

Novell Open Enterprise Server

www.novell.com

December 23, 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.

Any products or technical information provided under this Agreement may be subject to U.S. export controls and the trade laws of other countries. You agree to comply with all export control regulations and to obtain any required licenses or classification to export, re-export, or import deliverables. You agree not to export or re-export to entities on the current U.S. export exclusion lists or to any embargoed or terrorist countries as specified in the U.S. export laws. You agree to not use deliverables for prohibited nuclear, missile, or chemical biological weaponry end uses. Please refer to www.novell.com/info/exports/ for more information on exporting Novell software. Novell assumes no responsibility for your failure to obtain any necessary export approvals.

Copyright © 1999-2003, 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

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 Core Protocol and NCP are trademarks of Novell, Inc.

NetWare Loadable Module and NLM are trademarks 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.

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

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

SPX/IPX is a trademark of Novell, Inc.

Storage Management Services is a trademark of Novell, Inc.

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

Third-Party Materials

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

About This Guide

This guide describes how to use Storage Management Services™ (SMS) for OES.

This guide is divided into the following sections:

- [Chapter 1, “SMS Overview,” on page 7](#)
- [Chapter 2, “Installing and Configuring SMS,” on page 11](#)
- [Chapter 3, “Using SMS,” on page 19](#)
- [Chapter 4, “Optimizing SMS,” on page 29](#)
- [Chapter 5, “Coexistence and Migration Issues,” on page 35](#)
- [Chapter 6, “Troubleshooting SMS,” on page 39](#)
- [Appendix A, “TSA Features,” on page 45](#)
- [Appendix B, “Creating SMS Debug logs,” on page 51](#)
- [Appendix C, “POSIX File System Support in OES Linux,” on page 55](#)
- [Appendix D, “Documentation Updates,” on page 57](#)

Audience

The audience for this guide includes network administrators

Feedback

We want to hear your comments and suggestions about this manual and the other documentation included with this product. Please use the User Comments 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.

Additional Documentation

For documentation on the SBCON utility, the basic storage management engine for NetWare see the *Open Enterprise Server SBCON Administration Guide*.

Documentation Conventions

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

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

When a single pathname can be written with a backslash for some platforms or a forward slash for other platforms, the pathname is presented with a backslash. Users of platforms that require a forward slash, such as Linux* or UNIX* , should use forward slashes as required by your software.

Novell® Backup infrastructure (Storage Management Services or SMS) provides backup applications with the framework to develop complete backup and restore solutions. SMS helps back up file systems (such as NSS) or applications (such as GroupWise®) on NetWare® and SUSE® Linux Enterprise Server (SLES) to removable tape media or other media for off-site storage.

The following topics are discussed in this section:

- [Section 1.1, “Backup Properties and Operations,” on page 7](#)
- [Section 1.2, “SMS Components,” on page 7](#)
- [Section 1.3, “Backup Applications,” on page 9](#)
- [Section 1.4, “What’s New,” on page 10](#)

1.1 Backup Properties and Operations

A logical backup typically involves a backup server and a target server. The backup server hosts the backup application and possibly the tape device as well. The target server contains the data that needs to be protected and is also known as the backup target.

A backup target, in turn, can be a file system or an application. With a file system target, entities that are backed up are files and directories along with their associated metadata. With an application target, application-specific objects are exposed for backup. For example, a User object maybe exposed to determine backup of a particular mailbox.

A typical backup must allow for selection, filtering and control of what entities are backed up. This processing granularity provides tremendous benefits during a restore operation where an administrator has the ability to restore specific entities, such as a file or an application-specific object.


SMS provides a framework that can provide this functionality. The most significant property of SMS is its definition of a single consistent interface for all file systems and applications across NetWare and SLES. Backup applications can thus provide the backup administrator with selection and filtering operations in a consistent manner across all backup targets.

1.2 SMS Components

In order to achieve full functionality, SMS is implemented as two independent components that provide the following functional abstractions:

- **Storage Management Data Requester (SMDR)** provides remote connectivity and transfers data between the target and the backup servers.
See [Section 1.2.1, “Storage Management Data Requester,” on page 8](#) for more information.
- **Target Service Agent (TSA)** provides abstraction details of the specific target being backed up.
See [Section 1.2.2, “Target Service Agent \(TSA\),” on page 8](#) for more information.

Description: Description: SMS components on the backup server and the target servers



The backup process is explained below:

- A typical backup involves the backup application using the SMDR on the backup server to communicate with the target server. The SMDR on the target server uses a TSA to read and abstract the target data.
- The backup application uses a formatted buffer delivered by the TSA and the SMDR to send it to a storage medium such as a tape drive.
- Every target server needs to have its own TSA that understands the target-specific objects. If a new target needs to be backed up, only a new TSA needs to be added and the entire backup infrastructure can be reused.

1.2.1 Storage Management Data Requester

The Storage Management Data Requester (SMDR) is the communication module in the SMS architecture. The SMDR defines the API framework, provides remote connectivity, and abstracts the details of any communication between the servers. Thus, SMDR is capable of transferring any target data between the target and backup server. Most backup applications use the API exposed by SMDR to make use of functionality exposed by SMS. For information on configuring SMDR, See [Section 2.4, “Configuring SMDR,” on page 12](#).

1.2.2 Target Service Agent (TSA)

The Target Service Agent (TSA) provides an implementation of SMS APIs for a particular target. The TSA provides transparency by abstracting details of the specific service (such as GroupWise or NSS) being backed up. For example, various backup applications use file system TSA to back up and restore NSS file system data and metadata (trustee assignments, namespaces, and file attributes). A TSA understands the target and knows how to scan, read, and write a particular target's data. Each target needs a TSA.

Table 1-1 Target Services and Their Corresponding Target Service Agents in OES

| Target Service | Target Service Agent | Platforms |
|-----------------|----------------------|---------------|
| NSS file system | TSAFS | NetWare, SLES |

| Target Service | Target Service Agent | Platforms |
|---------------------------------|----------------------|-------------------|
| Cluster resources | TSAFS | NetWare, SLES |
| Traditional NetWare file system | TSAFS | NetWare 6.x |
| | TSA500 | NetWare 5.1 |
| VFS-compliant file systems | TSAFS | SLES |
| eDirectory™ | TSANDS | NetWare |
| GroupWise | TSAFS | NetWare 6.x, SLES |
| | GW TSA | NetWare 5.1 |

- TSAFS backs up a target server, so it services all file systems (and possibly cluster resources) on a particular target server.
- GroupWise backup functionality is included with the file system TSA.

However, the functionality does not provide object level backup, but simply ensures that GroupWise database backups are consistent by freezing the GroupWise database before a regular file system backup. This functionality is not turned on by default. See [“File System TSA \(TSAFS\)” on page 9](#) to turn on this functionality as required.

File System TSA (TSAFS)

The File System TSA (TSAFS) backs up the traditional NetWare file system as well as NSS on OES NetWare. On OES Linux, TSAFS supports all VFS-compliant file systems and NSS. Some of the salient features are:

- Implements a predictive data caching model that provides improved backup performance.
- Provides parameters to fine-tune performance.
- Provides parameters that can be used to fine-tune performance to the specific environment
- Ability to interpret OES NetWare and OES Linux data streams, so data can move interchangeably between the platforms.
- Cluster enabled
- Multiprocessor enabled
- Compatible with the data format used by existing versions of the TSA.
- Ability to ensure consistency while backing up GroupWise databases.
- Provides a NetWare emulation mode on OES Linux.
- Ability to handle data across locales by providing data in UTF-8 format.

1.3 Backup Applications

SMS provides two backup applications:

- SB CON on OES NetWare
- nbackup on OES Linux

These backup applications include the following functions:

- Provides a basic backup solution that is not designed to provide advanced session management and other features available with third-party backup applications.
- Can be used to create consistent backups, but are recommended for use as technology demonstrators, especially in backup performance.
- Used to troubleshoot problems

Several backup application vendors provide an enterprise backup solution using SMS.

For more information on SBCON, refer the *Open Enterprise Server SBCON Administration Guide*.

For more information on nbackup, refer to the nbackup(1) man page on OES Linux.

1.4 What's New

The following feature is included in OES SP2 release:

- Backup and Restore of Extended Attributes is supported in VFS system.
- UTF8 based trustee migration support for SCMU (Server Consolidation and Migration Utility).

Installing and Configuring SMS

During the installation of Open Enterprise Server (OES), Storage Management Services™ (SMS) is installed by default on both the NetWare® and SUSE® Linux platforms.

- [Section 2.1, “System Requirements,” on page 11](#)
- [Section 2.2, “Updating Existing Installations,” on page 11](#)
- [Section 2.3, “Starting SMS Services,” on page 11](#)
- [Section 2.4, “Configuring SMDR,” on page 12](#)
- [Section 2.5, “Configuring the Target Service Agent for File System,” on page 14](#)

2.1 System Requirements

SMS runs on any system where OES can be installed, so the minimum system requirements for SMS are the same as the requirements for an OES server.

For information regarding minimum server requirements for Linux and NetWare on OES, refer to the following Web sites:

- ❑ [OES Linux documentation Web site \(http://www.novell.com/documentation/oes/index.html?page=/documentation/oes/install_linux/data/front.html#bktitle\)](http://www.novell.com/documentation/oes/index.html?page=/documentation/oes/install_linux/data/front.html#bktitle)
- ❑ [OES NetWare documentation Web site \(http://www.novell.com/documentation/oes/index.html?page=/documentation/oes/install-nw/data/hz8pck9v.html#bktitle\)](http://www.novell.com/documentation/oes/index.html?page=/documentation/oes/install-nw/data/hz8pck9v.html#bktitle)

2.2 Updating Existing Installations

For upgrading an OES NetWare server, refer to the *OES for NetWare Installation Guide*.

