

Novell NetWare® 6.5

www.novell.com

February 28, 2005

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.

You may not use, export, or re-export this product in violation of any applicable laws or regulations including, without limitation, U.S. export regulations or the laws of the country in which you reside.

Copyright © 1999-2005 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.

Novell, Inc. has intellectual property rights relating to technology embodied in the product that is described in this document. In particular, and without limitation, these intellectual property rights may include one or more of the U.S. patents listed at <http://www.novell.com/company/legal/patents/> and one or more additional patents or pending patent applications in the U.S. and in other countries.

Novell, Inc.
404 Wyman Street, Suite 500
Waltham, MA 02451
U.S.A.

www.novell.com

NetWare 6.5 Storage Management Services Administration Guide

March 31, 2005

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 Materials

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

Contents

- About This Guide** **7**

- 1 What Is SMS** **9**
 - What's New 9
 - SMS Components 9
 - Storage Management Data Requester 10
 - Target Service Agents (TSAs) 11
 - Target Services and Their Corresponding Target Service Agents 11
 - Target Service Agent for NetWare 6.5 (TSAFS) 11
 - SBCON 12

- 2 Installing and Configuring SMS** **13**
 - Memory Requirements 13
 - Installing Backup Server Pattern 13
 - Prerequisites 13
 - Backing Up Open Files 14
 - Before Running the Backup Software. 14
 - Configuring SMDR 14
 - Using iManager 15
 - Using Server Console. 16
 - Configuring Target Service Agent for NetWare File System 16
 - Using iManager 16
 - Using Command Line Switches. 17

- 3 Using SMS** **19**
 - Backing Up Data Using SMS 19
 - Backing Up eDirectory 19
 - Backing Up the NetWare File System. 23
 - Backing Up Clusters 24
 - Restoring Data Using SMS 24
 - How SMS Restores Data 25
 - Restoring eDirectory and the File System 26
 - Restoring Cluster-Enabled resources. 28

- 4 Optimizing SMS** **29**
 - Prerequisites 29
 - Setting up the Right Hardware and Software 29
 - Executing Background Tasks. 30
 - Configuring SMS for Better Performance. 30
 - How to Test and Configure 31
 - Basic Configuration 31
 - Advanced Configuration 33
 - Guidelines to Back Up Compressed Files 33
 - Enhancing the Host Server Performance. 34
 - Setting Minimum Packet Receive Buffers. 34
 - Setting Reserved Buffers Below 16 MB. 34

Setting Maximum Subdirectory Level and Minimum Cache Buffers	35
5 Troubleshooting SMS	37
Cluster Support for Volume Backups	37
Common Backup and Restore Issues	37
Directory and User Quota	39
Login Errors Because of License Failures	40
A TSA Features	41
TSA Options	41
Backup Options	41
Restore Options	44
Backup Types	44

About This Guide

This guide describes how to use NetWare® Storage Management Services™ (SMS). The guide is intended for network administrators and is divided into the following sections:

- ♦ Chapter 1, “What Is SMS,” on page 9
- ♦ Chapter 2, “Installing and Configuring SMS,” on page 13
- ♦ Chapter 4, “Optimizing SMS,” on page 29
- ♦ Chapter 5, “Troubleshooting SMS,” on page 37

Additional Documentation

For documentation on the SBICON utility, the basic storage management engine for NetWare see the *SBICON Administration Guide*.

Documentation Conventions

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

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

User Comments

We want to hear your comments and suggestions about this manual and the other documentation included with this product. Please use the User Comment feature at the bottom of each page of the online documentation, or go to www.novell.com/documentation/feedback.html and enter your comments there.

1

What Is SMS

NetWare® Storage Management Services™ (SMS) is a collection of software programs that provides backup, restore, and data migration services.

SMS allows you to backup targets such as the file system, Novell® eDirectory™, and the GroupWise® on NetWare, to a removable tape media for off-site storage. SMS is cluster-enabled and supports failover or failback of cluster-enabled resources. The backup engines use this infrastructure to provide a complete backup solution.

This section discusses the following:

- ◆ [“What’s New” on page 9](#)
- ◆ [“SMS Components” on page 9](#)
- ◆ [“Storage Management Data Requester” on page 10](#)
- ◆ [“Target Service Agents \(TSAs\)” on page 11](#)
- ◆ [“SBCON” on page 12](#)

What’s New

The [Target Service Agent for NetWare 6.5 \(TSAFS\)](#) is redesigned to scale with optimized hardware configuration for better performance. Some of the important aspects of TSAFS are

- ◆ Multi-processor enabled and uses a predictive data caching mechanism to improve performance
- ◆ Compatible with all existing backup engines
- ◆ Compatible with the data format used by existing versions of the TSA for NetWare
- ◆ Provides tunable parameters, which can be used to fine-tune performance to the specific environment. They can be configured either using the NetWare server console or iManager 2.0

These changes have optimized SMS technology and the result is a fresh, scalable, and efficient solution for data storage and recovery.

NOTE: The existing backup engines are not impacted by these changes.

SMS Components

SMS provides the basic layer on which other services are built. It achieves this by having the following independent components that provide the following functional abstractions:

- ◆ **Storage Management Data Requester (SMDR)** provides remote connectivity and abstracts the details of any communication between the backup and the target servers.

See [“Storage Management Data Requester” on page 10](#) for more information.

- ◆ **Target Service Agents (TSAs)** provide transparency by abstracting details of the specific target being backed up. The TSA for the NetWare file system (TSAF5) is used by various engines to back up file system, and the TSANDS is used to backup eDirectory.

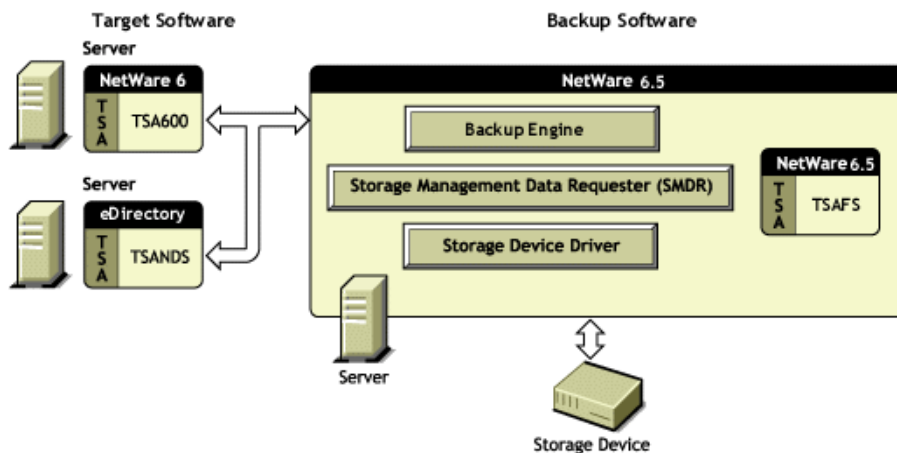
See “**Target Service Agents (TSAs)**” on page 11 for more information.

