See the online help for more information. |

Follow these general steps when backing up or loading the software:

- ◆ **Backing up the eDirectory database.** Load TSANDS.NLM once on the server with a replica of the largest partition.
- ◆ **Backing up the file system.** Load TSAFS.NLM for its server and on every server to be backed up.
- ◆ **Backing up cluster-enabled pools.** Load TSAFS.NLM on the target server and the backup servers.
- ◆ **Backing up workstations.** Load the appropriate TSA on the workstation.
- ◆ **Loading the files on startup.** Place the required commands in the server's STARTUP.NCF and in the workstation's NET.CFG file.

Loading SBCON

You can load SBCON from the server command line.

If you use NWBACK32 on a client, you do not need to load SBCON.

- 1** At the server console, enter

SMSSTART

TSAFS.NLM, TSAPROXY.NLM, and SMDR.NLM modules are loaded with default configuration values.

SBSC.NLM, QMAN.NLM, SMSDI.NLM and TSANDS.NLM must be manually loaded on need basis; SBSC.NLM, QMAN.NLM, SMSDI.NLM when storage devices are used, and TSANDS.NLM for backup and restore of eDirectory.

- 2** Load SBCON by entering

LOAD SBCON

Prerequisites

- ❑ Controller and storage device drivers are loaded on the server.
See “[Loading Controller and Storage Device Drivers on the Server](#)” on page 31.
- ❑ SMDR.NLM, TSAFS.NLM, SMSDI.NLM, QMAN.NLM, SME.NLM, SBSC.NLM, and SBCON.NLM are installed on the server.
They should have been installed automatically during NetWare 6 installation.
- ❑ Each NLM™ listed above is in the search path.
- ❑ TSA is loaded on each machine to be backed up.
See “[Loading the Target Service Agents](#)” on page 31.

Unloading SBCON

To ensure maximum memory utilization and throughput, you should unload backup software and other Storage Management Services™ (SMS™) modules when you are finished performing your backup. That way, more memory will become available on your host or target.

- 1** Exit SBCON from the Main Menu.

Press Esc until you reach the SBCON Main Menu; then press Esc again and confirm the prompt.

- 2** Unload all the modules by entering the following command at the server console:

SMSSTOP

Unloads all the previously loaded SMS modules.

WARNING: The system displays a warning message on the console if you attempt to unload a module being used by another NLM. If you continue, your server might abend.

Loading NWBACK32

- 1** Log in to the desired eDirectory™ tree structure.

If this is the first time you have used NWBACK32, go to [Step 2 on page 33](#). Otherwise, go to [Step 3 on page 33](#).

- 2** (Conditional) If this is the first time you have used NWBACK32, do the following:

- 2a** Log in to the server running the backup engine.

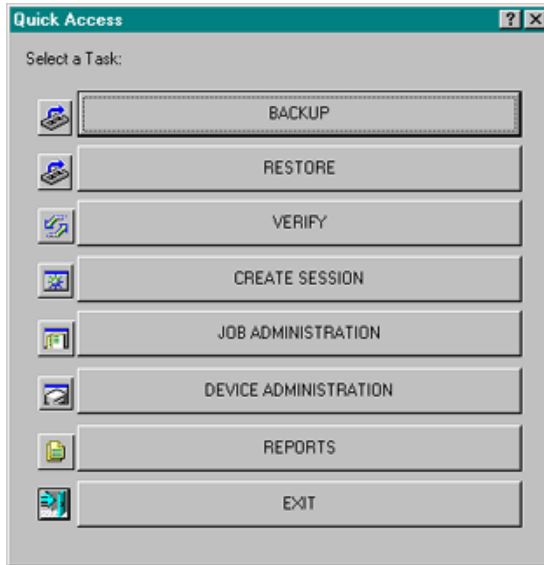
- 2b** Enter your user name and password.

- 2c** Configure eDirectory information.

- ◆ **Tree Name.** Enter the name of the eDirectory tree you will be using to backup and restore data.
- ◆ **SMDR Context.** Enter the context you created during the server configuration.
- ◆ **SMDR Group Context.** Enter the context you created during the server configuration.
- ◆ Check the SAP (IPX) box if you want to use IPX. Uncheck the box for IP.

- 3** Run NWBACK32.EXE from SYS:PUBLIC.

The following screen appears.



You are now ready to back up data. Go to [“Backing Up Data” on page 34](#) to continue.

Exiting NWBACK32

- 1 Exit NWBACK32 from the Quick Access window.
- 2 Unload all the modules loaded on the server by entering the following command at the server console:

```
SMSSTOP
```

Unloads all the previously loaded SMS modules.

Backing Up Data

You can use SBCON to create a tape backup of the eDirectory database or of the file system on your servers or workstations. The prerequisites listed here use SBCON as the Storage Management Engine (SME).

Prerequisites

- Understand the process for loading SBCON.
See [“Loading SBCON” on page 32](#)
- Ensure that you know the workstation password if you are backing up a workstation.
- Set the backup server and its tree as the primary resource using the Novell Client. This lets you view the tape devices in NWBACK32.
- If you are backing up the eDirectory database, ensure that eDirectory synchronization and communication are functioning properly.

Before starting the backup, make sure eDirectory is fully functional. This means that all partitions are synchronizing correctly.

If your backup host and targets communicate across a WAN, check the status of the WAN links to verify that they are operating properly.

- If you are backing up the file system, make sure you have Read and File Scan **rights** to the directories/files you are backing up.
- Ensure that you know what type of backup you want to perform: full, differential, or incremental.

For more information, see **“Backup Types” on page 11.**


- If you use the Advanced Options, ensure that you know the file system structure of the target you are backing up. You will be prompted for the paths to the volumes and directories that you want to include in or exclude from your backup.
- Ensure that you have loaded the drivers for your specific device and controller board.
- Ensure that the required files for your target are loaded.
- Ensure that you have loaded the SBCON files for your specific target.
See **“Loading the Target Service Agents” on page 31.**
- Ensure that media is inserted into your storage device.

You can back up your data two ways: from a server using SBCON or from a Windows workstation using NWBACK32.

Backing Up from the Server

When you back up your data from one server to a target server or workstation, you should use the command line.

- 1** Load the SBCON modules.
See **“Loading SBCON” on page 32.**
- 2** Click Main Menu > Job Administration > Backup.



- 3** Click Target Service to select the NetWare server running your backup or restore Target Service Agent.

To be able to select the server, you need to load the correct Target Service Agent (TSA). See **“Loading the Target Service Agents” on page 31.**

If you do not see the target you want on the list, check for these possible causes:

- ♦ The correct Target Service Agent is not loaded on either the host or the target. See **“Server Console Commands for Backup and Restore of Various Target Services” on page 31** for LOAD commands.
- ♦ Network traffic is heavy. Press Esc to return to the Main Menu, and then press Enter again. Your target should now be listed.

- 4** Select a target from the Target Services list.

The following table describes the targets that can be selected for backup.

| If you are Backing Up | Then Select |
|-----------------------------|---|
| Cluster-enabled pools | The cluster-enabled pool you want to back up |
| The file system | The server whose file system you want to back up |
| The eDirectory database | The server with TSANDS loaded |
| A workstation | The workstation's host server, and then the workstation |
| Server-specific information | The server whose server-specific information you want to back up, and then choose to back up the file system. |

- 5** When prompted, enter the username (and context, if required) for the target.

If SBICON rejects the username you entered, you probably need to include the context where the User object is located. For example, instead of entering **ADMIN** as the username, enter **.CN=ADMIN.O=company_name** or **.ADMIN.company_name**.

You must include the context in the username at this point if the username you use exists in an eDirectory container that is different from the bindery context set on the server running the engine.

- 6** (Conditional) If requested, enter the password for the target.

SBON will take a few moments to attach to that target.

If you select the default values on this form, a full backup of the source is performed. Otherwise, if you want to perform an incremental or differential backup, select Advanced Options.

- 7** Click What to Backup and press Enter.

- 8** Press Insert, select the volumes and files you want to back up from the screen, and press Esc.

The selected volumes and files appear in the List Resources screen.

- 9** Enter a descriptive name for the backup session in the Description field.

This is a descriptive name of the job, which will help you identify the specific backup session; if a restore is necessary, you can easily identify the session you need. There is no special requirement for what to enter as the session description.

TIP: You might want to include the full path of the data (for example, SYS:HOME\REPORTS\JULY.02) which you will need to know if a restore session is necessary.

- 10** Click Device/Media Name and select a device.

A list of available objects is displayed in the Select a Device screen.

If you accept the default value, the engine selects the device and media which you have designated as the default.

NOTE: If there is only one media type available, it will be treated as the default.

If your device has more than one storage device attached, select an available device from the list. If the device contains multiple media, select the media for the backup.

If the storage media does not have a label, SBICON displays a message saying that the media cannot be identified. Set the label before selecting it for a backup job.

11 Select Advanced Options.

For more details, see [“Customizing Backup” on page 37](#).

12 Select Append Session.

This option appears only if your device supports appending to previous sessions on the media.

Accept the default YES if you want to preserve all sessions on the media. The backup session is appended to the media at the end of the previous session. Each appended session has separate backup and error logs.

If you want to overwrite the data on the media, enter NO. The media rewinds and existing data is overwritten by the next backup.

TIP: Labeled media is not considered empty media. No sessions are necessary on the media for you to set the Append option to Yes.

Delete the log and error files for any session you overwrite to avoid trying to restore a session that no longer exists.

13 Press Enter to save the settings and continue with the backup.

WARNING: If you are appending a backup session to a media set (two or more tapes), use the tape with maximum space.