Existing SUSE SLES 9 servers can be upgraded to OES. As a part of the upgrade, SMS can also be selected for installation.

For more information on how to upgrade to OES Linux, refer to the Installing OES for Linux section in the *OES Linux Installation Guide*.

OES supports Red Carpet® updates to the servers. SMS is also available via Red Carpet updates. For more information on how to configure OES Linux server to receive updates via Red Carpet channels, refer to the Updating OES for Linux using Red Carpet Channel section in the *OES Linux Installation Guide*.

2.3 Starting SMS Services

This section provides information about starting SMS services on OES:

- [Section 2.3.1, “OES NetWare,” on page 12](#)
- [Section 2.3.2, “OES Linux,” on page 12](#)

2.3.1 OES NetWare

By default, the SMS services SMDR, TSAFS/TSA500, and TSANDS are loaded by the `autoexec.ncf` script. You can also load them by using the `SMSstart.ncf` script provided by the SMS installation.

To load SMS services manually, use the server console to execute `SMSstart.ncf` or load required NLM™ software manually.

2.3.2 OES Linux

On reboot of the server, SMS is started by default at run levels 3 and 5. You can also start SMS services by using a shell to execute the startup script `/etc/init.d/novell-smdrd`.

By default, the configuration file for SMDR autoloads TSAFS. To manually load TSAFS, use `smsconfig`. For more information on using `smsconfig`, refer to [Section 2.5.3, “Using the Command Line for OES Linux,”](#) on page 17.

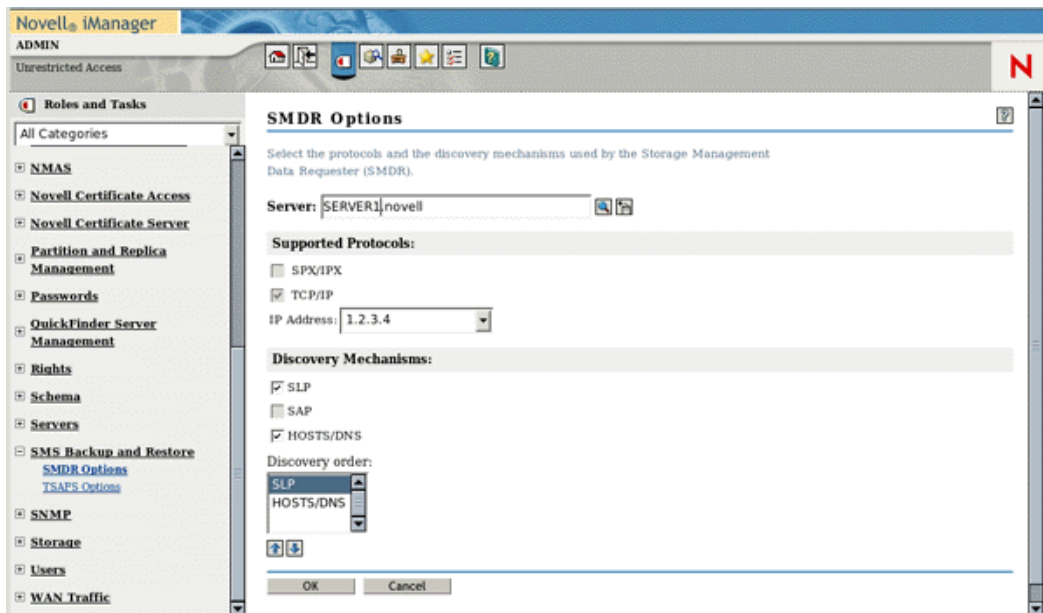
2.4 Configuring SMDR

- [“Using iManager”](#) on page 12
- [“Using the Server Console”](#) on page 13
- [Section 2.4.3, “Using the Command Line for OES Linux,”](#) on page 14
- [Section 2.4.4, “Using a Configuration File for OES Linux,”](#) on page 14

2.4.1 Using iManager

- 1 In iManager, click *SMS*, then click *SMDR Configuration*.

Description: Description: SMDR configuration window



- 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, then click OK.

- **Supported Protocols:** SMDR supports Internet Protocol (IP) and Sequenced Packet Exchange™ (SPX™) or Internet Packet Exchange (IPX™) protocols. On Linux only IP is supported.

On modifying the value for this option, the changed value takes effect the next time when you start SMDR .NLM for NetWare. In case of Linux, the daemon takes the values when you restart the daemon and the clients that use smdr library to take backup will take the values immediately except for the clients that already have established the connections.

- **IP Address:** SMDR can be configured to listen on the specified IP address on a multi-homed server.

On modifying the value for this option, it takes effect when you restart the SMDR .NLM for NetWare.

- **Discovery Mechanisms:** SMDR supports the following discovery and name resolution mechanisms:

The modified value for this takes effect the next time you start SMDR .NLM for NetWare. For Linux, the daemon takes the values when you restart the daemon and the clients that use smdr library to take backup will take the values immediately except for the clients that already have established the connections.

- **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 registers itself in the smdr.novell domain when loaded. The SLP-enabled SMDRs 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 (NetWare only):** 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. In an IP environment, Service Location Protocol (SLP) replaces SAP.
- **Name Resolution Through HOSTS File:** SMDR can be configured to use a hosts file (`sys:\etc\hosts` on NetWare and `/etc/hosts` on Linux) 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 uses these entries to resolve the IP address.

SMDR can be configured to alter the order of server name resolution by using the Discovery Order list box.

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

2.4.3 Using the Command Line for OES Linux

SMDR can be configured using the command line options on OES Linux:

```
smdrd [--(no)slp] [--(no)hosts] [--ip <local ip address>]
```

For more information, refer to the `smdrd(8)` man page in OES Linux.

2.4.4 Using a Configuration File for OES Linux

SMDR configuration file is located at `/etc/opt/novell/sms/smdrd.conf` on OES Linux servers. Each instance of SMDR reads the configuration file for its default configuration.

For more information on how to edit the configuration file, refer to the `smdrd.conf(5)` man page in OES Linux.

2.5 Configuring the Target Service Agent for File System

TSAFS provides configurable parameters to help tune its performance. It also provides configurable parameters to control certain specific features that are supported by it.

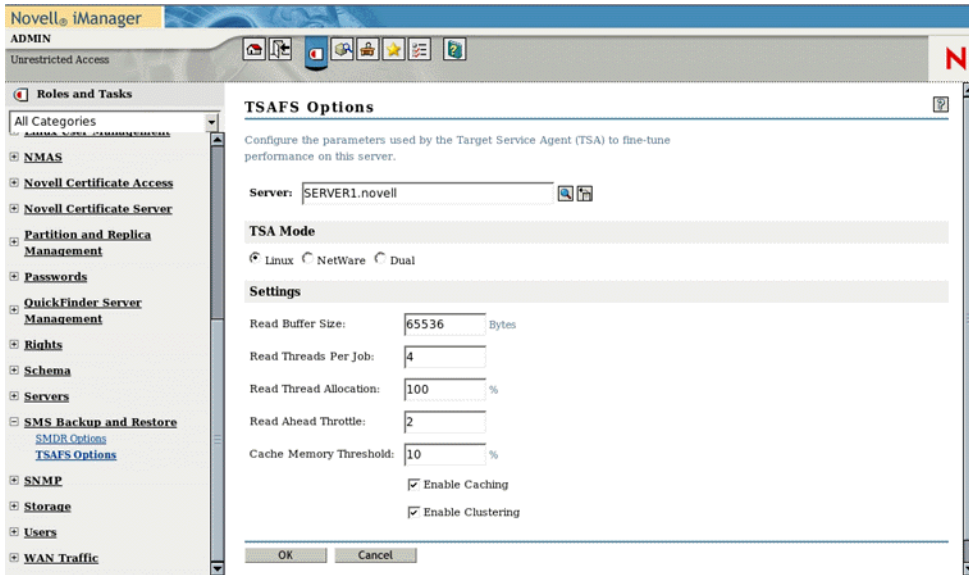
- [“Using iManager” on page 14](#)
- [“Using Server Console for OES NetWare” on page 16](#)
- [Section 2.5.3, “Using the Command Line for OES Linux,” on page 17](#)
- [Section 2.5.4, “Using a Configuration File for OES Linux,” on page 18](#)

2.5.1 Using iManager

Complete the following steps to configure the parameters used by TSA:

- 1 In iManager, click *SMS*, then click *TSA Configuration*.

Description: Description: TSA configuration screen



- 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*.

- **TSA Mode:** TSA Mode can be used to expose NSS as a native Linux file system (Linux Mode) or emulate NetWare file system semantics (NetWare Mode) on OES Linux.

If the Dual Mode is selected, both NetWare and Linux semantics are simultaneously exposed and can be used independent of each other. By default, the TSA Mode is set to Linux. This option is supported only on OES Linux.

The changed value takes effect when you reload the tsafs module.

- **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.
For OES Linux, the modified value takes effect when you reload the tsafs module.
For OES NetWare, the modified value takes effect from the next new connection onwards.
- **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.
For OES Linux, the changed value takes effect when you reload the tsafs module.
For OES NetWare, the modified value takes effect from the next new connection onwards.
- **Cache Memory Threshold:** The percentage of unallocated server memory that the TSA can utilize to store cached datasets. This represents a maximum percentage value of unallocated server memory that the TSA uses to store cached datasets. The default value is 25% of unallocated server memory. The cache memory utilized by the TSA is dynamically reset based on the available unallocated memory.