These components can be configured using either the NetWare server console or iManager 2.0.

All the third-party backup engines use these components to perform backup and restore tasks.

The backup process typically involves two machines:

- ◆ The backup server is typically the NetWare server running the backup program.
- ◆ 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 TSA must be loaded
- ◆ The backup server can also act as the target server.



The backup engine uses the SMS architecture on the host server to communicate and read data on the target server, and sends it to a storage medium such as a tape drive.

Storage Management Data Requester

The Storage Management Data Requester (SMDR) is the communication module in the SMS architecture. It provides transparent access to local and remote SMS services. The SMDR APIs are used by SBCON and other third-party applications as well.

The features of the SMDR are:

- ◆ **Supported Protocols**
- ◆ **Discovery and name resolution using SAP**
- ◆ **Discovery and name resolution using SLP**
- ◆ **Name resolution using hosts file**

Target Service Agents (TSAs)

A target is any service on the network that requires backup. Examples of targets include file system, eDirectory, GroupWise on NetWare.

Through specific TSAs, the backup engine allows you to back up the information that exists on the specific targets.

A Target Service Agent (TSA) understands the target and knows how to scan, read, and write the target data. The primary functions of a TSA are to prepare the target data for backup or restore, and to communicate with the backup engine. For example, TSAFS understands name spaces, file and directory attributes, security privileges, etc., for the files on target server.

The TSA encodes data in a generic format (SIDF). This allows one engine to interact with different TSAs.

Target Services and Their Corresponding Target Service Agents

Target Service	Corresponding Target Service Agent
NSS volumes and cluster-enabled resources	TSAFS
Traditional NetWare file system	TSAFS
eDirectory	TSANDS
GroupWise®	GW TSA

Target Service Agent for NetWare 6.5 (TSAFS)

The target service agent for NetWare 6.5 has been redesigned for better performance. Some of the salient features are:

- ◆ Compatible with all existing backup engines
- ◆ Compatible with the data format used by existing versions of TSA for NetWare
- ◆ Multiprocessor and cluster enabled
- ◆ Provides tunable parameters, which can be used to fine-tune performance to the specific environment
- ◆ Uses a data caching mechanism to improve performance

TSAFS Limitations


The limitations of TSAFS and other modules for the NetWare 6.5 release are documented in the *NetWare 6.5 Readme* (<http://www.novell.com/documentation/lg/nw65/readme/data/abvv0iw.html>).

The limitations for the Support Packs are documented in the readmes available with the Support Packs.

SBCON

SBCON is the native backup engine for NetWare. SBCON has three modules:

- ◆ User interface - Creates a job and submits to the eDirectory Queue.
- ◆ Q Manager (QMAN) - When a job is ready, QMAN takes the job from the eDirectory Queue. Q Manager facilitates multiple job scheduling.
- ◆ Engine - Executes and monitors the job.



- 1 User Interface creates a job and submits to the NDS queue
- 2 Q Manager takes the job from the NDS Queue which is ready for execution
- 3 The engine

For more information on SBCON, refer to the *SBCON Administration Guide*.

2 Installing and Configuring SMS

NetWare® Storage Management Services™ (SMS) is installed during the NetWare 6.5 installation.

This section provides information about the system requirements, prerequisites, and the procedure for running the backup software.

The following topics are discussed:

- ♦ “Memory Requirements” on page 13
- ♦ “Prerequisites” on page 13
- ♦ “Configuring SMDR” on page 14
- ♦ “Configuring Target Service Agent for NetWare File System” on page 16

Memory Requirements

The memory requirements are the same as that of NetWare. For more information, refer to *NetWare installation quickstart* (<http://www.novell.com/documentation/lg/nw65/pdfdoc/nw6install/nw6install.pdf>).

Installing Backup Server Pattern

Pattern deployment is a new concept introduced for installing services. Generally, when you install a certain service, you need a group of components or products. For a service to function properly, all the dependent products must be installed. Pattern deployment provides patterns for different services. Selecting a pattern automatically selects and installs its dependencies.

SMS is installed as a part of the NetWare 6.5 installation.

To install the server as a backup server, on the Choose a Pattern page, select the Backup Server pattern.

The patterns are listed along with the customized NetWare Server or the basic NetWare File Server.

Prerequisites

Meet the following prerequisites, before starting the backup software.

Backing Up Open Files

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

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

The Supervisor right is required to back up open files.

To enable CopyOnWrite on a single NSS volume, at the server console, enter **nss / FileCopyOnWrite=volume_name**.

Before Running the Backup Software

Before starting the backup process, you need to perform the following tasks:

- Load the controller and storage device drivers on the backup server.
- Load the SMDR and TSAs on the target server (see [“Loading the Target Service Agents” on page 14](#)).
- Load the SMDR and the TSAs on the backup server.

At the server console, enter **smsstart**, to load SMDR, TSAFS, and TSANDS.

By default, SMDR, TSAFS and TSANDS are loaded from the autoexec.ncf file.

Loading the Target Service Agents

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

To Back Up or Restore	At This Console	Enter This Command
NSS volumes and cluster-enabled resources	Target server	load TSAFS
NetWare Traditional file system	Target server	load TSAFS
eDirectory	NetWare 4 or later server (preferably a server with a replica of the largest partition)	load TSANDS
GroupWise	Target server	load GW TSA

Configuring SMDR

You can configure SMDR for SLP and SAP discovery, and Name resolution through HOSTS file, using iManager 2.0 or the server console.

This section provides information about:

- ◆ [“Using iManager” on page 15](#)
- ◆ [“Using Server Console” on page 16](#)

Using iManager

Complete the following steps to configure the discovery and name mechanisms used by the SMDR.