14 To return to the Main Menu, press Esc until you back out to the Main Menu.


15 To exit SBICON, press Esc again and answer the confirmation prompt.

Customizing Backup

You can customize your backup using the Advanced Backup Options form when you use SBICON at the server to launch your backup program. This form allows you to specify exactly what you want to back up data. It could be the server, the eDirectory database, workstation, directory, or file. You can even specify subsets of groups. The Advanced Backup Options form activates only after you have selected at least one of the resources (directory/file). Select Advanced Options from the Backup Options form to perform advanced backup sessions. See [“Backing Up from the Server” on page 35](#).

TIP: SBICON displays prompts and messages at the bottom of the screen. For example, you can press F1 for help at any time.

If you want to perform a series of backups and restores and need to change your target, see [“Changing Your Target” on page 43](#).



A custom backup allows you to specify exactly what you want to back up: server, volume, the eDirectory database, workstation, directory, or file.

To customize a backup, do the following:

1 Select Backup Options > Advanced Options.

2 Select Backup Type.

By default, Storage Management Engine (SME) performs a full backup. See [“Backup Types” on page 11](#). Differential and incremental backups are not interchangeable. See [“Planning a Backup Schedule” on page 15](#).

IMPORTANT: If you choose to include migrated data, ensure that your backup media capacity matches the jukebox media capacity.

3 Select Subsets of What To Backup.

The default setting means that everything in the Subsets of What You Want to Backup screen is backed up.

Include/Exclude TSA Resources—Includes or excludes resources such as servers or volumes.

Include/Exclude Directories—Includes or excludes directories.

Include/Exclude Files—Includes or excludes files.

Include/Exclude path/files—Includes or excludes pathnames and files.

4 Select Scan Options.

The How to Scan What you are Backing Up screen appears.

5 Do one of the following:

- ◆ If you do not want to accept the default (nothing will be excluded), type Y (Yes) and press Enter.
- ◆ Choose to exclude subdirectories, hidden files, system files, etc.

6 Press Esc to Finish.

For information about how SBCON scans data, see [“Scan Data Sets” on page 14](#).

7 Select Execution Time.

The Edit Date and Time screen appears. You can choose the default date and time shown for the backup to begin, or you can type a new date and time. You can change the month/day/year and hour/minute/second fields. The engine should initiate the rescheduled job.

To schedule the job to run another time, press Enter and go to [Step 8 on page 38](#). Otherwise, go to [Step 7 on page 38](#).

8 (Optional) Select Scheduling.

Reschedule—Select Yes to reschedule.

Rerun Interval—Specify the duration in terms of elapsed time and days, months or years to rerun an already scheduled job.

The Edit Date and Time screen appears. You can specify the number of days, months, or years after which an already scheduled job will run.

The day, month, or year you enter should not be the actual date when you want the job to run; rather, it is the duration after which the job gets executed.

Example

If you have scheduled a job to run on 03/20/1999 and want to rerun the job after two months and at the same time, then the Rerun Interval field should be the following:

Year: 0, Month: 2, Day: 0

Hour: 0, Minute: 0, Second: 0

Rerun Count—Specify the number of times you want to rerun a job after the current execution.

Keep Finished jobs—Select Yes to keep the job in the queue after execution so that it can be rescheduled when you need it.

NOTE: If you are performing a series of backups or restores and need to change your target, follow the procedures in [“Changing Your Target” on page 43](#).

Backing Up from a Workstation

When you load the backup software from a workstation to the target server or workstation, you are using the graphical interface.

To back up from a Windows workstation, do the following:

1 Run NWBACK32.

See [“Loading NWBACK32” on page 33](#).

2 At the Quick Access window, click Backup.

3 Select what you want to back up.

3a Double-click What to Backup.

3b Click eDirectory, NetWare Servers, or Workstations.

3c Double-click a server to back up.

3d Enter the username and password to authenticate to the server.

A list of resources to back up appears, such as volumes and files.

3e Click the resources to back up.

3f Click OK to finish.

4 Select where you want your backup data to go.

4a Double-click Where to Backup.

4b Change to the correct eDirectory tree structure and context by clicking the Change to Context button on the toolbar.

4c Double-click Queues.

4d Select a Queue object from the list.

4e Right-click the server to open the drop down box.

4f Click Submit the Job.

This step allows you to submit the job to a queue on the chosen server or on the chosen media type.

Queue. Double-click to expand the queue.

Server. Double-click to expand the queue, double-click to select a server, and then right-click to submit the job.

Media. Double-click to expand the queue, double-click to select a server, double-click to view a list of devices, double-click to view a media list, and then right-click to submit the job.

Backing Up Cluster-enabled Pools from the Server

- 1 Load the SBICON modules.

See [“Loading SBICON” on page 32](#).

- 2 Click Main Menu > Job Administration > Backup.

- 3 Click Target Service to select a cluster-enabled pool name.

To be able to select the target, you need to load the correct Target Service Agent (TSA). See [“Loading the Target Service Agents” on page 31](#).

If you do not see the target you want on the list, check for these possible causes:

- ♦ The correct Target Service Agent is not loaded on either the host or the target. See [“Server Console Commands for Backup and Restore of Various Target Services” on page 31](#) for LOAD commands.
- ♦ Network traffic is heavy. Press Esc to return to the Main Menu, and then press Enter again. Your target should now be listed.

- 4 Select a target from the Target Services list.

- 5 When prompted, enter the username (and context, if required) for the target.

If SBICON rejects the username you entered, you probably need to include the context where the User object is located. For example, instead of entering **ADMIN** as the username, enter **.CN=ADMIN.O=company_name** or **.ADMIN.company_name**.

You must include the context in the username at this point if the username you use exists in an eDirectory container that is different from the bindery context set on the server running the engine.

- 6 (Conditional) If requested, enter the password for the target.

SBICON will take a few moments to attach to that target.

If you select the default values on this form, a full backup of the source is performed.

Otherwise, if you want to perform an incremental or differential backup, select Advanced Options. For more details on the Advanced options, see [“Customizing Backup” on page 37](#).

- 7 Click What to Backup and press Enter.

- 8 Press Insert, select the volumes and files you want to back up from the screen, and press Esc.

The selected volumes and files appear in the List Resources screen.

- 9 Enter a descriptive name for the backup session in the Description field.

This is a descriptive name of the job, which will help you identify the specific backup session; if a restore is necessary, you can easily identify the session you need. There is no special requirement for what to enter as the session description.

TIP: You might want to include the full path of the data (for example, SYS:HOME\REPORTS\JULY.02), which you will need to know if a restore session is necessary.

- 10 Click Device/Media Name and select a device.

A list of available objects is displayed in the Select a Device screen.

If you accept the default value, the engine selects the device and media which you have designated as the default.

NOTE: If there is only one media type available, it will be treated as the default.

If your device has more than one storage device attached, select an available device from the list. If the device contains multiple media, select the media for the backup.

If the storage media does not have a label, SBICON displays a message saying that the media cannot be identified. Set the label before selecting it for a backup job.

- 11** Select Advanced Options and set the desired values for all the fields.

For more details, see “[Customizing Backup](#)” on page 37.

NOTE: If you set Enable Auto Recovery to NO in Cluster Options, the remaining options become invalid.

- 12** Select Append Session.

This option appears only if your device supports appending to previous sessions on the media.

Accept the default YES if you want to preserve all sessions on the media. The backup session is appended to the media at the end of the previous session. Each appended session has separate backup and error logs.

If you want to overwrite the data on the media, enter NO. The media rewinds and existing data is overwritten by the next backup.

TIP: Labeled media is not considered empty media. No sessions are necessary on the media for you to set the Append option to Yes.

Delete the log and error files for any session you overwrite to avoid trying to restore a session that no longer exists.

- 13** Press Enter to save the settings and continue with the backup.

WARNING: To append a backup session to a media set (two or more tapes), use the tape with maximum space.

- 14** To return to the Main Menu, press Enter and then press Esc until you back out to the Main Menu.

- 15** To exit SBICON, press Esc again and answer the confirmation prompt.

Backing Up a Workstation

To back up the workstation (the target), a Target Service Agent (TSA) must be loaded on the workstation. (See “[Loading the Target Service Agents](#)” on page 31.)

You can select the directories to back up or back up the entire workstation.

Windows 95 and 98 Workstation

Make sure that the W95TSA is installed and configured. It is a part of the Novell Client modules (checkbox for Target Service Agent).

NOTE: If W95TSA is installed and registered, a shield-like icon appears in the system tray and indicates Novell TSA (Listening) message. You will also see an SMS screen popup when you log in to NetWare.

- 1** At the target workstation, register with the W95TSA.

- 1a** Double-click the shield-like icon that appears in the system tray.

The Novell Target Service Agent for Windows Properties box appears.

- 1b** Enter your user name and password for TSA.

- 1c** Enter the server name that will be servicing this Windows client.

- 1d** Select the disk drive you want to register.

1e (Recommended) Check Auto Register. This automatically registers your workstation with the server after you restart.

1f Click OK.

1g Restart your workstation.

The TSA must be installed and registered with the host server, along with the TSA username and password.

2 At the host server, load TSAPROXY.

3 Select the NetWare server running the target service agent.

4 Follow the procedure from [Step 3 on page 35](#) to complete your backup.

Windows 2000 and Windows NT Workstation

Make sure that the NT TSA is installed and configured. It is a part of the Novell Client modules (checkbox for Novell Target Service Agent).

The Windows NT TSA has two parts:

- ♦ **TSA Service** (TSAMAIN.EXE) is a service that waits for and acts when there is a request from the backup engine.
- ♦ **TSAPrefs** (TSAPREFS.EXE) allows an NT administrator to set preferences and monitor the activities of the TSA Service.

1 Right-click Network Neighborhood and then click Properties > Novell Target Service Agent > Properties.

2 Click Preferences and fill in the following fields:

Workstation: Enter the name of the workstation.

Preferred Server: Enter the name of the server on which you want to connect. This server must have TSAProxy installed.