For OES Linux, the modified value takes effect when you reload the tsafs module.

For OES NetWare, the modified value takes effect when the TSAFS.NLM is reloaded.

- **Enable Caching:** This option can be used to specify if the TSA should do predictive caching during backups. Caching improves backup performance, on certain workloads, by prefetching files in memory. By default, *Enable Caching* is selected.

For OES Linux, the modified value takes effect when you reload the tsafs module.

For OES NetWare, the modified value takes effect from the next new connection onwards

- **Enable Clustering:** Determines cluster support. If the backup server does not support clusters, this option is disabled. Deselect this option if the TSA is running on a cluster node and the backup engine is cluster-enabled. Select it if the TSA is running on a non-cluster node or the backup engine is not cluster-enabled. This is deselected by default.

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

For OES Linux, the modified value takes effect when you reload the tsafs module.

For OES NetWare, the modified value takes effect when the TSAFS .NLM is reloaded.

2.5.2 Using Server Console for OES NetWare

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

Table 2-1 Basic Parameters

| Task | Purpose | Command |
|--|---|--|
| Set the number of read threads for a job | This value 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 / readthreadspersperjob=value |
| Set the read buffer size | This value 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 datasets. | This value is used to specify the percentage of total server memory that the TSA can utilize to store cached datasets. This represents a maximum percentage value of total server memory that the TSA uses to store cached datasets. The default value is 25% of the total server memory. | TSAFS / CacheMemoryThreshold |
| Set the TSA to support failover or fallback on cluster volumes. | Set this value to <code>cluster</code> if the TSA is running on a cluster node and the backup engine supports clusters. Set this value to <code>nocluster</code> if the TSA is running on a cluster node or if the backup engine does not support clusters. The default value is <code>cluster</code> . Running a cluster-enabled TSA on a non-cluster-enabled node does not affect functionality in any way. | TSAFS / <code>cluster</code> <code>nocluster</code> |
| Enable Caching in TSAFS | Set the value to <code>cachingMode</code> if the TSA should do predictive caching during backups. Caching improves backup performance, on certain workloads, by prefetching files in memory. Set the value to <code>noCachingMode</code> if the TSA should not do predictive caching during backups. The default value is <code>cachingMode</code> . | TSAFS / <code>CachingMode</code> <code>noCachingMode</code> |

For advanced configuration, refer to [Section 4.3, “Fine-Tuning SMS Performance,”](#) on page 31.

To enable additional backup features, refer to [Section 3.1.4, “Additional Backup Features,”](#) on page 23.

2.5.3 Using the Command Line for OES Linux

The Storage Management Data Requestor (SMDR) is a daemon process on Linux that holds the information regarding Target Service Agents (TSA) that have registered to it.

The process of registering or de-registering a TSA with SMDR is referred to as a loading or unloading the TSA. The TSAs should be registered with SMDR for the backup engines to access the specific target.

On OES Linux, `smsconfig` is provided to load or unload the (TSA) with the SMDR daemon:

```
smsconfig [ -l | -u | -t ] [TSA module name] [TSA options] ...
```

For example:

| | |
|---|---|
| <code>smsconfig -l tsafs</code> | Registers the TSA 'tsafs' with SMDR |
| <code>smsconfig -u tsafs</code> | De-registers the TSA 'tsafs' from SMDR |
| <code>smsconfig -t</code> | List the TSAs currently registered with SMDR |
| <code>smsconfig -l tsafs --tsaMode=Linux</code> | Set the <code>tsamode</code> to <code>linux</code> to expose NSS as a native Linux file system. |
| | Set the <code>tsamode</code> to <code>netware</code> to expose NSS with NetWare file system semantics. |
| | Set the <code>tsamode</code> to <code>dual</code> , if both NetWare and Linux semantics should be simultaneously exposed and be used independent of each other. |
| | The default value is <code>linux</code> . |

The TSAs also expose their command line configuration interface to the user through `smsconfig`, which enables the user to configure the TSAs. For example, the following command can be used to see TSAFS configurable parameters:

```
smsconfig -l tsafs --help
```

TSAFS can be configured during registration as shown in the following example:

```
smsconfig -l tsafs --ReadThreadsPerJob=6
```

This configuration persists till the TSA is unloaded.

For more information regarding `smsconfig` and `tsafs` configurable parameters, refer to the `smsconfig(1)` and `tsafs(1)` man pages.

For advanced configuration, refer to [Section 4.3, “Fine-Tuning SMS Performance,”](#) on page 31.

To enable additional backup features, refer to [Section 3.1.4, “Additional Backup Features,”](#) on page 23.

2.5.4 Using a Configuration File for OES Linux

The TSAFS configuration file is located at `/etc/opt/novell/sms/tsafs.conf` on OES Linux. When the TSA is loaded, it reads the configuration file for its default configuration.

For more information refer to the `tsafs.conf(5)` man page.

Using SMS

This section provides information on SMS features and how SMS can be used in various scenarios.

Backup applications use SMS as an infrastructure to provide a complete backup solution. For specific information on features available in the backup application, refer to documentation pertaining to the specific application.

The following topics are discussed in this section:

- [Section 3.1, “Backing Up Data Using SMS,” on page 19](#)
- [Section 3.2, “Restoring Data Using SMS,” on page 24](#)
- [Section 3.3, “Moving Data Across OES NetWare and OES Linux,” on page 26](#)

3.1 Backing Up Data Using SMS

This section provides information on how SMS backs up data from Novell®eDirectory™ and from the file system.

- [Section 3.1.1, “Prerequisites,” on page 19](#)
- [Section 3.1.2, “Backing Up the File Systems,” on page 20](#)
- [Section 3.1.3, “Backing Up Clusters,” on page 22](#)
- [Section 3.1.4, “Additional Backup Features,” on page 23](#)

3.1.1 Prerequisites

Meet the following prerequisites before starting the backup software.

Backing Up Open Files

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

To enable CopyOnWrite on a single NSS volume, refer to the [NSS Server Console Commands Web site \(http://www.novell.com/documentation/oes/index.html?page=/documentation/oes/nss_enu/data/ajhvfjd.html#ajhvfjd\)](http://www.novell.com/documentation/oes/index.html?page=/documentation/oes/nss_enu/data/ajhvfjd.html#ajhvfjd). Supervisor rights are required to back up the open files.

Backing Up Compressed files

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

Listed below are few guidelines to make this decision:

- Backups are faster if files are in compressed form. 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, because the compression does not begin immediately.

Backing Up Migrated Files

Files that are not frequently accessed can be moved to tertiary storage by any Hierarchical Storage Management software (HSM Software). These files continue to be available in the form of stubs in the primary storage device. The stubs contain information necessary to access the file contents from the tertiary storage using the HSM software.

During backup it is possible to back up these files in the following manner:

- Back up only the stubs
- Back up both the stubs and the data associated with the file

If the tertiary device itself is backed up independently, choosing to back up only the stub information helps reduce the amount of data. This, in turn, helps save space on tape and increase backup performance because data does not need to be restored from the tertiary device during backup. However, restores require the HSM software to be set up and ensure tertiary storage associations are maintained as they were during the backup.

When both the stub and the data are backed up, the data is restored for the backup process. On restore, either the stub or data or both can be restored. However, backing up migrated file data can impact the backup performance because the data needs to be demigrated from a tertiary storage device. In addition, the backup would include both the target server as well as the tertiary storage data, which requires adequate planning for tape storage.

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 backup and target server.

Refer to the [Section 2.3, “Starting SMS Services,” on page 11](#) for information on how to start SMS services.

3.1.2 Backing Up the File Systems

To back up file system data, TSAFS must be loaded on each target server for which a backup is to be created (see [“Before Running the Backup Software” on page 20](#)).

TSAFS supports backing up:

- File system metadata such as name spaces, extended attributes, trustee rights, and data streams on OES.
- All POSIX* compliant file systems on OES Linux, see [Appendix C, “POSIX File System Support in OES Linux,” on page 55](#).

- NSS file system and associated metadata on OES Linux which are not available through POSIX interfaces.

TSAFS uses the ECMA SIDF standard format to store the file system data. For more information, refer to the [Standard ECMA-208 Web site \(http://www.ecma-international.org/publications/standards/Ecma-208.htm\)](http://www.ecma-international.org/publications/standards/Ecma-208.htm).

This section discusses the following:

- “[Backing Up Trustee Assignments](#)” on page 21
- “[Backing Up Server-Specific Information on OES NetWare](#)” on page 21
- “[Backing Up Links](#)” on page 21
- “[Backing up NCP volumes on OES Linux](#)” on page 22

Backing Up Trustee Assignments

Trustee assignments are stored as part of the file system as an Identifier (ID).

TSAFS uses these IDs to determine the respective fully distinguished names (FDN) and backs up the FDNs. This allows trustee assignments to be restored even if a particular user object was deleted and re-created which would cause the ID to be different. Even if the User object is deleted and re-created with a new ID, the user’s trustee assignments in the file system are restored using the FDN.

For additional information about object ID and trustee issues, see “[Restoring Trustee or Owner Assignments](#)” on page 26.

Backing Up Server-Specific Information on OES NetWare

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.