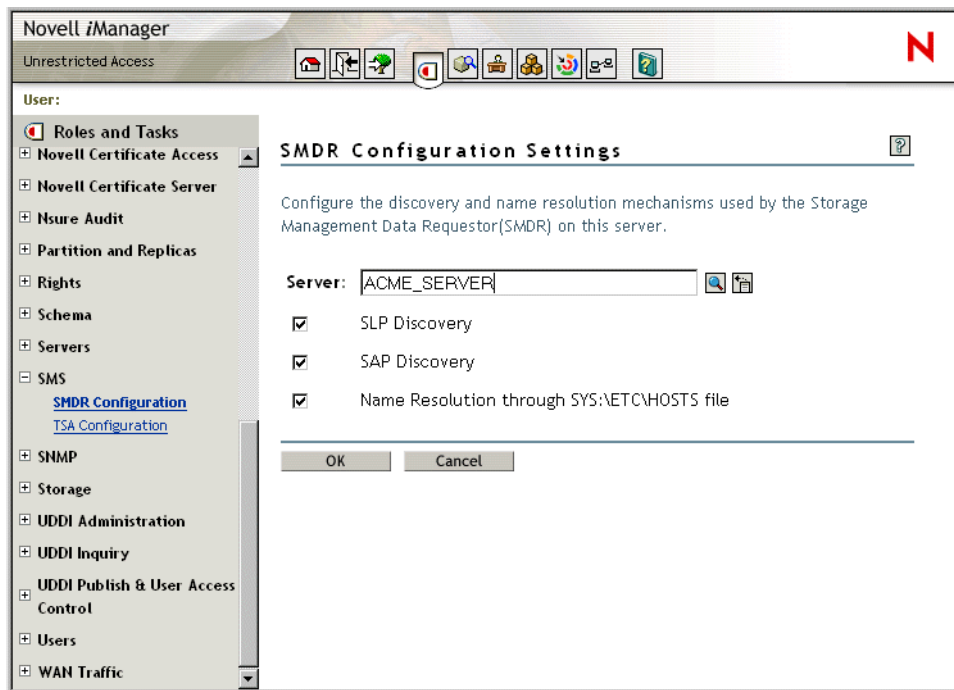
- 1 In the iManager screen click SMS and click SMDR Configuration.
- 2 Select the server for which you need to modify the SMDR configuration settings, using the eDirectory object selector.

The following options are displayed. Select the required options and click OK.

- ◆ **SLP Discovery:** SMDR can be configured to use SLP for discovery and name resolution. This enables SMDRs to locate other SMDRs running on other servers in the network. 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.

If cluster-enabled resources are to be backed up or restored, SLP should be used as the discovery mechanism.

- ◆ **SAP Discovery:** SMDR can also 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, Service Location Protocol (SLP) replaces SAP.
- ◆ **Name Resolution Through HOSTS File:** SMDR can also be configured to use sys:\etc\hosts file for IP address name resolution. The HOSTS file is automatically installed in the sys:\etc directory when you install TCP/IP. If entries are added to this file SMDR will use these entries to resolve the IP address.



Supported Protocols: SMDR supports Internet protocol (IP) and Sequenced Packet Exchange™ (SPX™) or Internet Packet Exchange (IPX™) protocols.

Using Server Console

To configure the SMDR using the server console, enter the following command.

```
load smdr new
```

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

The configuration file, smdr.cfg file is a text file located in the sys:\etc\sms directory on your server. The protocols are specified in this file.

Configuring Target Service Agent for NetWare File System

You can modify the tunable parameters of the target service agent using iManager 2.0 or the command line switches.

This section provides information about:

- ◆ [“Using iManager” on page 16](#)
- ◆ [“Using Command Line Switches” on page 17](#)

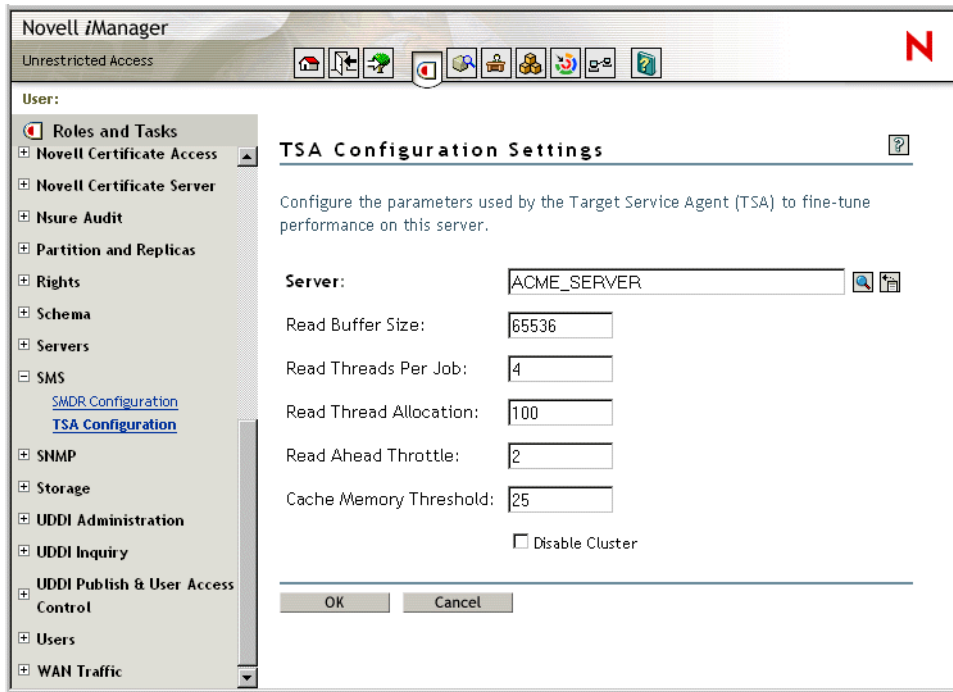
Using iManager

Complete the following steps to configure the tunable parameters used by the TSA.

- 1 In the iManager screen click SMS and click TSA Configuration.
- 2 Select the server for which you need to modify the TSA configuration settings, using the eDirectory object selector.
- 3 Enter the following values and click OK.
 - ◆ **ReadBufferSize:** The number of data bytes read from the file system by a single read operation.