Protocol: Select SPX/IPX or TCP/IP.

Events to Log: Check the items you want to log: Connection, Registration, Security, or Service.

Allow Backup User: This grants the backup user rights depending on the group.

Auto Register: This automatically registers the Windows workstation with the server after you restart.

3 Click Registration.

4 Check the fields in the Registration page and, if correct, click Register.

If information in any of the fields is incorrect, click Withdraw and return to the Preferences page to change the information. Then repeat Steps 4 and 5.

5 Click Connections.

6 Check the fields in the Connections page and then click OK.

If you need to add or change information, click Apply.

You can also delete connections in this page under Current Connections and then click Apply.

Changing Your Target

Any time you are performing a series of backup or restore sessions, you might need to change your target. This can be easily done from the SBCON Main Menu either at the command line or from the Windows workstation using NWBACK32.

If you do not specify a new target, SBCON automatically selects the target used during the last session and continues without asking you to make a selection.

- Ensure that you have met the prerequisites detailed in [“Prerequisites for SBCON Tasks” on page 57](#).
- Load the required TSA. (See [“Loading the Target Service Agents” on page 31](#).)

Changing a Target at a Server

- 1** Select Main Menu > Change Target to Back Up From or Restore To.
- 2** Select one server.
- 3** From the list of Target Service Agents, select a target.

If the target you select includes more than one TSA loaded, SBCON will show you a list of their full names.

- 4** Enter the username and password for the target you just selected.
The target is successfully changed.

Changing a Target at a Workstation

- 1** Run NWBACK32.
See [“Loading NWBACK32” on page 33](#).
- 2** At the Quick Access window, click Backup.
You can use the toolbar to open this window if preferred.
- 3** Double-click What to Backup.
- 4** Double-click Workstations (or NetWare Servers).
- 5** Select the workstation (or server) to use as the target.
The target is successfully changed.

Restoring Data

Use SBCON to restore information from tape backup. These procedures use SBCON as the Storage Management Engine (SME).

- Know the username and password for the target server or workstation to which you want to restore data.
- Know the session description that you want to restore. You can get this information from your session log and error files. (See [“Session Files” on page 26](#).)
- For an Advanced Options restore session, ensure that you know the file system structure of the data you are restoring. You will be prompted for specific paths and filenames while setting the restore options.

If you restore the server-specific information from the file system backup of the failed server to a functioning server or location, you can use the VOL\$INFO.TXT file as a reference for this information.

- ❑ If you want to restore data to a new location (different from where the original data was located), you must specify the full path to both the original data and the new location. If the new location does not exist, SBCON will create a new file system structure.

You can get the original path from your backup logbook or from the session log files if you noted the path at the time the backup was performed.

- ❑ Replace faulty hardware or correct the problems that caused data loss (if applicable).
- ❑ Ensure that media is inserted in your storage device.

You can restore data either from the server command line or from a Windows workstation using NWBACK32.

Restoring Data from a Server

- 1 Load the drivers for your specific device and controller board.
See [“Loading Controller and Storage Device Drivers on the Server” on page 31.](#)
- 2 Load the appropriate Target Service Agents.
See [“Loading the Target Service Agents” on page 31.](#)
- 3 Load SBCON.
See [“Loading SBCON” on page 32.](#)
- 4 Select Main Menu > Job Administration > Restore.
The Restore Options screen appears.



- 5 Select Target Service and then the NetWare server running your backup and restore Target Services Agent.
This is the name of the server or workstation on which you want to restore your data. Before selecting it, you must load the correct TSA. See [“Loading the Target Service Agents” on page 31.](#)
If you do not see the target you want on the list, check for these possible causes:
 - ◆ The proper TSA is not loaded on the target.
 - ◆ Network traffic is heavy. In these circumstances, it might take SBCON a few moments to detect all the targets. Press Esc to return to the Main Menu, and then press Enter again. Your target should now be listed. From the Target Services list, select a target.If the target you select has more than one TSA loaded, SBCON shows you a list of their full names in the Backup/Restore Target Services on This Server.

NOTE: When you are trying to connect to a different server, SBICON uses the same username and password you gave the first time. It prompts only if it fails.

The following table lists the target services and their associated datasets.

| Target Service to Restore | Dataset to Select |
|-----------------------------|--|
| Cluster-enabled pool | The cluster-enabled pool that you backed up |
| File system | The server whose file system you want to restore |
| eDirectory database | The server with TSANDS loaded |
| Workstation | The workstation's host server and then the workstation |
| Server-specific information | The server whose server-specific information you want to restore and then the file system to restore |

- 6 When prompted for the target username, enter your username (and context if required) as the network administrator for the target.

If SBICON rejects the username you entered, you probably need to include the context of where the user object is located. For example, instead of entering **ADMIN** as the username, enter a complete name such as **.CN=ADMIN.O=COMPANY_NAME** or **.ADMIN.COMPANY_NAME**.

You must include the context in the username at this point if the username exists in an eDirectory container that is different from the context set on the server running SBICON.

- 7 If a password is requested, enter the password for the target.

SBICON will take a few moments to attach to the target. Wait for the confirmation box, and then press any key to continue.

- 8 Enter a descriptive name for the restore session in the Description field.

The descriptive name of the job will help you identify the session.

- 9 Select the device and media for restore by pressing Enter at the Device/Media Name field and decide where you want to restore from.

A list of available objects is displayed in the Select a Device screen.

If you enter ***.* (DEVICE_NAME.MEDIA_NAME)**, the engine selects the device and media that it encounters first.

If your device has more than one storage device attached, select an available device from the list. If the device contains multiple media, select the media for the backup.

If the storage media does not have a label, SBICON displays a message indicating that the media cannot be identified. You should set the label before selecting it for a backup job.

- 10 On the Select a Session screen, select the Session to Restore.

WARNING: TSAFS.NLM backs up files in compressed or uncompressed format as specified by the user. However, if you try to restore a compressed file to a volume without compression, the file is corrupted and no error message is displayed.

- 11 Specify the path to the session log file of the session you want to restore.

Use one or more of the following methods:

- ◆ Press Enter to accept the default; then go to [Step 12 on page 46](#).
- ◆ Press Insert to select from a list of directories, press Esc to return to the previous window, and then go to [Step 12 on page 46](#).
- ◆ Press Backspace over the path shown (or a portion of the path), type in a new directory or path, and then go to [Step 12 on page 46](#).

Press Insert any time during this process to choose parts of the existing path from a list.

NOTE: At the present, it takes the default. This will be implemented in the future release.

- 12** When a list of sessions is displayed, select the session you want to restore.

If the media on which the data resides is not loaded on the device you choose, you will be prompted to insert the correct media.

- 13** Select Advanced Options to make the selective restore.

The Advanced Restore Options screen appears. See [“Customizing Restore” on page 46](#). You can perform the advanced restore, specify the execution time, reschedule the job, etc., using this form.

- 14** Complete the Advanced Restore Options form.

- 15** Press Esc and answer the prompt at the Submit Job screen to begin the restore session.

- 16** Press Alt+Esc to view the activity log screen.

If the job execution starts and is not successfully completed, the error messages appear on this screen. The activity log file ACTIVITY.LOG is located at the server's SYS:\SYSTEM\TSA\LOG directory.

- 17** To view the Run Time Status of a job, select Job Administration > Current Job List > *job name*.

- 18** When the restore session finishes, your screen display reads the following:

```
The restore process was completed normally.
```

- 19** To return to the Main Menu, press Esc as required.

The Esc key allows you to exit SBCON one screen at a time until you reach the Main Menu. If you want to exit SBCON, answer the confirmation prompt.

NOTE: To unload SBCON and applicable Target Service Agents, see [“Unloading SBCON” on page 33](#).

Customizing Restore

Using the Advanced Restore Options form at the server, you can perform a custom restore of the database you need.

- 1** Select Main Menu > Job Administration > Restore.

The Restore Options screen appears.

- 2** Select Advanced Options.

- 3** Select Rename Data Sets to restore data to a location different from the backup location.

- 4** At the Restore Datasets to a Different Location screen, do the following:

4a Press Insert to select the namespace type.

4b Enter the source path (where the data originally was backed up from).

4c Enter the destination path you want the data restored to, in the entry box.

Both the source and destination paths must include the volume names (for network server) or the hard disk drive letter (for example, C:\) for workstations.

IMPORTANT: When you specify a particular portion of the file structure as the source location, it does not necessarily mean that this will be the only data restored. What is restored can be influenced by the include and exclude settings on the Subsets of What You Want to Restore form. If you do not want to overwrite any subdirectories that might exist in the area you have included, you must specifically exclude them by their full path.

5 Select Subsets of What to Restore.

Use this screen to specify which parts of the file structure you want to restore. Various data sets to be included or excluded are listed in this form. You can include or exclude any or all items listed. For more information about including and excluding items, see [“Exclude and Include Options” on page 12](#).

The screen lists major and minor resources. *Default* indicates that none of that type of subset has been specified.

6 To specify a subset, do the following:

6a Select any item on the screen.

6b Edit the Include or Exclude screens using Insert and Enter.

6c Press Esc to save changes.

6d To view a list of values associated with major TSA resources (for example, server or volumes) select either the include or exclude line.

The Selection List Options form appears. Any value you previously selected is shown on the form. If you did not select a value, the form is blank.

6e To select a value for this major resource, press Insert and select the value from the list that appears by pressing Enter.

7 Select Open Mode Options and select parts of the backed up data structure to be restored.

The backed-up data set contains various types of data that are listed on the screen. By default, all the types of data are included in the restore.

7a To exclude a particular type of data from the restore, select that line and type Y.

7b When all the data items are selected correctly, press Esc.

8 Specify whether to overwrite the existing parent by typing Y (Yes) or N (No).

For more information, see [“Overwriting a Parent or Child” on page 15](#).