Backing Up Links

TSAFS supports backup of hard and soft links when backing up POSIX compliant file systems on OES Linux and it supports backup of hard links on OES NetWare.

In the case of hard links, a file is backed up for each instance of a hard link. TSAFS provides an option for backup applications that backs up the file data for only the first instance of the hard link and maintains stubs without backing up file data for subsequent instances. For a successful restore, ensure that the restore includes the first instance. In the case of soft links, TSAFS backs up soft link

information and also data indicating which file it is linked to. If the backup definition does not include the linked file, then the file's data is not backed up.

Backing up NCP volumes on OES Linux

NCP™ server for Linux allows administrators to create NCP volumes on OES Linux. These volumes contain additional meta data for files and directories as compared normal POSIX compliant file systems.

For more information on setting up NCP volumes and meta data maintained by the same refer to the *NCP Server of Linux Administration Guide* (http://www.novell.com/documentation/oes/ncp_lx/data/h9izvdye.html).

TSAFS supports backup and restore of additional meta data for files and directories under NCP volumes. When backing up an NCP volume file or directory, trustee assignments and inherited rights filters for the data set is additionally backed up using the Novell client libraries.

Both backup and restore operations use the Novell client libraries and hence require that the Novell client be installed on the server. For more information on how to install and configure the Novell client refer *Novell Client for Linux Installation and Administration Guide* (http://www.novell.com/documentation/linux_client/index.html)

3.1.3 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 OES 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/oes/index.html?page=/documentation/oes/oes_home/data/bs4p4gx.html#btto1hb).

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 restoration 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 application supports it.

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

- ❑ If cluster-enabled resources are to be backed up or restored, SLP should be used as the discovery mechanism.
- ❑ A cluster node will have clustered and one or more non-clustered volumes. When the particular cluster server is chosen for backup, only the clustered volumes will be listed. To backup non-clustered volumes, choose the physical server instead.

TIP: To treat all cluster volumes as non-clustered for backup, disable the cluster option in TSAFS, refer to **Section 2.5, “Configuring the Target Service Agent for File System,” on page 14**. This will enable listing of all cluster volumes as part of the cluster node instead of virtual server resource.

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

Backing Up Mixed Node Clusters

In OES, it is possible to have mixed node clusters, where different nodes in the cluster may run either OES NetWare or OES Linux. For more information regarding mixed node clusters, refer to the *OES Novell Cluster Services 1.8.1 Administration Guide for Linux*.

TSAFS supports mixed node cluster backup. Path names are represented differently on OES NetWare and OES Linux. In order to achieve a consistent backup, TSAFS supports an NetWare emulation mode on OES Linux. This mode is used to make TSAFS behave as a NetWare target on OES Linux. This resolves any path name conflicts that might arise because of mixed node clusters in the setup. To enable this feature, refer to the “[NetWare Emulation Mode on OES Linux](#)” on page 24.

3.1.4 Additional Backup Features

This section describes additional features supported by TSAFS:

- “[GroupWise Backup](#)” on page 23
- “[Non-caching Mode of Operation](#)” on page 23
- “[Code Page Support on OES Linux](#)” on page 24
- “[NetWare Emulation Mode on OES Linux](#)” on page 24

GroupWise Backup

TSAFS supports backing up GroupWise® database files from the OES. TSAFS is integrated with GroupWise on OES NetWare and OES Linux to provide consistent backups of GroupWise database files by locking them before a backup is taken.

NOTE: This feature ensures that a snapshot of the GroupWise database files are consistent are backed up. This backup cannot be used to restore GroupWise objects such as a particular mailbox or a user object.

To enable the GroupWise backup feature in TSAFS, use the following switches:

- For OES NetWare use:

```
TSAFS /EnableGW
```

- For OES Linux use:

```
smsconfig -l tsafs --EnableGW
```

Non-caching Mode of Operation

TSAFS by default uses a predictive caching mechanism to cache ahead data sets for backup operations. Some backup applications process incremental or differential backups by filtering the data sets themselves rather than use TSAFS options. Under such circumstances the cache built up by TSAFS is not used. This leads to slower backups as TSAFS spends more time caching unwanted data sets.

The non-caching mode of operation disables TSAFS predictive caching thus eliminating any performance issues when used with applications that do their own filtering.

To enable the non-caching mode of operation in TSAFS, use the following switch:

- For OES NetWare use:

```
TSAFS /noCachingMode
```

- For OES Linux use:

```
smsconfig -l tsafs --noCachingMode
```

Code Page Support on OES Linux

By default, TSAFS assumes that filenames on the disk are UTF-8 encoded. If they are not, TSAFS skips these files and reports them in the skipped data set log. In such cases, the following switch can be used to set the appropriate code set for backup and restore:

```
smsconfig -l tsafs --useCodeSet=codeset
```

For more information on codesets, refer to the `tsafs(1)` man page.

NetWare Emulation Mode on OES Linux

TSAFS on OES Linux, by default, exposes Linux File System as the target. TSAFS has a built-in switch that makes it possible to expose the TSA as a NetWare File System. This enables you to use TSAFS on OES Linux as if it is a NetWare target.

NetWare emulation mode can be turned on using:

```
smsconfig -l tsafs --tsaMode=mode
```

where *mode* is `linux`, `netware`, or `dual`. In `linux` mode, the TSA displays only the Linux File System target. In `netware` mode, the TSA displays only NetWare File System target. In `dual` mode, both the targets are displayed.

When connected to the NetWare File System target, you can see only NSS file system resources.

NOTE: NetWare Emulation mode is intended to provide a migration path for applications that are already NetWare-aware and might be deprecated in the future. However, all backups taken using the NetWare emulation mode will be valid and recoverable in all future releases.

3.2 Restoring Data Using SMS

This section provides information on how SMS restores data. For more information about options supported during a restore, refer to the respective backup application documentation.

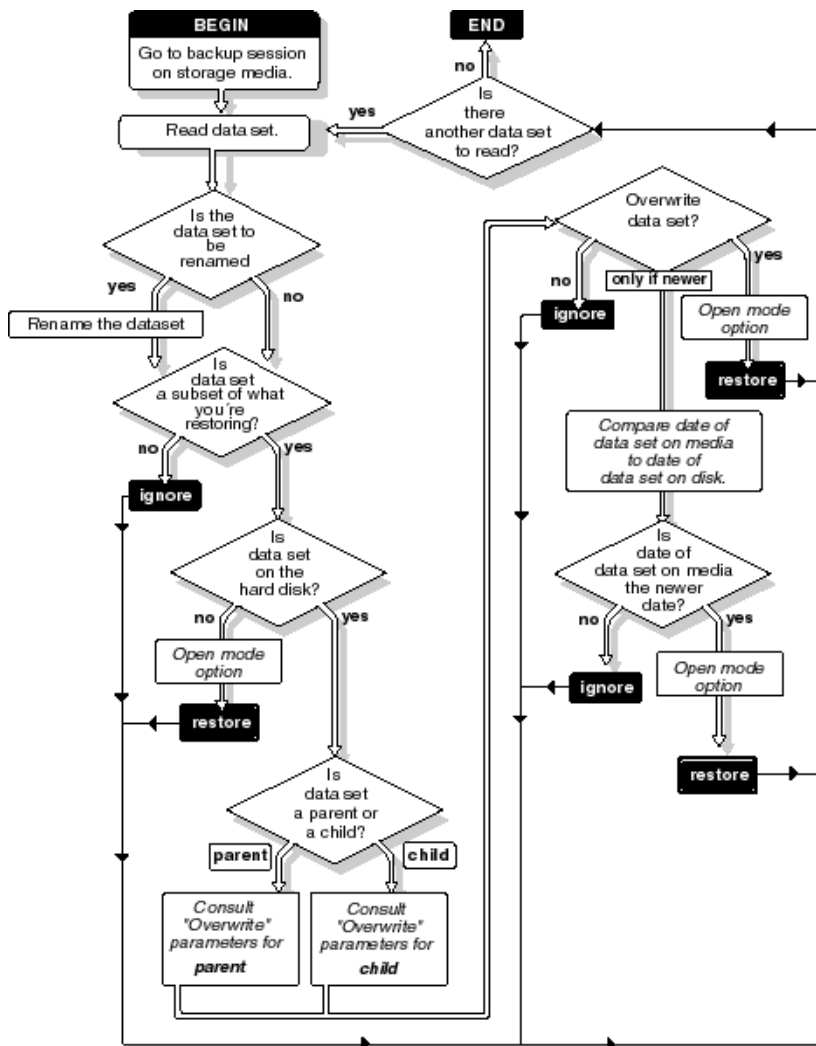
- [“How SMS Restores Data” on page 25](#)
- [“Restoring Trustee or Owner Assignments” on page 26](#)

3.2.1 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?

Description: Description: Flowchart illustrating how SMS restores data



3.2.2 Restoring Trustee or Owner Assignments

The file system backup contains the trustee or owner/group names for files and directories that were backed up. On restoration these names are used to map them back to the corresponding file system object IDs.

If the name-to-ID mapping is unavailable for any reason then the file is restored with the default connection ID. To ensure that the restoration preserves all ID information, update the relevant ID store (eDirectory or the user data base) on the system before you attempt the restore operation.

3.3 Moving Data Across OES NetWare and OES Linux

This section describes steps to move data transparently between OES NetWare and OES Linux:

- [Section 3.3.1, “Prerequisites,” on page 26](#)
- [Section 3.3.2, “Restoring NetWare Backups to NSS on OES Linux,” on page 26](#)
- [Section 3.3.3, “Restoring NetWare Backups to Non-NSS File Systems on Linux,” on page 27](#)

3.3.1 Prerequisites

The following section assumes that on OES Linux TSAFS is running in the default mode of operation and is not using NetWare emulation.