This parameter is based on the buffer size requested by the engine. For example, if the engine requests 64 KB of data for each read operation, set the buffer size to 64 KB to allow the TSAFS to service the engine better. By default this is set to 65536 Bytes.
 - ◆ **Read Threads Per Job:** The number of read-ahead threads for a job. This enables the TSAFS to read data ahead of the engine request during backup. This switch is based on the number of processors in the system. The default value is 4 for a single or dual processor system. Set the read threads to a higher value if the system has more processors.
 - ◆ **Cache Memory Threshold:** The percentage of unallocated server memory that the TSA can utilize to store cached data sets. This represents a maximum percentage value of unallocated server memory that the TSA uses to store cached data sets. Default value is 25% of unallocated server memory. The cache memory utilized by the TSA is dynamically reset based on the available unallocated memory.
 - ◆ **Disable Cluster:** Specific to clusters. If the backup server does not support clusters, this option will be disabled. Determines cluster support. Uncheck if the TSA is running on a cluster node and the backup engine is cluster enabled. Check if the TSA is running on a non-cluster node or the backup engine is not cluster enabled. This is unchecked by default.

Running a cluster-enabled TSA on a non-cluster node does not affect functionality in any way.



Using Command Line Switches

The following table provides information about the basic parameters and their usage.

Task	Purpose	Command
Set the number of read threads for a job	This enables the TSA to read data ahead of the engine request during backup. This switch is based on the number of processors in the system. The default value is 4. Set the read threads to a higher value if you have more processors.	TSAFS / readthreadperjob=value
Set the read buffer size	This is the amount of data (Bytes) read from the file system by a single read operation. This switch is based on the buffer sizes used by the engine. For example, if the engine requests 32 KB of data for each read operation, set the buffer size to 32 KB to allow the TSA to service the engine better. This value is set in multiples of 512 Bytes (sector size or the actual block size). Another aspect to consider while setting the buffer is the mean size of the dataset being backed up. For example, if the mean size of the dataset is 28 KB, set the buffer size to 64 KB so additional memory is added to the mean size of the dataset. This is required for backup of file characteristics and SIDF encoding. The default value is 64 KB.	TSAFS / readbuffersize=value

Task	Purpose	Command
Set the percentage of total server memory to store cached data sets.	<p>This is used to specify the percentage of total server memory that the TSA can utilize to store cached data sets. This represents a maximum percentage value of total server memory that the TSA uses to store cached data sets.</p> <p>The default value is 25% of the total server memory.</p>	TSAFS /Cache Memory Threshold
Set the TSA to support failover or failback on cluster volumes.	<p>Set this to cluster if the TSA is running on a clusternode and the backup engine supports clusters.</p> <p>Set this to nocluster if the TSA is running on an on cluster node or the backup engine does not support clusters.</p> <p>The default value is cluster.</p> <p>Running a cluster-enabled TSA on a non-cluster-enabled node does not affect functionality in any way.</p>	TSAFS /cluster nocluster

For performing advanced configuration refer to [“Configuring SMS for Better Performance” on page 30](#).

3

Using SMS

This section provides information about how to use SMS to backup and restore Novell® eDirectory™, NetWare® file system, and Novell Storage Services (NSS™).

- ◆ [“Backing Up Data Using SMS” on page 19](#)
- ◆ [“Restoring Data Using SMS” on page 24](#)

Backing Up Data Using SMS

This section provides information on how SMS backups data from eDirectory and file system.

- ◆ [“Backing Up eDirectory” on page 19](#)
- ◆ [“Backing Up the NetWare File System” on page 23](#)
- ◆ [“Backing Up Clusters” on page 24](#)

Backing Up 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:

- ◆ [“Backing Up a Distributed Database” on page 20](#)
- ◆ [“Backing Up Object ID Numbers” on page 20](#)
- ◆ [“Backing Up Placeholder \(Unknown\) Objects” on page 20](#)
- ◆ [“Backing Up the Schema” on page 20](#)
- ◆ [“Server Interdependence” on page 21](#)
- ◆ [“Backup Software for eDirectory” on page 21](#)
- ◆ [“Setting Rights to Back Up Portions of the eDirectory Tree” on page 22](#)
- ◆ [“Backup Frequency” on page 22](#)

Backing Up a 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.

Backing Up 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 name 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.

Backing Up the Schema

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.

Backing Up Placeholder (Unknown) Objects

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

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.

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.

Backup Software for eDirectory

In order to back up the eDirectory database, the TSANDS 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 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 that ships with NetWare allows selective backup and restoration of an eDirectory tree structure.

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

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. For more information see, [Partial eDirectory Backups and Restores](#).

Partial eDirectory Backups and Restores

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

The TSANDS 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. In other cases, an object might be restored but not be functional unless you first restore its dependent objects.

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.

Backup Frequency

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.

Backing Up the NetWare File System

To back up file system data, TSAFS must be loaded on each target server for which a backup is to be created (see [“Loading the Target Service Agents” on page 14](#)). TSAFS supports backup of NetWare file system name spaces, extended attributes, trustee rights, compression etc.

After the device drivers for your backup hardware and the TSAFS is loaded, you can run the backup program of your choice. The backed up information is stored in the SIDF format.

Whenever you back up or restore files or directories, corresponding log and 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 engine can restore it. These session IDs are created from the .log files.

For more information about how to backup refer to the respective backup engine documentation.

This section discusses the following:

- ◆ [“Backing Up Trustee Assignments” on page 23](#)
- ◆ [“Backing Up Server-Specific Information” on page 23](#)

Backing Up 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 TSAFS. If a User object is deleted and then re-created or restored, its object ID changes. 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.

For additional information about object ID and trustee issues, see [“Restoring eDirectory and the File System” on page 26](#).

Backing Up Server-Specific Information

Server-specific information such as the replica information, ID information, name spaces loaded, and system configuration is stored on the sys volume. 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 the replica list and replica types on the backup server during backup.
- ◆ **startup.ncf** contains the disk driver, name spaces, and SET parameters.
- ◆ **autoexec.ncf** contains load modules and the server configuration.
- ◆ **vol\$info.txt** contains volumes on the server, name spaces loaded, compression, and migration information.

The server specific information does not need to be restored unless you have lost the sys: volume. For more information, see [“Restoring eDirectory and the File System” on page 26](#).

Backing Up Clusters

Novell Cluster Services™ is a server clustering system that ensures high availability and manageability of critical network resources including data (volumes), applications, and services. It is a multinode clustering product for NetWare that is enabled for eDirectory and supports failover, failback, and migration (load balancing) of individually managed cluster resources. For more information, see the [Novell Cluster Services documentation \(http://www.novell.com/documentation/lg/ncs65/index.html\)](http://www.novell.com/documentation/lg/ncs65/index.html).