9 Specify whether to overwrite the existing child by typing Y (Yes) or N (No).

10 Specify the execution time.

11 Select Scheduling to reschedule the job.

12 Press Esc as required and then select Yes to confirm the Submit Job? prompt.

The job is submitted and is listed in the Current Job List menu.

Restoring Data from a Workstation

When you restore data, you are copying your backed up data to a server or workstation.

1 Run NWBACK32.

See [“Loading NWBACK32” on page 33](#).

You can run the NWBACK32 program (NWBACK32.EXE) from SYS:PUBLIC.

- 2** At the Quick Access window, click Restore.
- 3** Click What you Want to Restore.
 - 3a** Double-click What to Restore.
 - 3b** Select the context.
 - 3c** Double-click Queues and select the queue.
 - 3d** Double-click Servers and select a server.
 - 3e** Double-click Devices and select a device.
 - 3f** Double-click the device to view the media list.
 - 3g** Select a medium.
- 4** Click where you want the restored data to go.
 - 4a** Double-click Where to Restore.
 - 4b** Select eDirectory, NetWare Servers, or Workstations.
 - 4c** Click the server to which you want to restore your data.
 - 4d** Enter your username and password to authenticate to the server.
 - 4e** Click OK.
- 5** Submit the restore job by clicking Restore > Submit the Job on the toolbar.
- 6** (Optional) Click Next to display more information about the restore.

Restoring eDirectory

The procedures in the following sections are performed at the beginning of the restore process.

- ◆ [“Restoring a Volume Other Than SYS:” on page 48](#)
- ◆ [“Restoring Volume SYS: in a Single-Server Network” on page 49](#)
- ◆ [“Restoring Volume SYS: in a Multiple-Server Network” on page 49](#)
- ◆ [“Restoring the Entire eDirectory Tree Structure” on page 50](#)

Restoring a Volume Other Than SYS:

WARNING: Do not delete the Volume object for the failed volume from the eDirectory tree structure. The Volume object preserves any references that other objects (such as Directory Map and Queue objects) might have to that volume. If the Volume object is deleted and you have objects that depend on this volume, you must re-establish the relationships through a selective eDirectory restore. See [“Custom eDirectory Restore” on page 51](#).

- 1** Bring down the server.
- 2** Correct the problem.
- 3** Bring the server back up and re-create the volume.
- 4** Perform an advanced file system restore.

See [“Customizing Restore” on page 46](#) for the procedure.

Restoring Volume SYS: in a Single-Server Network

Since there are no replicas in a single-server network, you cannot recover any eDirectory information from a replica. After repairing or replacing the failed hardware, you must restore the entire NetWare environment, including eDirectory, from an SMS backup.

To restore a single SYS: volume, do the following.

- 1** Correct the problem that caused volume SYS: to fail.
- 2** Reinstall NetWare and eDirectory.

When NWCONFIG asks for the names for Organization objects immediately under the [Root] object, use the same names that existed before in the eDirectory tree. Otherwise you will end up with new, empty containers in the restored eDirectory tree structure.

Make sure disk partitions are at least as large as they were before, and that the volumes are defined as before.

- 3** Restore eDirectory using [“Customizing Restore” on page 46](#).

Restoring Volume SYS: in a Multiple-Server Network

To restore SYS: volumes on several servers, do the following.

- 1** Restore the server-specific eDirectory information file (SERVDATA.NDS) to another server on the network following the custom (advanced) file system restore (see [“Customizing Restore” on page 46](#)).

By default, the file is placed into the SYS:SYSTEM*server_name* directory on the server you have selected.

- 2** Use the DSMISC.LOG file (part of the information restored with the server specific information) to see if the failed server had a master replica of any partition.
- 3** If the failed server had a master replica of any partition, select the server you want to hold the master replica and load DSREPAIR on that machine.

3a Select Available Options > Advanced Options Menu.

3b Select Replica and Partition Operations.

NOTE: Use the NDS Manager utility to perform regular partition operations.

3c Select the partition you want to edit.

3d To see the list of servers that have replicas of that partition, select View Replica Ring.

3e Press Esc.

3f Select Designate This Server as the New Master Replica.

3g Repeat Steps 3a through 3f for each master replica that the failed server contained.

- 4** (Conditional) If the failed server contained any replicas, remove them.

4a Select Available Options > Advanced Options > Replica and Partition Operations.

4b Select the partition the server held replicas of.

4c Select View Replica Ring.

4d Select the name of the failed server.

4e Select Remove This Server from the Replica Ring.

4f Enter the supervisor name and password.

- 4g** Select Yes to continue.
- 5** Exit DSREPAIR.
- 6** Reinstall NetWare.
 - 6a** Begin the installation.
 - 6b** When prompted for DS install, select Create a New DS Tree and complete the installation.
 - 6c** Load NWConfig and remove DS.
 - 6d** Install DS, when prompted for DS install and select Install into an Existing Tree.
 - 6e** When prompted to choose a eDirectory tree structure, press F5.

This allows you to choose to restore the SERVDATA.NDS file to the failed server.
 - 6f** To specify the path where the SERVDATA.NDS file exists, press F3 and enter the path.

TIP: If the information fits on a diskette, you can save the information on a diskette instead of another server on the network.
 - 6g** Press Enter to accept the path.
 - 6h** Log in to the server that contains the SERVDATA.NDS file.
 - 6i** Log in to the eDirectory database.
 - 6j** If the START.NCF file of the failed server is different from the default, you can edit the default file.
 - 6k** You now can choose from the following:
 - ◆ To continue the install and not copy the files from install, but use the backup tapes to finish restoring the information, press Enter.
 - ◆ To go back to the previous screen, press Esc.
 - ◆ To copy the files from install and then restore the files from tape, press F3.
- 7** Restore the file system using “Restoring Data” on page 43.

Restoring the Entire eDirectory Tree Structure

To restore an entire network from a full backup in a multiple-server environment, do the following:

- 1** Reinstall NetWare on the first server.

By default, this server will hold the master replica of the [Root] partition.

When NWCONFIG asks for the names of Organization objects immediately under the [Root] object, use the same names that existed before in the eDirectory tree.

In NWCONFIG, select Directory Options > Directory backup and restore options to obtain the information.

Otherwise you will end up with new empty containers in the restored eDirectory tree.

When this installation is complete, you will have a working eDirectory tree containing one NetWare server with a master [Root] partition.
- 2** Add name spaces as required.
- 3** Install the remaining servers to complete a skeleton of your network.

Before restoring a full eDirectory session, you should create a skeleton of your network.

- ◆ All servers and volumes should be up and running.
- ◆ The eDirectory objects should exist in the eDirectory tree in the same context as before. (NWCONFIG prompts you for the container in which you want each server to be placed.)
- ◆ Make sure all servers are communicating with one another.
- ◆ Make sure time synchronization is working properly.

NOTE: If you can get some, but not all of the servers backed up, you can still proceed with the restoration. However, you might see errors and experience problems due to eDirectory objects having dependencies that cannot be resolved.

The User object (Admin or equivalent) used to create the backup session must exist in the same container with the same password and eDirectory rights as when the backup was performed.

Once this step is completed, you still have just one partition: [Root]. Because of the NWCONFIG program defaults, you now have two replicas of that partition. These are stored on the second and third servers you installed.

- 4** Restore eDirectory on one server using [“Restoring Data” on page 43](#).

eDirectory Restore Session Types

You can choose to restore eDirectory in one of the following ways:

- ◆ [“Single Server or an Entire Tree Restore” on page 51](#)
- ◆ [“Custom eDirectory Restore” on page 51](#)
- ◆ [“Custom File System Restore” on page 52](#)
- ◆ [“Server-Specific Information Restore” on page 53](#)
- ◆ [“Single File or Directory Restore” on page 54](#)

Single Server or an Entire Tree Restore

- 1** Restore the entire eDirectory session from your SMS backup.

You will have no session file to work from at this point. Choose the option to restore without session files.

- 2** Continue through [Step 15 on page 46](#) to [Step 19 on page 46](#) to restore eDirectory.

Custom eDirectory Restore

To perform a custom eDirectory restore, do the following.

- 1** Enter SBCON at the server console prompt.
- 2** Select Job Administration.
- 3** Select Restore Options > Advanced Restore Options > Subsets of What to Restore.
- 4** To specify a subset, do the following.
 - 4a** Select any item on the screen.
 - 4b** Edit the Include or Exclude screen using Insert and Enter.
 - 4c** Press Esc to save your changes.
- 5** Specify whether or not to overwrite the existing parent or child.

A parent is a container. A child is an object, which is the lowest level of the eDirectory tree.

If you want to overwrite the parent or child that exists on the hard disk, regardless of whether the existing version or the backed-up version has the latest date, accept YES for parent, child, or both.

If you do not want to overwrite the parent or child that exists on the hard disk, regardless of whether the existing version or the backed-up version has the latest date, select the option (parent or child or both), and enter N.

- 6 Continue through [Step 14 on page 46](#) to [Step 19 on page 46](#).

Custom File System Restore

To perform a custom file system restore, do the following.

- 1 Enter SBCON at the server console prompt.
- 2 Select Job Administration > Restore Options > Advanced Restore Options > Subsets of What to Restore.

The screen lists major and minor resources. The word *DEFAULT* indicates that none of that type of subset has been specified.

- 3 To specify a subset, do the following.
 - 3a Select any item on the screen.
 - 3b Edit the Include or Exclude screen using Insert and Enter.
 - 3c Press Esc to save your changes.
- 4 (Conditional) If you want to exclude certain types of data from the restore session, select Job Administration > Restore Options > Advanced Restore Options > Open Mode Options, and then edit the form.

The backed-up data set contains the various types of data items listed on the screen. By default, all types of data items are included in the restore.