If NetWare emulation mode (see [“NetWare Emulation Mode on OES Linux” on page 24](#)) is being used, then TSAFS on OES Linux can be treated like a NetWare backup or restore.

NFS Name Space Support on OES Linux

OES Linux TSAFS supports only the NFS name space. This is to ensure consistency of pathnames for both NSS and non-NSS file systems on the same server.

LUM-Enabled Trustees and Owners

All trustees and owners of files and folders on the OES Linux NSS file system should be LUM-enabled. This enables setting trustees and owners for files and folders by NSS.

For more information on LUM-enabling users, refer to the [Linux User Management Guide \(http://www.novell.com/documentation/oes/pdfdoc/lumadgd/lumadgd.pdf\)](http://www.novell.com/documentation/oes/pdfdoc/lumadgd/lumadgd.pdf)

For more information regarding trustee and access rights for LUM-enabled users on NSS for OES Linux, refer to the [File System Management Guide \(http://www.novell.com/documentation/oes/index.html?page=/documentation/oes/stor_filesys/data/hn0r5fzo.html#bktitle\)](http://www.novell.com/documentation/oes/index.html?page=/documentation/oes/stor_filesys/data/hn0r5fzo.html#bktitle)

3.3.2 Restoring NetWare Backups to NSS on OES Linux

When restoring data that was backed up from a NetWare source, use the following guidelines

- The NFS name space should be enabled on the source from which the backup was taken.
- The backup should be restored in the NFS name space on the OES Linux server.

- If trustees and owners for files and folders need to be preserved, these users should be LUM-enabled.

On restoring data to NSS volumes on OES Linux, there is no data loss. All attributes and streams that are supported on NetWare are supported on OES Linux NSS volume.

3.3.3 Restoring NetWare Backups to Non-NSS File Systems on Linux

When restoring to NCP volumes under non-NSS file systems, the trustee assignments and inherited rights filters are preserved.

Both backup and restore operations use the Novell client libraries and hence require that the Novell client be installed on the server. For more information on how to install and configure the Novell client refer, *Novell Client for Linux Installation and Administration Guide* (http://www.novell.com/documentation/linux_client/index.html)

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.

- [Section 4.1, “Prerequisites,” on page 29](#)
- [Section 4.2, “Troubleshooting Backup Performance,” on page 30](#)
- [Section 4.3, “Fine-Tuning SMS Performance,” on page 31](#)

4.1 Prerequisites

Backup depends on the combined performance characteristics of the following entities:

- Storage subsystem
- File system
- SMS
- Backup application
- Tape subsystem

You need to optimize each of these entities to ensure that they do not impact the throughput of the system.

4.1.1 Storage Subsystem

For improved performance, it is necessary that all the components should meet the throughput requirements.

If backup is critical and a nontime consuming process, the disk subsystem should be configured to deliver high throughput. Doing parallel I/Os at the disk subsystem level improves the overall disk throughput, it helps the disk/RAID controllers to group the requests better, which reduces the overall seek time and improves the throughput as multiple heads are working at the same time. It is also important to ensure that components do not limit performance throughput delivered by other components.

Connecting Ultra320 disks to an Ultra160 controller or connecting both the Network and disk controllers to the same IO bus limits the backup performance.

Creating this parallelism through optimal configuration ranges from, setting up appropriate RAID levels to the load balancing of the multiple peer-to-peer buses at different levels, from SCSI to PCI. For details on optimizing the storage subsystem, refer to respective hardware reference guides.

File System (NSS) and Applications

- The file system performance tuning and networking parameters should be configured for improved performance. Refer to the operating system documentation for more information on the system tunable parameters.

- 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.
- Different types of files have different impacts on the backup performance. For example, 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 the existing files are overwritten. Compression is not supported in some environments (such as a NetWare® 3.11 server or Novell® Storage Services™ 2.0 volumes or ReiserFS). 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.
- Anti-virus software running at the time of backup significantly slows down the backup process due to checks made on each file access. Most anti-virus software provides options to either ignore backup applications accessing the file system or are tuned to validate modify or write operations alone during a backup process. Since backup is read-centric, the performance is improved significantly.

Refer to the Managing Software RAID Devices section in the *Novell Storage Services File System Administration Guide for OES* for a detailed discussion on NSS tuning parameters.

SMS

SMS can be configured to optimally exploit the underlying subsystem capability by finetuning its working parameters. For more information, see [Section 4.3, “Fine-Tuning SMS Performance,” on page 31](#).

Backup Applications

Backup applications typically process and transfer data obtained from the SMS components to the tape sub-system. Backup applications employ different processing models which have different performance characteristics and features. Most backup applications provide parameters that can be used to optimize performance. For more information, refer to the respective backup application documentation.

Tape Subsystem

The tape subsystem typically consists of the tape drivers, devices and media. It is important to consider the throughput of the device and employ appropriate devices based on the performance needs. In many cases, having a good disk subsystem and a poor tape subsystem limits backup performance. For more information, refer to the appropriate vendor documentation.

4.2 Troubleshooting Backup Performance

This section provides troubleshooting information that you can use to optimize the Backup Performance. [Suggest some info here.](#)

4.2.1 Identify Bottlenecks

TSATEST is a performance analyzer whose main purpose is to aid troubleshooting backup performance bottlenecks. TSATEST, reads data from SMS and discards it, simulating an infinite

performance of the tape system. It also incorporates recommended processing models for backup applications and is used to distinguish performance bottlenecks.

- If backup performance is poor, and TSATEST delivers as much as the backup application, then the bottleneck mostly points to the storage sub-system. For improving storage subsystem performance, refer to [“Storage Subsystem” on page 29](#).
- If backup performance is poor, but TSATEST displays very high throughput on the same storage sub-system, then bottleneck points to the backup application or the tape sub-system. For more information on backup applications and tape sub-systems, refer to [Section 4.1, “Prerequisites,” on page 29](#).

Monitor the disk throughput using Novell Remote Manager on OES NetWare. This shows the efficiency of the disk system to provide data for backup. As backup is an disk I/O bound problem, the above helps to check if SMS delivers comparable performance as the disk system.

Fine-Tune Performance

Once the storage subsystem has been optimized, use the SMS tunable parameters to fine-tune performance. To accomplish and verify if the appropriate tuning parameter does influence performance, TSATEST can be used in various test runs in multiple combinations of TSAFS switches. This provides an inductive way towards detecting optimal sweet spots for your system. For more information on configuration parameters, refer to [Section 4.3, “Fine-Tuning SMS Performance,” on page 31](#).

For more information on TSATEST, refer to the [Novell Developer Site \(http://developer.novell.com/ndk/doc/samplecode/smscomp_sample/tsatest/tsatest.html\)](http://developer.novell.com/ndk/doc/samplecode/smscomp_sample/tsatest/tsatest.html)

4.3 Fine-Tuning SMS Performance

TSATEST is used to determine bottlenecks in the backup systems. For more information, see [Section 4.2, “Troubleshooting Backup Performance,” on page 30](#). Using this information, the following tunable switches can be used to identify sweet spots that helps improve throughput of your backup systems.

4.3.1 Basic Configuration

Configure the following basic tunable parameters to enhance the SMS performance. For more information about configuring the switches, see [Section 2.5, “Configuring the Target Service Agent for File System,” on page 14](#).

Table 4-1 Basic Tunable Parameters to Enhance SMS Performance

| Task | Purpose | Field Name in the iManager Interface | Command |
|---|--|--------------------------------------|---|
| Set the number of read ahead threads for a backup 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 and the system load due to other processes 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 or less system load during backup. Also, monitor the disk I/O performance and set the switch to higher values to check if the disk I/O improves and strike a balance between high backup performance and system utilization.</p> | Read Threads Per Job | <p>TSAFS / readthreadsperjob= value</p> |
| 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 64 KB 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 data set being backed up. For example, if the mean size of the data set is 55 KB, set the buffer size to 64 KB so additional buffer is added to the mean size of the data set. This is required for backup of file characteristics and SIDF encoding.</p> <p>The default value is 65536 bytes. This value ranges from 32 KB to 256 KB.</p> | Read Buffer Size | <p>TSAFS / readbuffersize= value</p> |

| Task | Purpose | Field Name in the iManager Interface | Command |
|---|---|--------------------------------------|---|
| Set the percentage of server's free 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's free memory that the TSA uses to store cached data sets.</p> <p>The default value is 10% 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 |

4.3.2 Advanced Configuration

Configure the following advanced tunable parameters to enhance the SMS performance. For more information about configuring the switches, see [Section 2.5, “Configuring the Target Service Agent for File System,”](#) on page 14.

Table 4-2 *Advanced Tunable Parameters to Enhance SMS Performance*

| 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 readthreadsperjob 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 readthreadallocation 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= value |

Coexistence and Migration Issues

One of the top priorities in designing Novell® Open Enterprise Server (OES) was to ensure that new OES components, running on either NetWare® or Linux, can be introduced into an existing network environment without disrupting any of the products and services that are in place. It was also deemed important that there be a clear migration path for moving existing products or services and related data onto the OES platform. This section discusses the issues involved in the coexistence and migration of SMS in OES. It is divided into the following sections:

- [Section 5.1, “Coexistence,” on page 35](#)
- [Section 5.2, “Migration,” on page 37](#)

For a general discussion of coexistence and migration issues in OES, see the *OES Coexistence and Migration Guide*. [OES Coexistence and Migration Guide](#).