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. SMS supports backup and restore of cluster-enabled resources. In addition, the backup session can be automatically recovered in case of a failover or failback of the target cluster-enabled resources, if the backup engine supports it.

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

Consider the following before preparing for backup and restore of cluster-enabled resources. These conditions are applicable only if the backup engine is cluster-enabled.

- Use TSAFS to backup or restore cluster-enabled resources. (See “[Target Service Agents \(TSAs\)](#)” on page 11 for details.)
- If a file system target service is chosen on a cluster server, only non-cluster volumes can be backed up.
- If cluster-enabled resources are to be backed up or restored, SLP should be used as the discovery mechanism.
- A cluster node will have one or more non-cluster volumes. When the cluster node is chosen, only these volumes will be listed. To backup the cluster volumes, choose the cluster resources corresponding to those cluster volumes. The resources are listed as virtual servers along with other servers. To treat all cluster volumes as non-clustered for backup, so that they will be listed under the cluster node instead of virtual server, use [TSAFS command line switches](#).

Restoring Data Using SMS

This section provides information on how SMS restores data from eDirectory and NetWare file system.

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

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

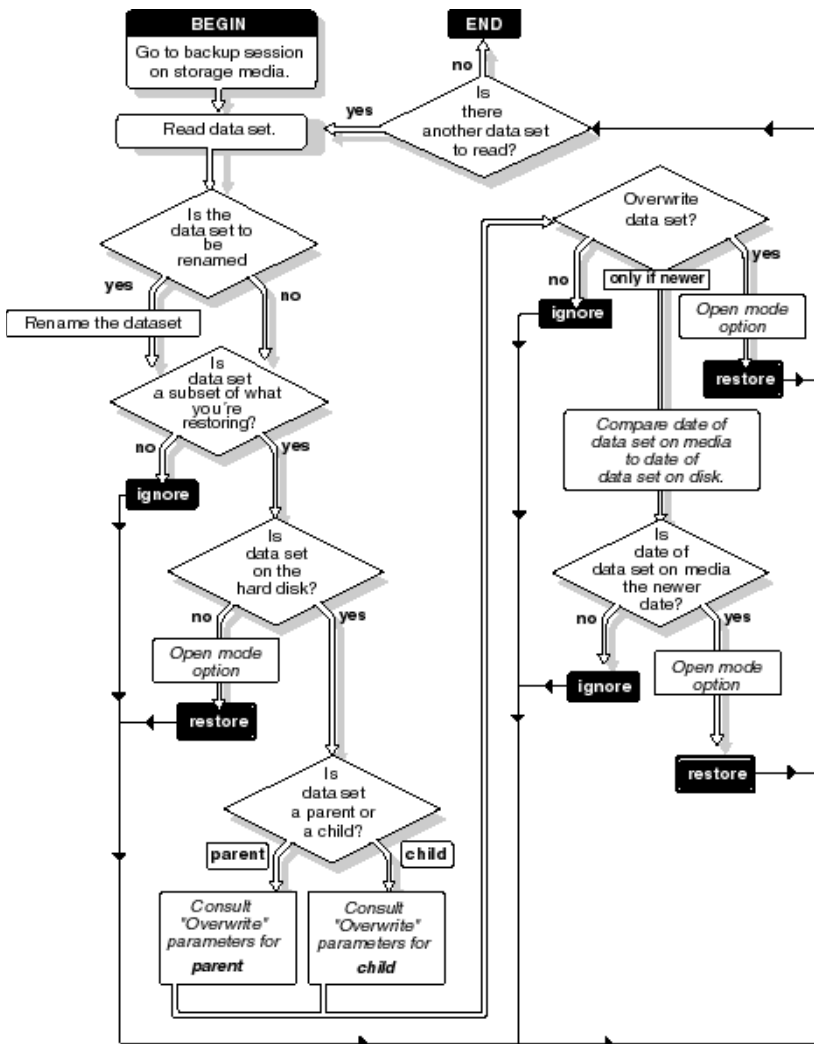
The following topics are discussed in this section:

- ◆ [“How SMS Restores Data” on page 25](#)
- ◆ [“Restoring eDirectory and the File System” on page 26](#)
- ◆ [“Restoring Cluster-Enabled resources” on page 28](#)

How SMS Restores Data

During a restore session, the backup engine 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?



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, exists when the eDirectory database information is restored, those partitions and replicas will also be restored, and you will not need to repartition the tree.

This section provides specific recovery procedures for the following scenarios:

- ◆ [“Corrupted Data” on page 26](#)
- ◆ [“Loss of a Volume Other Than Sys:” on page 26](#)
- ◆ [“Loss of the Sys: Volume or an Entire Server” on page 26](#)
- ◆ [“Loss of the Entire eDirectory Tree” on page 27](#)

Corrupted Data

For 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 deleted, 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 the [“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.

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 metadata such as trustee rights and file attributes.

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.](#)

Loss of the Only Server in a 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.

Because 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.
4. Restore the file system.

Loss of a Single Server in a 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 is damaged, causing the failure of the entire server.

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.
3. Reinstall NetWare, including restoring the servdata.nds file (located in sys:) to the original server.
4. Restore eDirectory.
5. Restore the file system.
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.

To restore the eDirectory tree structure:

- 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 previously in the eDirectory tree.

In NWCONFIG, select Directory Options > Directory Backup and Restore Options to obtain the information.

If you do not use the same names, 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 because of 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.

After 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.

Restoring Cluster-Enabled resources

Restoring of cluster-enabled resources is similar to a normal restore session. See “[Restoring cluster-enabled resources](#)” in the *SBCON Administration Guide*.

NOTE: Auto-recovery of the restore session on failover or failback will be performed based on the backup engine used.

4 Optimizing SMS

The need for faster backup solutions has grown as the data storage needs have increased and the backup window remains unchanged. This section provides a brief overview and specific information on enhancing the backup performance.

- ◆ [“Prerequisites” on page 29](#)
- ◆ [“Configuring SMS for Better Performance” on page 30](#)
- ◆ [“Guidelines to Back Up Compressed Files” on page 33](#)
- ◆ [“Enhancing the Host Server Performance” on page 34](#)

Prerequisites

The factors that affect the backup performance are:

- ◆ Disk subsystem and configuration
- ◆ Software
- ◆ Archive medium

Before you proceed with tuning the TSAFS, ensure to complete the following tasks:

- ◆ [“Setting up the Right Hardware and Software” on page 29](#)
- ◆ [“Executing Background Tasks” on page 30](#)

Setting up the Right Hardware and Software

Performance tuning depends on the hardware and the software used in that setup. Ensure to use the right hardware and software to improve the performance.