If you want to exclude any item, select it and enter Y. Press Esc when you are finished editing the form.

- 5 (Optional) Restore data to a different workstation or location on a different NetWare server.
 - 5a Back up the data from the server where it resides.
 - 5b Exit to the Main Menu.
 - 5c Select Change Target to Back Up From or Restore To screen from the Main Menu.
See [“Changing Your Target” on page 43](#).
- 6 (Optional) Restore data to a different workstation or location within the backed-up server's file system structure.

Follow these important rules about restoring to a new location:

- ◆ When you specify a particular portion of the file system structure as the source location, it does not necessarily mean that this will be the only data restored. You can influence what is restored by using the Include and Exclude options on the Choose Subsets Of What You're Restoring screen (Job Administration > Restore Options > Advanced Restore Options > Subsets of What to Restore).
- ◆ You can restore only to a new location in the name space that is supported by the drive you are restoring to.

- ◆ If you do not want to overwrite any subdirectories that might exist in the area you have included, you must specifically exclude them by their full path.

6a Go to Subsets of what you want to restore, press Enter, and then Insert, select the name space you want, and enter the name you used for your backup at the Enter the Data Set Name field.

6b Enter the full directory path of the source including the volume name.

6c Enter the full directory path to the new destination.

6d Press Esc to save your changes.

NOTE: The eDirectory Target Service Agent (TSANDS) does not support restoration of data to a different location.

7 Specify whether to overwrite the existing parent or child.

A parent might be a server, a volume, or a directory. A child is a file, which is the lowest level of the file system structure.

If you want to overwrite the parent or child that exists on the hard disk, regardless of whether the existing version or the backed-up version has the latest date, accept YES for parent or child or both.

If you don't want to overwrite the parent or child that exists on the hard disk, regardless of whether the existing version or the backed-up version has the latest date, select the option (parent or child or both), enter N.

If you want to overwrite the child that exists on the hard disk only if the date of the backed-up copy on the media is later than the date of the copy on the hard disk, select Overwrite Existing Child and select the option.

8 Continue through [Step 15 on page 46](#) to [Step 19 on page 46](#).

Server-Specific Information Restore

To view and restore server information, do the following.

1 Select Main Menu > Job Administration > Restore Options > Advanced Options > Subsets of What to Restore > Include Major TSA Resources.

The Selection List Options screen appears and is empty.

2 Press Enter to see the Selection List Options. Press Insert to see the server specific information.

3 Select the Server-Specific Info and press Enter.

Esc returns you to the Selection List Options.

This will return you to the Choose Subset of the Session to be Restored screen.

4 Complete the form and press Esc.

This will return you to the Restore Options screen.

5 Complete the Restore Options form and press Esc.

The server-specific information will be restored to the SYS:SYSTEM\failed_server_name directory of the server on which you are running SBCON.

6 Continue through [Step 15 on page 46](#) to [Step 19 on page 46](#).

Single File or Directory Restore

- 1 Select Main Menu > Job Administration > Restore Options > Advanced Options > Subsets of What to Restore.
- 2 Enter the full path and name of the file or directory, including full directory and volume names.

The following options appear:

Include/Exclude directories (full path)

Include/Exclude files

Include/Exclude path/files

- 3 Press Insert and enter the name space for the name space format you used when entering the file or directory name.

The name space allows SBCON to correctly interpret the path information entered earlier in this menu in either the directory field or the location field.

IMPORTANT: If files are restored to a new location, the name spaces must be the same or an error message will occur.

You can restore to a new location only in the name space that is supported by the drive to which you are restoring.

The supported name spaces are DOS, Macintosh, NFS, and Long.

SBCON provides a list of the name spaces currently loaded on the target to which you are restoring.

For example, if you are restoring to a new Macintosh location, enter a path, (for example, **SYS::test:monday**) and then enter **Macintosh** as your name space.

- 4 (Optional) Restore to a new location (located in Rename Data Sets).
 - 4a Enter the source path, including the volume name, to the target directory (a location different from where the data was backed up). Enter **Volume:/directory/directory**
 - 4b If the new location you are restoring to is a different NetWare server, do the following:
 1. Back up the data from the server where it resides.
 2. Exit to the Main Menu.
 3. Change your target to the server you want to restore the data to. See [“Changing Your Target” on page 43](#).
- 5 Continue through [Step 15 on page 46](#) to [Step 19 on page 46](#).

Completing the eDirectory Restore

If you restored the eDirectory database to volume SYS:, you should complete the restoration for the following to make sure the eDirectory database is working properly:

Volume Other than SYS:

To finish restoring eDirectory on a volume other than SYS:, do the following.

- 1 Bring the server up.

- 2 (Optional) Verify proper restoration of the data, trustee assignments, file ownership, and other related information by spot-checking some of the restored directories and files.

Single-Server Network

To finish restoring eDirectory on a server, do the following.

- 1 Restore the full file system following “Restoring Data” on page 43.
- 2 (Optional) Verify proper restoration of the data, trustee assignments, file ownership, and other related information by spot-checking some of the restored directories and files.

Multiple-Server Network

To finish restoring eDirectory on multiple servers, do the following.

- 1 Restore replicas to the server.

The DSMISC.LOG file, which is created when the SERVDATA.NDS is restored, contains the partition and replica information that resided on the failed server at the time of the backup. This will show you what replicas were on the server when it failed.
- 2 (Optional) Verify proper restoration of the data, trustee assignments, file ownership, and other related information by spot-checking some of the restored directories and files.

Commands that might be helpful include RIGHTS /T /S (displays users, groups and other objects that have explicit trustee assignments in a directory and its subdirectories) and NDIR (displays owners and other NetWare file information).

Entire Tree

To finish restoring eDirectory on the tree, do the following.

- 1 Restore the file system information to each server following “Restoring Data” on page 43.
- 2 Re-create partitions and replicas.
- 3 (Optional) Verify proper restoration of the data, trustee assignments, file ownership, and other related information by spot-checking some of the restored directories and files.

Commands that might be helpful include RIGHTS /T /S (displays users, groups and other objects that have explicit trustee assignments in a directory and its subdirectories) and NDIR (displays owners and other NetWare file information).

Partial eDirectory Restores

Certain conditions might arise in special cases involving the backup and restoration of eDirectory information.

The SMS TSA software allows you to do selective restores from the backup media. However, partial restoration of eDirectory from a backup can have many subtle consequences, particularly when only a single object or a selected group of objects is restored.

For partial eDirectory restores, keep these two main issues in mind:

- ♦ **Object ID numbers**—If you restore objects that no longer exist in the eDirectory tree, those objects receive new ID numbers when restored. New object IDs affect file system trustees, print queue directories, user mail directories, etc.

If you restore objects on top of objects that exist in the eDirectory tree, the objects do not receive new ID numbers. These objects' current attribute and property information is overwritten with previous information from the SMS backup.

- ◆ **Objects that depend on other objects**—In the eDirectory schema, objects are defined to have certain attributes. Some of these attributes are mandatory (meaning they must contain a value); others are optional.

For some eDirectory objects, the value for a particular attribute is a reference to another object upon which the object depends. For example, Queue object has a Queue Directory attribute that contains the file system path to the queue directory. It also has a Host Server attribute that identifies the file server on which the queue directory resides. This information is used to determine the physical location of the resource.

The specifics of restoring objects vary depending on what type of object is involved and whether the object's dependencies are physical entities (servers and volumes) or logical entities.

In some cases, you can simply restore an object and everything will work fine. In other cases, an object might be restored but not be functional unless you first restore its dependent objects.

Restoring Cluster-enabled Pools

Restore of cluster-enabled pools is similar to a normal restore session.

See [“Restoring Data” on page 43](#).

NOTE: Auto recovery of the restore session on failover or failback is currently not supported.

3

Managing

After you install and load SBICON, you can manage your jobs, data, devices, reports, and so on either from the server or workstation. Though the screens might look different, you can perform the same tasks in each interface.

Be sure to read [“Prerequisites for SBICON Tasks” on page 57](#) before performing any tasks in this section.

This section discusses the following:

- ◆ [“Administering Jobs” on page 57](#)
- ◆ [“Storage Devices” on page 61](#)
- ◆ [“Erasing or Retaining Media” on page 63](#)
- ◆ [“Spanning Storage Media” on page 65](#)
- ◆ [“Log and Error Files” on page 66](#)
- ◆ [“Creating Session Files” on page 69](#)
- ◆ [“Verifying Backup Data” on page 70](#)

Prerequisites for SBICON Tasks

Before you can manage any backup tasks, ensure the following:

- Drivers for your specific device and controller board are loaded.
- SBICON files for your specific target are loaded.
- The media is inserted into your storage device.
- You can see the SBICON Main Menu on your server console or the Quick Access window on your Windows 95, 98, 2000 or Windows NT workstation.

Jobs

The backup or restore jobs created by you can be administered anytime. You can delete, hold, reschedule, abort, or enable a job.

Administering Jobs

Jobs can be managed either from the server or from the workstation.

From the Server

On the SBCON Main Menu at the server, the Current Job List option gives you a job management facility where you can delete a job, hold a job, or enable it if it has been disabled after a job failure. You can even monitor the status of the backup/restore jobs you have submitted.

Current Job List provides the following information about all jobs submitted.

- ◆ **The name of the server from which the job was submitted.** If you initiated the job from the client, locate the workstation name. See [“Viewing the Job Queue” on page 58](#).
- ◆ **The job name that you entered on the Backup Options form.** The job name can be renamed using the Job Details screen if you are rescheduling it to run another time and you want to change the name. See [“Viewing Job Details” on page 59](#).
- ◆ **The Hold status.** If you want to prevent a prescheduled job from execution, you can put it on Hold. See [“Postponing a Job” on page 60](#).
- ◆ **The job status, whether it finished the execution of the request, or if it was submitted for execution at a later time.** Other statuses are: active currently, could not run, processing, failed, and success. The log of each of these status reports can be viewed using the log and error reporting. For more details see [“Log and Error Files” on page 66](#).