5.1 Coexistence

This section provides information regarding the coexistence of the OES version of SMS with existing NetWare or Linux networks, and with previous versions of the product. The following topics are discussed:

- [Compatibility](#)
- [Coexistence Issues](#)

5.1.1 Compatibility

- [SMDR](#)
- [Using iManager](#)
- [TSAFS](#)

SMDR

- [“Wire Compatibility of the SMDR Protocol” on page 35](#)
- [“Discovery Protocols Used in SMDR” on page 35](#)

Wire Compatibility of the SMDR Protocol

The SMDR (see [Section 1.2.1, “Storage Management Data Requester,” on page 8](#)) wire protocol is fully compatible between Linux and NetWare platforms. In other words, there are no changes in the wire protocol. This enables SMDR to communicate with other SMDRs on the same network irrespective of the platform that it is running on.

Discovery Protocols Used in SMDR

SMDR uses standard discovery and name resolution protocols on both NetWare and Linux. SMDR is enabled to SLP protocol version 1 upwards. SMDR also uses the hosts file to discover other SMDRs on the network and supports a policy ladder implementation to describe the order of priority of using any one mechanism over the other. All these methods are consistent and compatible on both

NetWare and Linux.SMDR on NetWare uses SAP as an additional discovery mechanism. See [“OES Linux on IPX-Based Networks” on page 37](#) for more information on how compatibility for SAP can be achieved. For more information, see [Section 2.4, “Configuring SMDR,” on page 12](#).

Using iManager

You can use iManager to configure SMS services on OES. You can also use the SMS plug-in to configure SMS in an existing network with older NetWare servers. The following table shows which versions of iManager are compatible with SMS on OES and pre-OES NetWare servers,

| Target Servers / iManager Versions and platform | OES Linux | OES NetWare | NetWare 6.x |
|---|---------------|-------------|-------------|
| iManager 2.5 on OES Linux | Supported | Supported * | Supported# |
| iManager 2.5 on OES NetWare | Supported | Supported | Supported |
| iManager 2.0.x on NetWare | Not Supported | Supported | Supported |

* Needs an openwbem running on the NetWare server or sys: volume of the NetWare server to be exported as sys in CIFS exports# Needs the sys: volume on the NetWare server to be exported as sys in CIFS exports

TSAFS

- [“Data Stream Compatibility” on page 36](#)

Data Stream Compatibility

TSAFS (see [“File System TSA \(TSAFS\)” on page 9](#)) uses the ECMA standard SIDF to format file system information into data streams. These streams are supplied to a backup application during backup. Backup applications usually present these streams during a restore, and the TSAFS interprets them. In OES, TSAFS provides full data stream compatibility between NetWare and Linux. In other words, TSAFS on Linux continues to maintain backward compatibility with all existing backups. That is, if the backup application provided TSAFS with a NetWare data stream from an older backup, it is capable of restoring this data correctly to NSS on OES Linux without any data loss. However, if an attempt is made to restore data from a NetWare file system or NSS file system, backing up to a non-NSS file system on OES Linux would create data loss due to the inherent differences in file system semantics. The following list indicates some of the metadata that is lost during a restore of NetWare file system or NSS file system data to non-NSS file systems on OES Linux:

- Secondary data streams
- Extended attributes
- Trustees
- File owner/modifier/archiver information
- Inherited rights filters
- Directory quotas
- User space restrictions

- File attributes such as hidden, rename inhibit or copy inhibit
- File characteristics such as compressed, migrated and sparse

5.1.2 Coexistence Issues

- [Backup Application Support for OES Linux File System Backup](#)
- [OES Linux on IPX-Based Networks](#)

Backup Application Support for OES Linux File System Backup

SMS services are consumed by various commercial backup applications. Backup applications might need to be upgraded to enable backing up of OES Linux. For more information, refer to the backup application documentation. SMS also supports a NetWare emulation mode (see [“NetWare Emulation Mode on OES Linux” on page 24](#)) where the Linux TSAFS exposes the system as though it were a native NetWare system. Some backup applications might use this option in the interim, while they move to a broader solution. Although the emulation mode itself might be deprecated in future (After all backup applications have moved to backing up OES Linux natively), data backed up using this option would be recoverable by all future TSAs. SBCON and nbackup (see [Section 1.3, “Backup Applications,” on page 9](#)) support backing up of OES Linux. However, these backup applications are technology demonstrators and are not positioned as enterprise backup applications.

OES Linux on IPX-Based Networks

SMDR on NetWare can be configured to use Service Advertising Protocol (SAP*) for locating other SMDRs in an IPX™ environment. SAP is not supported on OES Linux, so in a pure IPX environment SMDR on OES Linux cannot discover or resolve SMDRs on NetWare and vice versa. For SMS services on OES to work independent of platforms, discovery and name resolution protocols that are supported by SMDR must be common across all the platforms. For more information on supported protocols by SMDR, see [Section 2.4, “Configuring SMDR,” on page 12](#).

Using SMS Across Mixed Node Clusters

Novell cluster services provides a migration path wherein a cluster can have a mix of NetWare and Linux nodes. For more information regarding mixed node clusters, refer to the [OES Novell Cluster Services 1.8.1 Administration Guide for Linux](#).

TSAFS supports backing up of NSS file system resources across a mixed node cluster with failover/failback support for the same. To backup cluster resources in a mixed node environment use the TSAFS on OES Linux in the emulation mode of operation. Refer, [“NetWare Emulation Mode on OES Linux” on page 24](#) for more information on how to use the emulation mode.

5.2 Migration

This section provides information on migrating to OES version of SMS with existing previous versions of the product. The following topic is discussed: [Please suggest some info here](#).

5.2.1 Migration Path

As SMS does not store any data on the server apart from its configuration files, you just need to install SMS components to bring up SMS services on the migrated server. If you are migrating an existing NetWare server to Linux, remember the following:

- SMS must be selected as part of OES installation of the new server. This ensures that all relevant software components are installed.
- Configuration files are not migrated and adapted to the OES Linux environment. Any configuration updates must be done using iManager (See [Section 2.4, “Configuring SMDR,”](#) on page 12 and [Section 2.5, “Configuring the Target Service Agent for File System,”](#) on page 14).

Troubleshooting SMS

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

- [Section 6.1, “Startup and Connection Issues on OES NetWare,” on page 39](#)
- [Section 6.2, “Startup and Connection Issues on OES Linux,” on page 39](#)
- [Section 6.3, “Common Backup and Restore Issues,” on page 40](#)
- [Section 6.4, “Backup and Restore Issues on OES NetWare,” on page 42](#)
- [Section 6.5, “Backup and Restore Issues on OES Linux,” on page 43](#)
- [Section 6.6, “Cluster Related Issues,” on page 43](#)

6.1 Startup and Connection Issues on OES NetWare

[“Unable to list or connect to target server or service” on page 39](#)

Unable to list or connect to target server or service

Possible Cause: SLP is not configured properly.

Action: To list the SMDRs with the target service's IP address, enter `display slp services smdr.novell` at the server console. 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.

6.2 Startup and Connection Issues on OES Linux

[“Unable to connect to target server or service” on page 39](#)

[“Unable to modify, SMDR listener IP address using iManager” on page 40](#)

Unable to connect to target server or service

Possible Cause: SLP is not configured properly.

Action: Check SLP DA configuration and restart SLP services, followed by SMDR.

Possible Cause: SMDR failed to register to SLP because SMDR was loaded before SLP services were started.

Action: Restart SLP services followed by SMDR.

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

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

Unable to modify, SMDR listener IP address using iManager

Possible Cause: Changing the SMDR listener IP address is not supported through iManager.

Action: Manually change the IP address for SMDR listener in `smdrd.conf` file located at `/etc/opt/novell/sms/`. See `smdrd.conf(5)` man page for details on editing the configuration file.

6.3 Common Backup and Restore Issues

[“Backup does not include the modifications” on page 40](#)

[“Backup is slow” on page 40](#)

[“Files were restored but the error file contains a message specifying the name space formats not restored” on page 41](#)

[“Restore is slow” on page 41](#)

[“Unable to backup open files in NSS Volumes” on page 41](#)

[“Unable to restore compressed files” on page 41](#)

[“Unable to set the data set name” on page 41](#)

[“Out of disk space” on page 42](#)

[“DOS namespace based restores fail for certain files” on page 42](#)

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 [Section A.2, “Backup Types,” on page 49](#).

Possible Cause: The modified date filter for backup was set incorrectly.

Action: Set the modified date filter to the last full or differential backup based on the backup type.

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

Action: Modification are backed up based on the modify bit or modified date filter. Check for applications on the server that could be clearing the bit.

Backup is slow

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

Action: See [“Backing Up Compressed files” on page 19](#).

Possible Cause: Migrated files are being backed up by demigrating the data.

Action: See [“Backing Up Migrated Files” on page 20](#).

Possible Cause: Background file compression and backup are running at the same time.

Action: Schedule background file compression before or after backup is scheduled. Based on how compressed files are backed up (see [“Backing Up Compressed files” on page 19](#)) the order of scheduling can be decided.

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.

Restore is slow

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

Action: Schedule restore and background compression of files at different times.

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

Action: Back up compressed files as compressed in the future, if the restore is going to be to a compressed file system. See [“Backing Up Compressed files” on page 19](#).