This section provides information about:

- ◆ [“Hardware Specifications” on page 29](#)
- ◆ [“Software Specifications” on page 30](#)

Hardware Specifications

- ◆ A single disk IDE based system will not provide optimal performance.
- ◆ A single LUN (even when striped across multiple disks) will not provide optimal performance.

To improve performance, the disk subsystem should be capable of aggregating or providing multiple paths to access the data. This would improve the possibility of doing overlapped I/O to disk and hence would improve the overall disk capability. This is why a single LUN or a single

IDE disk cannot perform at high speeds. It is required that if backup is critical and has to be a non-time consuming process, the disk subsystem should be capable of extracting high amounts of data.

Using multiple LUNs and storing each LUN across different disk chassis units help in aggregation as each of the disk chassis has the potential to be busy reading or writing data at the same time thereby improving total throughput from the disks.

Software Specifications

- ◆ Set the number of LUNs per pool to the number of processors up to a limit of around 4
- ◆ Use software RAID-0 on hardware RAID-5

Using RAIDs that are not backup friendly is a sure way of killing performance at the media manager level. To reduce disk failures, set the disks at a RAID 5 configuration. This may also not necessarily improve the situation when the disk start failing.

On the other hand, using software RAID 0 on Hardware RAID 5 provides better disk I/O from the media layer thereby improving a disk bound activity like backup go faster.

Enhancing the data fault tolerance and performance of the NSS volumes using RAID will also improve the SMS performance. For more information see [NSS documentation \(http://www.novell.com/documentation/lg/nw65/nss_enu/data/ajhvbk8.html#ajhvbk8\)](http://www.novell.com/documentation/lg/nw65/nss_enu/data/ajhvbk8.html#ajhvbk8).

SMS provides optimal performance with 4 LUNs.

Executing Background Tasks

The backup performance will improve when:

- ◆ Anti-virus software is disabled during backup.

Most of the anti-virus software programs have options to run checks on write or modify access to files. Setting these options at least during backups would help completing the process, as the backup software performs only open and read operations. There are some backup engines that do not perform well, when the anti-virus software is set to modify and open files in write mode.

Also, if a disk is affected by virus it should be detected before starting the backup. The backup process is just a read operation and hence will not induce a virus on the system. Based on this fact anti-virus software can also (if tunable) be set to ignore backup applications accessing the file system.

- ◆ Compression and backup are not run together.

Configuring SMS for Better Performance

This section provides information about how to test and configure your setup to optimize the backup performance.

- ◆ [“How to Test and Configure” on page 31](#)
- ◆ [“Basic Configuration” on page 31](#)
- ◆ [“Advanced Configuration” on page 33](#)

How to Test and Configure

You can use TSATEST to diagnose the performance related issues. TSATEST executes backup jobs and performs statistical analysis of the behavior of the backup APIs being used. It helps to demonstrate what the SMS components are capable of delivering.

For more information see [Novell support site \(http://support.novell.com/cgi-bin/search/searchtid.cgi?/2962002.htm\)](http://support.novell.com/cgi-bin/search/searchtid.cgi?/2962002.htm).

Using TSATEST in multiple combinations of TSAFS switches to check how tuning the following parameters can affect performance would provide a inductive way towards detecting optimal sweet spots for these switches on your systems.

You can also monitor the disk throughput using NetWare[®] Remote Manager. This would show the efficiency of the disk system to churn out data. This information can be compared with the SMS throughput, to isolate the issues. This can be due to other competing processes running on the system while backup is being run. As backup is an I/O bound problem, the above would help to check if SMS is able to deliver the same performance as the disk system.

Basic Configuration

Configure the following basic tunable parameters to enhance the SMS performance. For more information about configuring the switches, see “[Configuring Target Service Agent for NetWare File System](#)” on page 16.

Task	Purpose	Field Name in the iManager Interface	Command
Set the number of read threads for a job	<p>This enables the TSA to read data ahead of the engine request during backup. This switch is based on the number of processors in the system.</p> <p>The default value is 4. This value ranges from 1 - 32.</p> <p>Set the read threads to a higher value if you have more processors.</p> <p>Also, monitor the disk I/O performance and set the switch to higher values to check if the disk I/O improves.</p>	Read Threads Per Job	TSAFS / readthreadperjob=value

Task	Purpose	Field Name in the iManager Interface	Command
Set the read buffer size	<p>This is the number of data bytes read from the file system by a single read operation. This switch is based on the buffer sizes requested by the engine. For example, if the engine requests 64KB of data for each read operation, set the buffer size to 64 KB to allow the TSA to service the engine better.</p> <p>Another aspect to consider while setting the buffer is the mean size of the dataset being backed up. For example, if the mean size of the dataset is 28 KB, set the buffer size to 64 KB so additional memory is added to the mean size of the dataset. This is required for backup of file characteristics and SIDF encoding. The default value is 64768 bytes.</p>	Read Buffer Size	TSAFS / readbuffersize=value
Set the percentage of total server memory to store cached data sets.	<p>This is used to specify the percentage of total server memory that the TSA can utilize to store cached data sets. This represents a maximum percentage value of total server memory that the TSA uses to store cached data sets.</p> <p>The default value is 25% of the total server memory.</p> <p>Set it to a higher value to enable the TSA to cache more data sets and improve the backup performance of TSA.</p>	Cache Memory Threshold	TSAFS / cachememorythreshold=value

Advanced Configuration

Configure the following advanced tunable parameters to enhance the SMS performance. For more information about configuring the switches, see “[Configuring Target Service Agent for NetWare File System](#)” on page 16.

Task	Purpose	Field Name in the iManager Interface	Command
Set the percentage of read threads to process a data set.	<p>This sets the maximum number of read threads that process a data set at a given time. This determines the percentage of <i>readthreadsperjob</i> that should be allocated to a data set before proceeding to cache another data set.</p> <p>This enables the TSA to build a cache of data sets in a nonsequential manner. Engines reading data sets simultaneously have the advantage of improved performance if the TSA builds a nonsequential cache rather than a sequential cache.</p> <p>The default value is 100. This sets all read threads to completely process a data set before proceeding to another data set.</p> <p>Set this value lower than 100 if the backup engine reads multiple data sets from the TSA simultaneously.</p>	Read Thread Allocation	TSAFS / readthreadallocation =value
Set the maximum threshold for data sets that can be processed simultaneously	<p>This sets the maximum number of data sets that the TSA caches simultaneously. This prevents the TSA from caching parts of data sets and enables complete caching of data sets instead.</p> <p>Use this switch along with the <i>readthreadallocation</i> switch.</p> <p>Set this value to reflect the number of data sets that the backup engine processes simultaneously. The default value is 2.</p>	Read Ahead Throttle	TSAFS / readaheadthrottle=val ue

Guidelines to Back Up Compressed Files

WARNING: Backup and file compression operation should not be run simultaneously. For example, if the default time for both scheduled backup and restore sessions, and compression is midnight, set one of these defaults to different time.