From a Workstation

From Job Administration at the Quick Access window of a Windows workstation or at the toolbar, you can access the current list of jobs. If you select a job, you can see the properties, which include the following:

- ◆ **General properties** contains fields for the Job Name, Job Type, Submitted By, Entry Open, and Serviced By. These fields are the same as the Job Details screen at the server.
- ◆ **Static properties** contains fields for the Job Status, Job Submitted At, Execution Time, and Job's Rescheduled Count.
- ◆ **Dynamic properties** contains data that indicates how much data has been backed up along with the filenames. This dialog box appears when a job is running and SBSC is loaded on the server.
- ◆ **Enable job** allows you to start a job that has been disabled.
- ◆ **Delete the job** allows you to remove a job.
- ◆ **Start/Hold the job** allows you to start a job that is in the hold state or place it on hold if it is not in a hold state.
- ◆ **Reschedule options** contains fields for changing the backup schedule.
- ◆ **Abort** cancels an active job.

This section discusses the following:

- ◆ [“Viewing the Job Queue” on page 58](#)
- ◆ [“Viewing Job Details” on page 59](#)
- ◆ [“Postponing a Job” on page 60](#)

Viewing the Job Queue

You can view the job queue from the server or a workstation.

Ensure that you have met the prerequisites detailed in [“Prerequisites for SBCON Tasks” on page 57](#).

From the Server

- 1 Select Main Menu > Job Administration > Current Job List.

The Queue Jobs screen displays a list of all jobs, showing each job name, hold status, status of the job, and name of the server from where the job was submitted.

From a Workstation

- 1 At the Quick Access window, click Job Administration
- 2 Double-click the context of the queue location.
- 3 Double-click Queues.

A list of jobs appears.

- 4 (Optional) Select a job from the list > right-click and then click Dynamic Properties to view the dynamic properties of the job.

You can also view the properties using the toolbar. Click Job Administration > Properties.

Viewing Job Details

This screen gives you all the details of the job you submitted including job name, the time the job was submitted, and the time the job should be executed. You can also schedule the job to run at a predefined time again after its execution.

You can view job details from the server or a workstation.

Ensure that you have met the prerequisites detailed in [“Prerequisites for SBCON Tasks” on page 57](#).

From the Server

- 1 Select Main Menu > Job Administration > Current Job List.
- 2 Select one of the jobs listed and press Enter.

The Job Details screen appears displaying the following information:

- ♦ **Job Name**—A descriptive name of the job for your identification. You can change the name if you want to reschedule the job and run it again.
- ♦ **Disabled**— Indicates whether the job is disabled. If it is, this field displays **Yes**; if enabled, it displays **No**. Updates to this field depend on the job status. If the job status is **Could Not Finish**, the message is **Yes**. You can perform the required troubleshooting and change this field to **No**, subsequently allowing the job to execute. After completion of the job, the Disabled field indicates **Yes**.
- ♦ **Submitted By**—The name of the server where you have submitted the job. If you submitted the job from a client, it displays the name of the workstation and is the same as the Owner in the Queue Jobs screen.
- ♦ **Submitted at**—The date and time the job was submitted. The execution time could be different depending upon the specifications you defined in Execution Time and Scheduling Options in the same screen.

- ◆ **Entry Open**—The status is `Yes` when the job is active and `No` when it is inactive. This is one of the fields that is updated automatically.
- ◆ **Execution Time**—An option to give the time and date when you want the job to run. It is the same as the Execution Time in the Advanced Backup Options form.
- ◆ **Serviced By**—This indicates the name of the server where the job is to be executed. From a single server, you can submit jobs to various servers. This field indicates which job is serviced by which server. This field is active only when the job is active; otherwise, it indicates `None`.
- ◆ **Scheduling Options**—You can schedule the same job to run at another time. For information, see [“Customizing Backup” on page 37](#).
- ◆ **Session Report**—Press Enter at this field to view the session log report or the session error report for the particular session you selected. The same information can be seen in the Error Log File where it gives a log of all the jobs. For information, see [“Viewing a Log File” on page 66](#).

From a Workstation

- 1** At the Quick Access window, click Job Administration.
- 2** Double-click the context.
- 3** Double-click Queues.
- 4** Select the job from the list.
- 5** Right-click the job and then Dynamic Properties to view the dynamic properties of the job.
(Optional) You can view the properties using the toolbar. Click Job Administration > Properties.

Postponing a Job

This task lets you to place a job on hold from either the server or a workstation.

Ensure that you have met the prerequisites detailed in [“Prerequisites for SBCON Tasks” on page 57](#).

From the Server

- 1** Select Main Menu > Job Administration > Current Job List.
The Queue Jobs screen appears with a list of jobs.
- 2** Select a job and press Tab.
The job is on hold if `Hold` is displayed in the column. If no status is displayed, it is ready for execution at the time you configured it. Press Tab to toggle holding a job.

From a Workstation

- 1** At the Quick Access window, click Job Administration.
- 2** Double-click the context.
- 3** Double-click Queues.
- 4** Select the job to place on hold.
- 5** Right-click the job and then click Start/Hold the Job.

(Optional) You can also hold the job using the toolbar. Click Job Administration > Start/Hold the Job.

Storage Devices

The Storage Device Administration option of the SBCON Main Menu at the server or the Device Administration button at a Windows workstation lists the devices and the media you can access.

You can perform the following administrative tasks:

- ◆ “Checking the Device Status” on page 61
- ◆ “Changing the Device Label” on page 62
- ◆ “Viewing the Media List” on page 62
- ◆ “Spanning Storage Media” on page 65

SBCON supports autoloader, which holds multiple media within a magazine. Every magazine has slot numbers given by the manufacturer and each slot holds a single medium. This is extremely useful when you are backing up a large server. If one medium is fully written, the SME automatically writes to the other available media.

IMPORTANT: Any media under a device becomes inaccessible if a job is running on it because it is reserved by the engine. Any job scheduled to be executed at that instance will fail.

Checking the Device Status

The Device Status screen displays information about the current status of a device, such as read, write, and format from the server or workstation. It also gives information on the storage capacity.

Ensure that you have met the prerequisites detailed in “Prerequisites for SBCON Tasks” on page 57.

From the Server

- 1 Select Main Menu > Storage Device Administration.
- 2 Select a device and press Insert.
- 3 Select Utilities > Device Status.

The Device Status screen is displayed listing the following information:

- ◆ **Current Operation**—Indicates whether the device is currently reading, writing, formatting the media, or none of these.
- ◆ **Device Mode**—Indicates whether the device is selected for reading, writing, or both.
- ◆ **Maximum Capacity**—Indicates the total storage capacity of the media currently present in the device.

- 4 Press Esc when you are finished to return to the Main Menu.

From a Workstation

- 1 At the Quick Access window, click Device Administration.
- 2 Double-click the context.
- 3 Double-click Queues.
- 4 Double-click the preferred queue.

- 5** Double-click Servers.
- 6** Double-click the preferred server.
- 7** Double-click Devices.
- 8** Right-click the device and then click Properties to view the Device Label or Device Type.
(Optional) You can also view the properties using the toolbar. Click Device Administration > Properties.

Changing the Device Label

The default names for storage devices are controlled by the device manufacturers. These names are often not descriptive, so you change the name of the device during multiple backup sessions. However, this change is temporary, and it is lost after a reboot. If you have several devices, it might be difficult to remember which name goes with which device. For this reason, renaming storage devices is helpful.

You can change a device label from the server or a workstation.

Ensure that you have met the prerequisites detailed in [“Prerequisites for SBCON Tasks” on page 57](#).

From the Server

- 1** Select Main Menu > Storage Device Administration.
- 2** Select a device name and press Insert.
- 3** Select Utilities > Change the Device Label.
- 4** Enter the new name in the field.
- 5** Press Esc as required to return to the Main Menu.

From a Workstation

- 1** At the Quick Access window, click Device Administration.
- 2** Double-click the context.
- 3** Double-click Queues.
- 4** Double-click the preferred queue.
- 5** Double-click Servers.
- 6** Double-click the preferred server.
- 7** Double-click Devices.
- 8** Right-click the device to change.
- 9** Click Properties > Device Label.
- 10** Change the name in the Device Label field.
- 11** Click OK.

Viewing the Media List

This task allows you to view your list of devices from the server or a workstation.

Ensure that you have met the prerequisites detailed in [“Prerequisites for SBCON Tasks” on page 57](#).

From the Server

- 1 Select Main Menu > Storage Device Administration.

A list of devices is displayed.

If no device is displayed in the list, run **LIST DEVICES** at the server prompt.

- 2 Select a device name and press Enter.

The List of Media screen appears displaying the following information:

- ♦ **Slot**—The slot number given by the manufacturer.
- ♦ **Media name**—The name of the media, which can be changed using the Utilities menu. See [“Erasing or Retaining Media” on page 63](#).
- ♦ **Number**—The sequential number of the media (tape) in a particular backup set. SBCON automatically labels and appends an incremental number to any backup session that spans multiple media.

TIP: It takes a lot of time to display the media in an Autoloader. However, once the media list is displayed, press Tab to access it. This will not rescan the device for new media. Press Enter only when you remove or add media.

From a Workstation

- 1 At the Quick Access window, click Device Administration.
- 2 Double-click the context.
- 3 Double-click Queues.
- 4 Double-click the preferred queue.
- 5 Double-click Servers.
- 6 Double-click the preferred server.
- 7 Double-click Devices.
- 8 Right-click the preferred device.
- 9 View the media list.

Erasing or Retaining Media

This task erases media headers, erases all data on the media, or retentions media in tape devices.