Unable to backup 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, see the [“Backing Up Open Files” on page 19](#).

Unable to restore compressed files

Possible Cause: Compression is not supported in some environments such as Novell Storage Services 2.0 or ReiserFS on OES Linux.

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

Possible Cause: The file system that the data is restored to does not have the compression feature enabled.

Action: Enable the compression feature on the file system that data is being restored to.

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, but while restoring names in all name spaces there were name conflicts on the non-primary name spaces. This causes a failure, with an error message similar to the following:

```
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 does not impact the restore process.

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 users who are affected.

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

Action: Increase the directory quota on the affected directories then restart the restore.

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

Action: Compress the files before the restore.

DOS namespace based restores fail for certain files

Possible Cause: DOS names are usually name mangled forms of their LONG or other namespace formats. These names are often auto-generated by the file system. There could be a name clash because of existing files or directories on the system having similar mangled DOS names as that of the data sets being restored.

Action: Restore using non-DOS name spaces.

Backed up filenames display characters of the form [xxxx]

Possible Cause: When converting characters from Unicode to MBCS for display or to return to the backup application it is possible to have characters that do not map to any valid MBCS character on the server locale. Such unmappable characters are displayed in the square bracket notation.

Action: This does not impact the backup process. File names are stored in Unicode* and MBCS formats during backup and hence on restoring such files the restored files will have the correct name convention.

6.4 Backup and Restore Issues on OES NetWare

[“Out of Memory” on page 42](#)

[“Restore of zero byte extended attributes on Traditional NetWare volumes fail” on page 43](#)

[“Moving data between Traditional volumes gives errors” on page 43](#)

Out of Memory

Possible Cause: The server is running on lower memory than the NetWare® system requirements.

Possible Cause: Many applications are running when backups are scheduled.

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

Action: Set the **CacheMemoryThreshold** parameter to a lower value. By default, this is set to 10% of the unallocated server memory.

Restore of zero byte extended attributes on Traditional NetWare volumes fail

Possible Cause: Extended attributes are named streams that are created to store additional metadata that is not processed by the NetWare file systems. Traditional volumes on NetWare do not support zero byte extended attributes.

Action: Restore data to NSS volumes on NetWare for the extended attribute to be available.

Moving data between Traditional volumes gives errors

Possible Cause: The source and destination Traditional volumes are on different server locales. Traditional file systems store names in MBCS format based on the server locale, so these names might not be converted correctly to different locales.

Action: Restore data to NSS volumes on NetWare if multiple locale support is required on the file system.

6.5 Backup and Restore Issues on OES Linux

[“Certain files do not get backed up on non-NSS file systems” on page 43](#)

[“Backup or restore hangs on submitting a requestshifted from \(Startup and Connection Issues on OES Linux\) - smsconfig hangs while loading TSAFS” on page 43](#)

Certain files do not get backed up on non-NSS file systems

Possible Cause: The connection user might not have access to read the files that are not being backed up.

Action: Connect as a user with higher privileges to back up these files.

Possible Cause: These files might have been created using a non-UTF-8 locale. TSAFS uses a UTF-8 locale to work with file names on non-NSS file systems.

Action: Use the TSAFS option as detailed in [“Code Page Support on OES Linux” on page 24](#).

Backup or restore hangs on submitting a requestshifted from (Startup and Connection Issues on OES Linux) - smsconfig hangs while loading TSAFS

Possible Cause: Stale mount points on the Linux server. This causes TSAFS to wait indefinitely on the file system APIs.

Action: Check for stale mount points on the system and fix them by either remounting or unmounting the mount point.

6.6 Cluster Related Issues

[“Clustered volumes are not backed up during full server backups” on page 44](#)

[“Cluster pools are not listed on Linux” on page 44](#)

“Reconnect to TSAFS fails when a cluster resource migrates to another cluster node in a mixed mode cluster.” on page 44

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 because the TSA was loaded as cluster-enabled.

Action: Disable cluster support in TSAFS and rerun the backup job. See [Section 2.5, “Configuring the Target Service Agent for File System,”](#) on page 14.

Cluster pools are not listed on Linux

Possible Cause: TSAFS is configured to run in non-cluster mode.

Action: Enable cluster support in TSAFS and rerun the backup job. See [Section 2.5, “Configuring the Target Service Agent for File System,”](#) on page 14.

Possible Cause: SLP configuration is incorrect or SMDR failed to register its services to SLP.

Action: See [“Unable to connect to target server or service”](#) on page 39.

Reconnect to TSAFS fails when a cluster resource migrates to another cluster node in a mixed mode cluster.

Possible Cause: The backup application may require the same user name and password to be available on all nodes in the cluster. During reconnection the same information used for the first node is reused by the application to connect to the other node as well. At times, the user names or their passwords may be different on the two nodes.

Action: Connect using a user name and password that is common for all cluster nodes.

TSA Features

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

- [Section A.1, “TSA Options,” on page 45](#)
- [Section A.2, “Backup Types,” on page 49](#)

A.1 TSA Options

- [“Backup Options” on page 45](#)
- [“Restore Options” on page 48](#)

A.1.1 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 45](#)

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 47](#)

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.

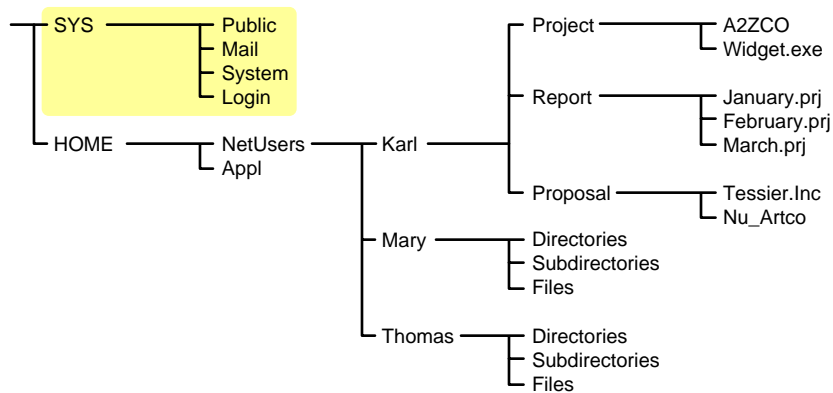
Include

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.

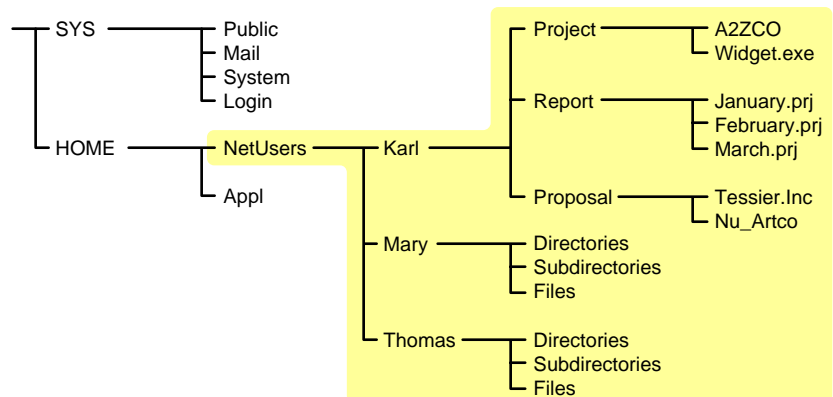
In the figure shown below, volume sys: is selected as an *Include* option. All other areas of the file system structure are excluded from the backup. You can exclude some subdirectories or files beneath your selection if necessary.

Figure A-1 *Include option with a specific volume included, and all others excluded*



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.

Figure A-2 *Include option with a specific directory included, and all others excluded*



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.

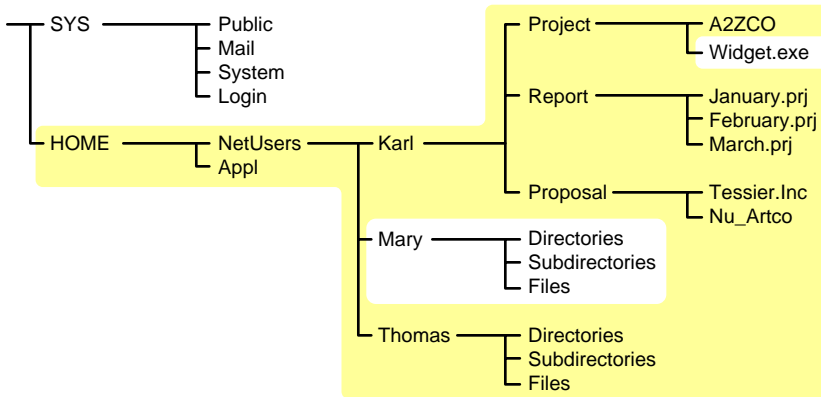
For example, the following command sequence results in volume home being included in the backup with the exception of the mary directory and the widget.exe file.