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 you perform a backup, you can specify whether to keep compressed files in the same state or back them up in a decompressed state.

To decide about whether to compress or decompress your backed up data, follow these guidelines:

- ◆ Backups are faster if compressed files are backed up in the same state.
If volume compression is turned on and you back up compressed files in a decompressed state, restore speed is degraded when existing files are overwritten.
- ◆ Compression is not supported in some environments (such as a NetWare® 3.11 server 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.
- ◆ You might run out of disk space if you restore decompressed files to a volume, as the compression does not begin immediately.

Enhancing the Host Server Performance

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

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

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

- ◆ “Setting Minimum Packet Receive Buffers” on page 34
- ◆ “Setting Reserved Buffers Below 16 MB” on page 34
- ◆ “Setting Maximum Subdirectory Level and Minimum Cache Buffers” on page 35

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 might 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 SMS

This section provides troubleshooting information that you can use to rectify some of the issues that might arise during backup or restore operations.

The issues are categorized under the following sections:

- ♦ [“Cluster Support for Volume Backups” on page 37](#)
- ♦ [“Common Backup and Restore Issues” on page 37](#)
- ♦ [“Directory and User Quota” on page 39](#)
- ♦ [“Login Errors Because of License Failures” on page 40](#)

Cluster Support for Volume Backups

This section provides information about the following troubleshooting scenarios:

1. [“Clustered volumes are not backed up during full server backups” on page 37](#)

Clustered volumes are not backed up during full server backups

Possible Cause: The application used does not handle SMS clustered resources, causing clustered volumes to be skipped as the TSA was loaded cluster enabled.

Action: Use `TSAFS /NoCluster` command to disable the cluster support.

Common Backup and Restore Issues

This section provides information about the following troubleshooting scenarios:

1. [“Backup does not include the modifications” on page 38](#)
2. [“Backup is slow” on page 38](#)
3. [“Files were restored but the error file contains a message specifying the name space formats not restored” on page 38](#)
4. [“Modified files are not recognized” on page 38](#)
5. [“Out of Memory” on page 38](#)
6. [“Restore is slow” on page 38](#)
7. [“The target service agents you want is not listed” on page 38](#)
8. [“Unable to back up open files in NSS Volumes” on page 39](#)
9. [“Unable to restore compressed files” on page 39](#)
10. [“Unable to set the data set name” on page 39](#)

Backup does not include the modifications

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 44](#).

Backup is slow

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

Action: See [“Guidelines to Back Up Compressed Files” on page 33](#).

Files were restored but the error file contains a message specifying the name space formats 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 the name space is configured on the required volume.

Modified 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 44](#).

Out of Memory

Possible Cause: The server is running on lower memory than what NetWare 6.5 system requirements recommend

Possible Cause: A large number of applications are running when backups are scheduled.

Possible Cause: The server reports low memory condition while backups are ongoing.

Action: Set the [CacheMemoryThreshold](#) parameter to a lower value. By default this is set to 25% of the unallocated server memory.

Restore is slow

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

Possible Cause: Compressed files are being overwritten with decompressed files.

Action: See [“Enhancing the Host Server Performance” on page 34](#).

The target service agents you want is not listed

Possible Cause: SLP is not configured properly.

Action: Enter `display slp services smdr.novell` at the server console, to list the SMDRs with the target service's IP address. Check if the required target's SMDR is listed.

Possible Cause: The SMDR and the TSA are not loaded.

Action: Ensure that the SMDR and the TSA are loaded.

Unable to back up open files in NSS Volumes

Possible Cause: TSAFS supports backup of open files on Novell Storage Services (NSS) volumes if the CopyOnWrite feature is enabled. The Supervisor right is required to back up open files

Action: To enable CopyOnWrite on a single NSS volume, at the server console, enter **nss / FileCopyOnWrite=volume_name**.

To know about the various options provided by the backup applications refer to the respective documentation.

Unable to restore compressed files

Possible Cause: Compression is not supported in some environments such as a NetWare® 3.11 server or Novell® Storage Services™ 2.0.

Possible Cause: The volume that the data is restored to does not support compression.

Action: Restore the file to a volume that supports compression.

Unable to set the data set name

Possible Cause: This happens when the TSA is unable to restore a particular name in a name space. The data set is restored to completeness, but while restoring names in all name spaces there were name conflicts on the non primary name spaces. This causes a failure to set the names in those name spaces.

The following is an example of such an error:

```
Unable to set the data set name in MACnamespace for sys:/abc/def.txt. Restore will continue processing the data set name in other name spaces.
```

Action: This is a warning and will not impact the restore process.

Directory and User Quota

This section provides information about the following troubleshooting scenarios:

1. [“Out of disk space” on page 39](#)

Out of disk space

Possible Cause: There is not enough disk space on the volume to which the data is restored.

Action: Increase the volume size and restart the restore.

Possible Cause: A user space restriction is set for certain users and the application is attempting to restore data that exceeds this space restriction.

Action: Increase the user space restriction on the restored volume for required users.

Possible Cause: There are directory quotas set on specific directories that are being restored, and the application is trying to restore data that exceeds this quota specification.

Action: Increase the directory quota on the specific directories restart restore.

Possible Cause: You might run out of disk space if you restore decompressed files to a volume, as the compression does not begin immediately.

Action: Compress the files before restore.

Login Errors Because of License Failures

This section provides information about the following troubleshooting scenarios:

1. [“Unable to log in due to license failures” on page 40](#)

Unable to log in due to license failures

Possible Cause: SMS connections to the file system TSA occupy a licensed connection on the server for file system access. Having multiple connections on the server would cause authentication failures when the number of licenses used equals the number installed.

Action: Increase the number of licenses on the server.



TSA Features

This section provides information about the various options to modify the backup settings, and the types of backup.