Ensure that you have met the prerequisites detailed in [“Prerequisites for SBCON Tasks” on page 57](#).

The following table describes all media options.

| Option | Explanation |
|------------------------|--|
| Change the Media Label | This is permanent. Blank media (tape) does not have a label. Before submitting a job to the media, you need to label it. |
| Erase the Media | This option erases the tape completely. It might take up to two hours, depending on the size of the media. |

| Option | Explanation |
|------------------------|---|
| Erase the Media Header | This is a quick method of making the media appear blank. Although it is less secure than erasing all the data, this method is quick and effective. |
| Media Status | This option lets you view the status of the media. |
| Move the Media | This option moves the media from one slot of an autoloader to another slot. It is not a valid operation on a single media device. |
| Retention of Media | This option is for tape devices that allow retention (for example, 0.25-inch cartridges). This rewinds the tape, which might resolve a problem reading the tape. An error appears if the driver does not support this method on the device. |

From the Server

- 1** Select Main Menu > Storage Device Administration.
A list of devices is displayed.
- 2** Select one device from the list.
- 3** Select one media and press Insert.
- 4** Select Utilities > Erase the Media [or] Retention of Media.

From a Workstation

- 1** At the Quick Access window, click Device Administration.
- 2** Double-click the context.
- 3** Double-click Queues.
- 4** Double-click the preferred queue.
- 5** Double-click Servers.
- 6** Double-click the preferred server.
- 7** Double-click Devices.
- 8** Right-click the media > click Erase Media [or] Retention.
(Optional) You can also use the toolbar. Click Device Administration > Media > Erase Media [or] Retention.

Checking the Media Status in a Device

This task lets you check the status of a particular media type from either the server or a workstation.

Ensure that you have met the prerequisites detailed in [“Prerequisites for SBCON Tasks” on page 57](#).

From the Server

- 1** Select Main Menu > Storage Device Administration.
- 2** Select one device from the list.
- 3** Select one media and press Insert.

4 Select Utilities > Media Status.

A screen is displayed listing the following information:

- ◆ **Number in Media Set**—The sequential number of the media within a particular media set.

SBCON automatically labels and appends an incremental number to any backup sessions that span multiple media.

- ◆ **Creation time**—The time the current media label was first used on this media.
- ◆ **Mount status**—Whether the media is Mounted (ready for reading or writing). There are two status messages:

Mount Pending: The request to mount is waiting to be processed.

Not Mounted: The media not ready to access.

- ◆ **Media mode**—Whether the media is selected for reading, writing or both, or whether the mount request is pending.
- ◆ **Media type**—The type of media in the device, such as 4mm DDS (Digital Data Storage) tape, or 8mm tape.
- ◆ **Total capacity**—The total capacity of the media, if known. Sometimes it is not displayed, depending on the media manager.

The type of media manager controls the display of the total capacity of the media.

5 Press Esc when you are finished to return to the Main Menu.

From a Workstation

- 1 Click Device Administration at the Quick Access window.
- 2 Double-click the context.
- 3 Double-click Queues.
- 4 Double-click the preferred queue.
- 5 Double-click Servers.
- 6 Double-click the preferred server.
- 7 Double-click Devices.
- 8 Right-click the preferred media > Properties > Media Status.

Spanning Storage Media

If the storage media you are using does not have enough space for the entire backup, SBCON will span the media, prompting you to insert additional media when the first medium is full.

TIP: Always have extra media on hand in case the backup session spans more than one medium.

SBCON designates the first media as label #1 and increments any subsequent medium with the same label.

WARNING: To help safeguard your network, SBCON disconnects from the target when a delayed backup is complete. If the delayed backup session cannot fit on the media, SBCON prompts you to insert additional media. If additional media are not inserted, the backup does not finish and SBCON does not disconnect, thereby compromising security.

Log and Error Files

SBCON keeps a list of all the log and error files, which show specific information about the backup and restore session. See [“Session Files” on page 26](#) for more information.

You can manage these files as follows:

- ◆ View information contained in the log and error files.
See [“Viewing a Log File” on page 66](#) and [“Viewing an Error File” on page 67](#) for information.
- ◆ Create these files.
See [“Creating Log and Error Files” on page 68](#) for information.
- ◆ Change the location of these files.
See [“Setting the Location of Backup Log and Error Files” on page 68](#) for information.
- ◆ Delete these files.
See [“Deleting the Log Files” on page 68](#) for information.

Viewing a Log File

This task lets allows you to view your log file for backup and restore information from either the server or a workstation. The log file is created on the host server the first time data is backed up or restored and contains details about the session.

Ensure that you have met the prerequisites detailed in [“Prerequisites for SBCON Tasks” on page 57](#).

From the Server

- 1** From the SBCON Main Menu, select Log File Administration > View a Log File.
- 2** If necessary, set the location of the log and error files for a backup session. (This option is not supported for a restore session.)

To set the location, use one or more of the following methods:

- ◆ Press Enter to accept the default location, and then go to [Step 3 on page 66](#).
- ◆ Press Insert to select from a list of network directories, select a directory, and then press Esc to return to the previous window, and then go to [Step 3 on page 67](#).
- ◆ To change the directory, delete the existing directory path and type a new one, and then go to [Step 3 on page 66](#).

You can also press Insert during this process to select existing parts of the path from a list.

The backup log displays all supported name space types. The creator name space is the first in the group and a right-angle bracket (>) appears next to it.

If you have linked UNIX files, SBCON retains both symbolic and hard links created under the NFS name space type, but links are not listed in the backup log. Restoring the original file restores associated links.

- 3** To return to the Main Menu, press Enter and then Esc.

The Esc key allows you to exit SBCON one screen at a time until you reach the Main Menu. If you want to exit SBCON, press Esc again and answer the confirmation prompt.

From a Workstation

- 1** At the Quick Access window, click Reports
- 2** Double-click the context.
- 3** Double-click Queues.
- 4** Double-click the preferred queue.
- 5** Double-click Servers.
- 6** Right-click the preferred server.
- 7** Select Session Report.

The server displays where the report is located.

Viewing an Error File

An error file records any errors that occurred during the backup or restore session. (See “[Session Files](#)” on page 26.) It can be viewed from the server or a workstation.

Ensure that you have met the prerequisites detailed in “[Prerequisites for SBCON Tasks](#)” on page 57.

From the Server

- 1** From the SBCON Main Menu, select Log File Administration > View an Error File.
- 2** If necessary, set the location of the log and error files for a backup session. (This option is not supported for a restore session.)

To set the location, use one or more of the following methods:

- ◆ Press Enter to accept the default location, and then go to [Step 3 on page 67](#).
- ◆ Press Insert to choose from a list of network directories, choose a directory, press Esc to exit and then go to [Step 3 on page 67](#).
- ◆ To change the directory, delete the existing directory path and type a new one, and then go to [Step 3 on page 66](#).

You can also press Insert during this process to select existing parts of the path from a list.

Restore errors are appended to a session's backup error file. A `Restore Session Begins` message marks the beginning of the restore errors, if any exist.

If you linked UNIX files, symbolic and hard links created under the NFS name space type might not be restored if the pathname is not recognized. When this happens, the error file contains messages similar to the following:

```
Error restoring namespace specific information of XXXX:tmp/hosts in NFS namespace, error 0x7!
```

```
Unable to allocate directory handle for XXXX:tmp/test!
```

- 3** When you are finished viewing the file, press Esc.
- 4** To return to the Main Menu, press Esc as required.

The Esc key allows you to exit SBCON one screen at a time until you reach the Main Menu. If you want to exit SBCON, press Esc again and respond at the confirmation prompt.

From a Workstation

- 1 At the Quick Access window, click Reports.
- 2 Double-click the context.
- 3 Double-click Queues.
- 4 Double-click the preferred queue.
- 5 Double-click Servers.
- 6 Right-click the preferred server.
- 7 Select Error Report.

Creating Log and Error Files

If you do not know what is on the media, or you have lost or deleted the log and error files, you can re-create them.

Ensure that you have met the prerequisites detailed in [“Prerequisites for SBCON Tasks” on page 57](#).

To create a log or error file, see [“Creating Session Files” on page 69](#).

Setting the Location of Backup Log and Error Files

Every time you do a backup, you are required to set a location (or directory path) for the log and error files. The restore log and error files are always located in the directory, SYS:SYSTEM/TSA/RESTORE.

You can set the location when requested in the backup procedures, or you can use the following method prior to beginning the backup procedures.

Ensure that you have met the prerequisites detailed in [“Prerequisites for SBCON Tasks” on page 57](#).

- 1 From the SBCON Main Menu, select Log File Administration > Set Location of Log and Error Files.
- 2 Set the location using one or more of the following methods:
 - ◆ Press Enter to accept the default location, and then press Esc to return to the Main Menu.
 - ◆ Press Insert to select from a list of network directories, select a directory, and press Esc twice to return to the Main Menu.
 - ◆ Press Backspace over the path shown (or a portion of the path), type a new directory or path, press Enter, and then press Esc to return to the Main Menu.
You can also press Insert during this process to select existing parts of the path from a list.
- 3 Press Esc when finished to return to the Main Menu.

Deleting the Log Files

SBCON allows you to overwrite sessions and reuse storage media.

When you delete the log file for a session using the following procedure, you automatically delete the corresponding error file as well.

To prevent outdated session information from taking up disk space and to help prevent confusion when you select sessions to view or restore, delete the session log files for any session that has been overwritten.

Ensure that you have met the prerequisites detailed in “Prerequisites for SBCON Tasks” on page 57.

NOTE: You cannot delete log and error files at the workstation.

To delete the log and error files from the server, do the following.