```
Include major TSA resources home:
```

```
Exclude directories (full path): home:netusers/mary
```

Exclude path/files home:netusers/karl/project/widget.exe

Figure A-3 Example combining SBCON Include and Exclude options



Scanning Data Sets

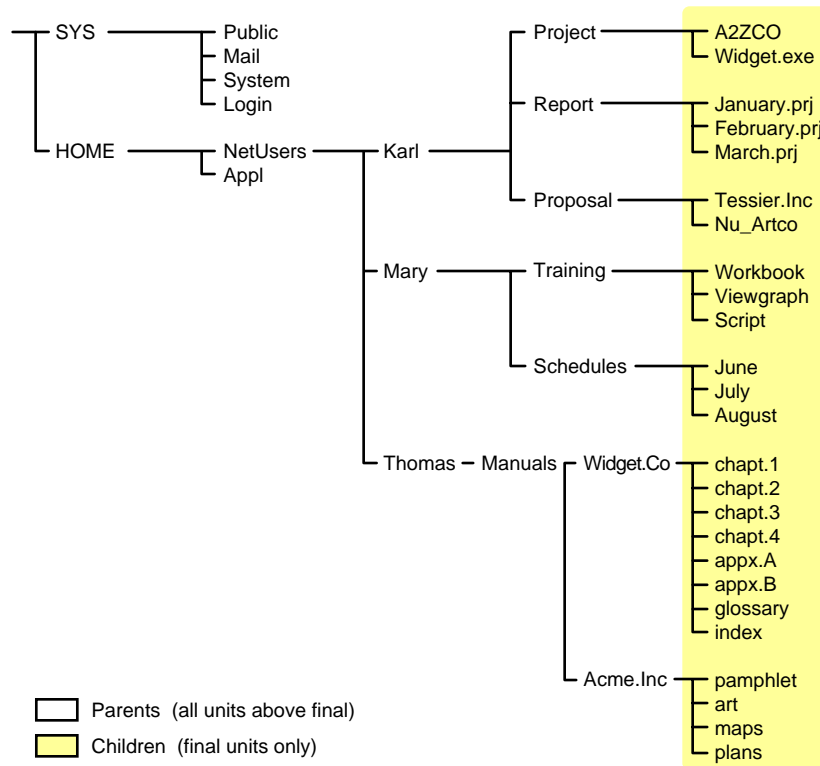
You can specify a different type of data set to be scanned.

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.

Figure A-4 Parent and child levels in a file system



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.

A.1.2 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 45](#).

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.

A.2 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).

NOTE: Incremental and Differential backups are not supported for non-NSS file systems on OES Linux, as these file systems do not support the Modify bit. Use the SBCON option to filter data sets based on the modification time to perform incremental or differential backups.

Creating SMS Debug logs

If SMS issues require technical support from Novell®, it is useful to have a log that provides additional information on the problem. It is especially useful when problems can be reproduced only on specific configurations. SMS provides for creation of debug logs that help technical support provide quicker resolutions.

Currently `smdr.nlm` and `tsafs.nlm` on OES NetWare, and `smdrd` and `tsafs` on OES Linux support creation of debug logs.

- [Section B.1, “Deciding Which Module to Enable for Debug Logging,” on page 51](#)
- [Section B.2, “Enabling Debug Logging,” on page 51](#)
- [Section B.3, “Location of the Debug Log,” on page 52](#)

B.1 Deciding Which Module to Enable for Debug Logging

Although it is useful to have debug logs of all operations, it is necessary to control the size of the logs. In order to achieve this, SMS provides module specific debug logging:

- For problems related to connecting to remote targets, debug logging for SMDR should be enabled. For example, if a backup application is unable to see a remote target SMDR, then enable debug logging for SMDR at the local and remote servers.
- For problems that are related to backup or restore to a particular target service, debug logging should be enabled for that particular target service. For example, debug logging can be enabled for TSAFS to trace a problem where errors are received for a set of files.

As a general rule, if an issue is seen before connecting to a target service, enable debug logging for SMDR. If issue is seen after connecting to a target service, enable debug logging for the target service.

B.2 Enabling Debug Logging

Debug logging uses two switches, `SmsDebug` and `SmsDebug2`, to control the amount of logging. `SmsDebug` is used to log information about a particular feature, and `SmsDebug2` is used to control the level of debug messages required. Although the details of all possible values for these switches are beyond the scope of this document, it is important to note that both these switches must be supplied in the command line to enable debug logging. Novell support might request additional debug logs to be generated for specific features based on the complexity of the problem.

The following command lines produce a debug log that provides general information on all features.

- To enable debug logging on OES Linux:

```
smsconfig -l tsafs --smsDebug=FFFFFFFFC --smsDebug2= FFFFDFFC
```

- To disable debug logging on OES Linux:

```
smsconfig -l tsafs -smsDebug=0 -smsDebug2=0
```

- To enable debug logging on OES NetWare:

```
tsafs /smsDebug=FFFFFFFFC /smsDebug2=FFFFDFFC
```

- To disable debug logging on OES NetWare:

```
tsafs /smsDebug=0 /smsDebug2=0
```

NOTE: On OES NetWare debug logging can be enabled at runtime. On OES Linux, debug logging must be enabled during loading of smdr or tsafs. To ensure that the debug logs have been created and all the information has been flushed to the files, unload the modules before collecting the log files.

B.3 Location of the Debug Log

- [Section B.3.1, “Debug Log Location on OES NetWare,” on page 52](#)
- [Section B.3.2, “Debug Log Location on OES Linux,” on page 52](#)

B.3.1 Debug Log Location on OES NetWare

By default, the debug files are created in `sys:\system\tsa`. The debug file name for `smdr.nlm` is `smdrdbg.log` and for `tsafs.nlm` is `tsadbq.log`. Debug messages are appended to these files each time the NLM™ programs are run with debug options.

To change the location where the debug logs are created, use the `DebugFileName` switch. For example, to change the debug file location to `TESTVOL:DebugDir\tsafs.log` for TSAFS, enter the following command:

```
tsafs /DebugFileName=TESTVOL:DebugDir\tsafs.log
```

B.3.2 Debug Log Location on OES Linux

By default, the debug files are created in `/var/opt/novell/log/sms/`. The debug file name for `smdrd` is `smdrd_debug_pid of smdrd.log`. The debug file name for `tsafs` is `tsafs_debug_pid of smdrd.log`.

Each time `smdrd` is run with debug options, new debug log files are created. If TSAFS is loaded and unloaded with debug options without restarting SMDR, debug messages are appended to the same TSAFS debug file.

To change the location where the debug logs are created, use the `DebugFileName` switch. For example, to change the debug file location to `/home/testuser/smdrd`, enter the following command:

```
smdrd -debugfilename=/home/testuser/smdrd -smsdebug=ffffffffc -  
smsdebug2=ffffdfc
```

B.3.3 Reducing the Debug Log Size

Debug logs can take a large amount of space depending on the length and nature of operations performed with SMS. The following procedures can help reduce the size of the debug log:

- Ensure that SMS modules are run in debug mode for only those operations that are causing problems. If multiple backup or restore jobs are running at the same time, this would increase the size of the debug logs.

- Ensure that only a few files or directories are included. To reduce the size and time taken to generate log files, try to narrow down a problem to a particular file or directory you suspect is causing a problem.
- Turn off debug logging after the debug process is complete. Debug logging can impact performance as well as generate unnecessary debug log information that occupies disk space.
- Compressing the log files before sending them to Novell Support also helps in getting the logs to arrive quickly even if slower links are involved in the transfer.

POSIX File System Support in OES Linux

This section provides information regarding the TSAFS support of POSIX-compliant file systems on OES Linux. These file systems are sometimes referred to in the document as non-NSS file systems.

POSIX-compliant means file systems that comply to the IEEE Std 1003.1 system interfaces. For more information, See [Open Group Publications Web site \(http://www.unix.org/single_unix_specification\)](http://www.unix.org/single_unix_specification).

The following table lists metadata that is backed up or restored to non-NSS file systems on OES Linux. The table uses the definition of metadata structure fields from the definition of the structure stat. (See man page stat(2) for more information)

| Metadata | Description |
|----------|--|
| st_mode | Mode of the file, including File types and File access permissions |
| st_nlink | Number of hard links to the file |
| st_uid | User ID of the file |
| st_gid | Group ID of the file |
| off_t | Size of the file |
| st_atime | Time of last access |
| st_mtime | Time of last data modification |

File Types

- Block special files
- Character special files
- Regular files
- Directories
- Symbolic links
- Socket files

Additionally, TSAFS also backs up the following information for a file or directory (when applicable),

- Symbolic link information
- Data stream
- Extended ACLs (POSIX Draft ACLs)
- Extended attribute streams
- File attributes on a Linux second extended file system

For more information on extended ACLs, see [POSIX Access Control Lists on Linux \(http://www.suse.de/~agruen/acl/linux-acls/online/\)](http://www.suse.de/~agruen/acl/linux-acls/online/).

For more information on file attributes on a Linux second extended file system, see man page `chattr(1)`, installed by the RPM `e2fsprogs`.

Documentation Updates

D

D.1 December 23, 2005

- Added the Troubleshooting section [“Backup or restore hangs on submitting a request shifted from \(Startup and Connection Issues on OES Linux\) - smsconfig hangs while loading TSAFS”](#) on page 43
- Updated the section [Appendix C, “POSIX File System Support in OES Linux,”](#) on page 55
- Page design reformatted to comply with revised Novell[®] documentation standards.

D.2 August 19, 2005

- Added the section [“Backing up NCP volumes on OES Linux”](#) on page 22.
- Added the section [“Using SMS Across Mixed Node Clusters”](#) on page 37.
- Updated the sections [“Using iManager”](#) on page 12, [“Using the Server Console”](#) on page 13, [“Common Backup and Restore Issues”](#) on page 40, and [Appendix C, “POSIX File System Support in OES Linux,”](#) on page 55 for the OES Support Pack 1 release.

D.3 May 26, 2005

- Updated the links for several manuals referred in the book.
- Added an appendix with Documentation Updates information.