TSA Options

The following TSA options are discussed in this section.

- ◆ [“Backup Options” on page 41](#)
- ◆ [“Restore Options” on page 44](#)

Backup Options

All backup types contain advanced options to allow you to customize your backup. These options allow you to perform the following tasks.

- ◆ [“Choosing Subsets of Data to Back Up” on page 41](#)

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.

- ◆ [“Scanning Data Sets” on page 43](#)

You can specify how to scan what you are backing up.

Choosing Subsets of Data to Back Up

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 Novell[®] 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.

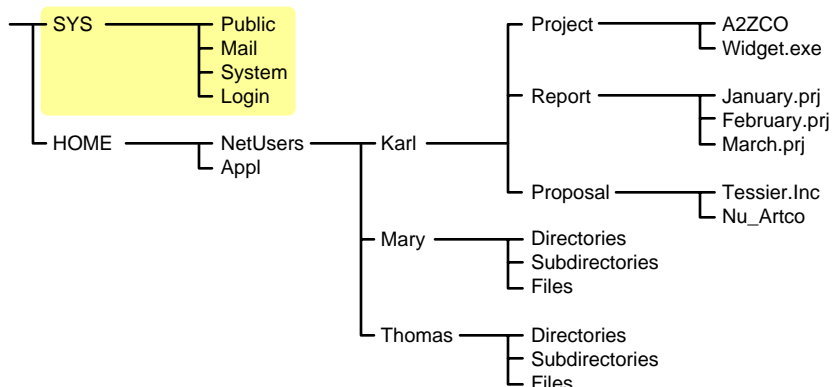
After you exclude part of the structure such as a volume, directory, or container, you cannot include any subdirectories, files, or objects beneath that excluded volume, directory, or container.

Include

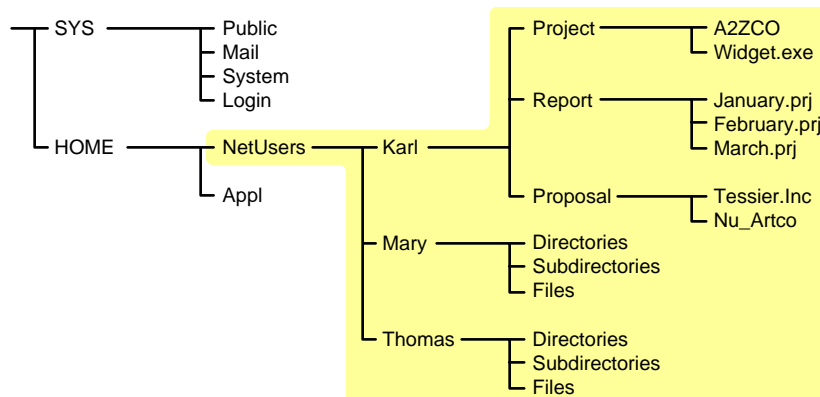
To back up a small part of the file system structure, use the Include option to specify the data you want. Everything you do not specifically include is excluded.

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 shown 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

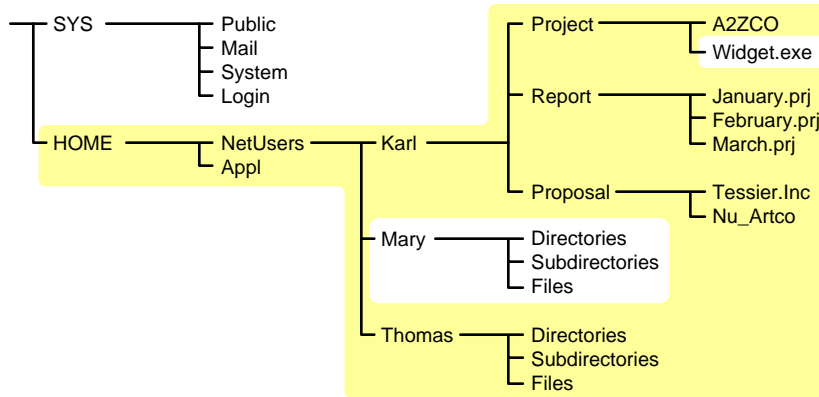
By combining the include and exclude options, you can control what is backed up.

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



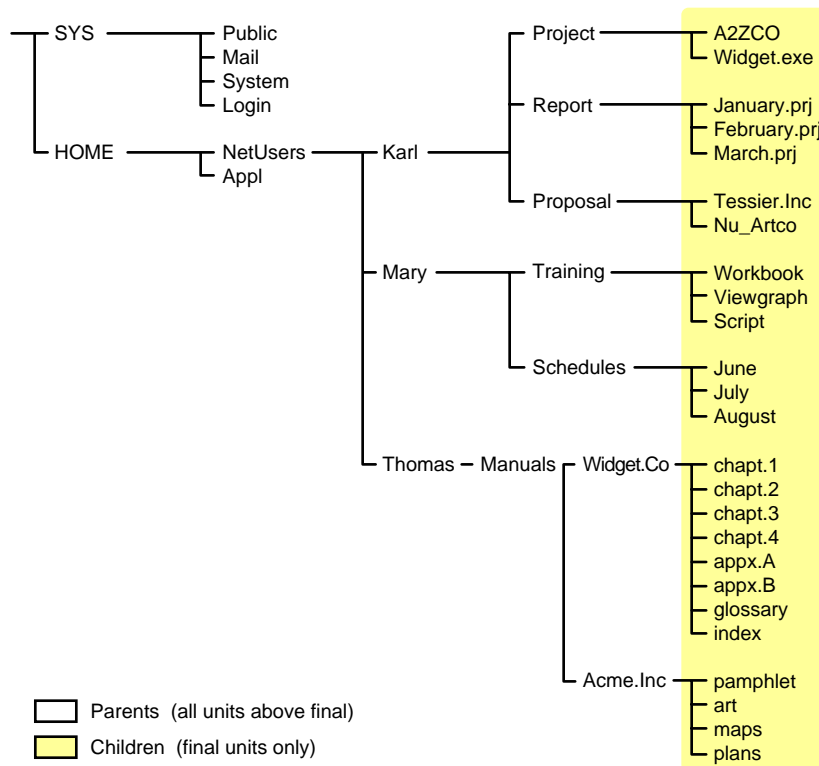
Scanning Data Sets

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.



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.

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 [“Backup Options” on page 41](#).

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 when 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.

Backup Types

SBCON 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).