- 1** From the SBCON Main Menu, select Log File Administration > View a Log File.
- 2** If you have not set the location of log and error files, do so now using one or more of the following methods:
 - ♦ Press Enter to accept the default location, and then go to [Step 3 on page 69](#).
 - ♦ Press Insert to select from a list of directories, select a directory, press Esc to return to the previous window, and then go to [Step 3 on page 69](#).
 - ♦ Press Backspace over the path shown (or a portion of the path), type a new directory or path, and then go to [Step 3 on page 69](#).

You can also press Insert during this process to select existing parts of the path from a list.

- 3** Select the description of the session that you have overwritten (or plan to overwrite) and press Delete.

A confirmation prompt is displayed.

- 4** To delete the log and error file for that session, click Yes.
- 5** Press Esc when finished to return to the Main Menu.

Creating Session Files

If you do not know what is on the media, or you have lost or deleted the log and error files, you can re-create them from either the server or a workstation.

Ensure that you have met the prerequisites detailed in “Prerequisites for SBCON Tasks” on page 57.

From the Server

- 1** From the SBCON Main Menu, select Job Administration > Create Session File.

The Session Files Options screen appears.

- 2** Enter a descriptive name.
- 3** Select the media you will be using.

If you have only one loaded, press Enter to accept the default.

- 4** To verify the media, press Enter.

SBCON reads the media and shows you the selected session.

- 5** Confirm the creation of the session files.

If the session selected is the session you want for log and error file creation, click Yes, Create Files for This Session.

If the session selected is not the session you want for log and error file creation, click No, Go to the Next Session.

If you elect to create session files, they will be created in the location you chose.

From a Workstation

- 1** At the Quick Access window, click Create Session.
- 2** Double-click the context.
- 3** Double-click Queues.
- 4** Select the preferred queue.
- 5** Double-click Servers.
- 6** Select the preferred server.
- 7** Double-click Devices.
- 8** Select the preferred device.
- 9** Select the preferred media.
- 10** Right-click the job and select Submit the Job.
- 11** Check the appropriate fields in the Submission of Job dialog box.

Verifying Backup Data

This option allows you to check the data on the media from either the server or a workstation to make sure it has been backed up properly, and to check later that the data is valid and can be restored.

You can verify the backup data in two ways:

- ◆ If you know what sessions are on the media and you have the log and error files, you can select the session from the list.
- ◆ If you do not know what sessions are on the media, and if you are not sure the log and error files still exist, you should create the session files.

Ensure that you have met the prerequisites detailed in [“Prerequisites for SBCON Tasks” on page 57](#).

From the Server

- 1** Select Main Menu > Job Administration > Verify

The Verify Options screen appears.

- 2** Enter a descriptive name for the verify job.
- 3** Select the device and media of the job.

You can also manually enter this in the form *device name.media name*.

- 4** Press Enter in the Session to Restore field.

A list of all backup sessions is displayed in the Select a Session screen.

- 5** Select a session from the list.
- 6** Submit the job after confirming the Submit Job? prompt.
- 7** Select Job Administration > Current Job List.

- 8** Select the job you have submitted.

If there were no problems with the verification or the data, the Run Time Status screen displays the following message:

The verification process was completed normally.

From a Workstation

- 1** At the Quick Access window, click Verify.
- 2** Double-click the context if necessary.
- 3** Double-click Queues.
- 4** Double-click the preferred queue.
- 5** Double-click Servers.
- 6** Double-click the preferred server.
- 7** Double-click Devices.
- 8** Double-click the preferred device.
- 9** Right-click the preferred media.
- 10** Select Submit the Job.
- 11** In the Submission of Job dialog box, click the session.
- 12** Click Finish.

4

Optimizing

This chapter discusses the following topics:

- ◆ Whether you should compress your files before backup. See “Compressed Files” on page 73.
- ◆ How to increase speed when backing up your files. See “Host Server Performance” on page 74.
- ◆ How to enhance or optimize server performance. See “Setting Minimum Packet Receive Buffers” on page 74.
- ◆ How to synchronize memory between your media and your server. See “Setting Reserved Buffers Below 16 MB” on page 74.

Compressed Files

WARNING: SBCON and file compression should not be run simultaneously. Because the default time for both delayed SBCON sessions and compression is midnight of the current day, change one of these defaults to another time immediately.

If you want to perform a delayed backup that includes files flagged for compression, schedule the delayed backup after the compression time to allow time for the compression to be completed.

When a NetWare® 6 volume is mounted, file compression is set to On by default. When you perform a backup, you can specify whether to keep already compressed files in a compressed state for the backup, or back them up in a decompressed state.

To help you make a decision about whether to compress or decompress your backed up data, follow these guidelines:

- ◆ Backups are faster if files already compressed are left compressed.
- ◆ Compression is not supported in some environments (such as a NetWare 3.11 server, a DOS workstation, or Novell Storage Services™2.0). If you intend to restore a file that is currently compressed to an environment that does not support compression, back it up in a decompressed state.
- ◆ The SBCON utility has no compression feature, so it cannot compress a file that is currently decompressed.
- ◆ If volume compression is turned on and you back up compressed files in a decompressed state, restore speed is degraded if you overwrite existing files.

To improve restore speed, delete the files you no longer want from the hard disk before restoring them from the backup media. If you back up compressed files in a compressed state, restore speed is not affected.

- ◆ You might run out of disk space if you restore decompressed files to a volume that compresses files.

Host Server Performance

The speed of the SBCON utility depends upon the configuration of the host server and whether the server is backing up its own data or that of another server or workstation.

For information on server configuration and setting parameters, see Communications Parameters and File Caching Parameters under **SET** in NetWare Utilities Reference.

A server backing up its own data runs about twice as fast as a server backing up data from another server.

This section includes four options using LOAD and SET commands that can enhance your host server's performance.

Setting Minimum Packet Receive Buffers

Packet receive buffers represent the space in the server memory dedicated to handling network traffic. If the buffers are set too low, the server performance may be degraded.

To reset the packet receive buffers at the server console prompt, include the following SET command in the STARTUP.NCF file:

```
SET MINIMUM PACKET RECEIVE BUFFERS = X
```

The *x* represents the number of buffers. Set two buffers for each user on the system. The minimum is 10. The changes will take effect immediately.

Setting Reserved Buffers Below 16 MB

If your storage device driver requires memory below 16 MB and the server has more than 16 MB available to it, you must reserve memory below 16 MB for the driver.

To change the reserve memory at the server console prompt, include the following command in the STARTUP.NCF file:

```
SET RESERVED BUFFERS BELOW 16 MEG=X
```

Replace the *x* with a number between 8 and 300. The default is 200.

Setting Maximum Subdirectory Level and Minimum Cache Buffers

If you change the maximum subdirectory level in the server's STARTUP.NCF file, you must also change the minimum **cache buffer**. The default maximum subdirectory level is 25.

To change the cache buffers at the server console prompt, include the following commands in the AUTOEXEC.BAT file:

```
SET MINIMUM FILE CACHE BUFFERS=X  
SET MAXIMUM DIRECTORY CACHE BUFFERS=X
```

This command changes the minimum and maximum directory in MONITOR > Available Options > Server Parameters. The changes will take effect immediately.

5

Troubleshooting

This section provides information about the various troubleshooting scenarios.

Common Backup Problems

This section provides solution for some of the back up issues.

The media owner is unidentified

Possible Cause: Media is blank or was written in a non-SIDF format.

Action: See [“Checking the Media Status in a Device” on page 64.](#)

Backup speed is slow

Possible Cause: Compressed files are being backed up in a decompressed format.

Action: See [“Compressed Files” on page 73.](#)

Restore speed is slow

Possible Cause: File compression and SBICON are running at the same time.

Compressed files are being overwritten with decompressed files.

Action: See [“Host Server Performance” on page 74.](#)

The target you want is not listed

Possible Cause: The Target Service Agent isn't loaded on the desired target.

Action: See [“Backing Up Data” on page 34.](#)

Possible Cause: Insufficient time to list the targets.

Action: Allow more time for the software to find the target.

Possible Cause: The host and target are physically far apart and network traffic is heavy.

Action: Wait to perform a backup during nonpeak hours or when fewer users are logged in.

Data is corrupted, but no error message was given

Possible Cause: Compressed files were restored to a volume without compression.

Action: See [Step 10 on page 45.](#)

You (the network administrator) can't back up a workstation

Possible Cause: The /Password option is set and is denying access to data on the workstation.

Action: See [“Setting Rights to Back Up Portions of the eDirectory Tree” on page 18.](#)

Possible Cause: Backup was attempted from two different servers at the same time. One succeeded and the other one failed.

Action: Check to make sure all necessary backup software is installed loaded.

A backup does not contain all changes

Possible Cause: Differential and incremental backups were combined.

Action: Use one or the other of these types in conjunction with full backups. See [“Backup Types” on page 11.](#)

Changed files are not recognized

Possible Cause: The modify bit was cleared after the last customized backup, so changed files are not recognized.

Action: See [“Backup Types” on page 11.](#)

Unable to find the session log file you want

Possible Cause: It might be in a different directory or you might have deleted it accidentally.

Action: See [“Deleting the Log Files” on page 68.](#)

Pressing Insert on a device list does not produce a media list at the server

Possible Cause: No medium has been selected.

Action: Use the Select option at the bottom of the screen.

The Time Form Is Invalid message displays when you try to schedule a delayed backup

Possible Cause: The date has been entered incorrectly in the time form.

Action: Check that the scheduling has been entered correctly.

Files were restored but the error file contains a message specifying which name space formats were not restored

Possible Cause: The file attributes and name space formats are not configured on the volume you restored to.

Action: Check to make sure you selected the correct name space.

Can't list devices or see the Session/Error Report from the client and the message Unable to Establish Communication with Server appears

Possible Cause: A client-server communication problem.

Action: Check whether SBSC is loaded on the server. If not, load it.
If SBSC is already loaded, unload and reload it again.