
Open Enterprise Server 2015 SP1

Storage Management Services

Administration Guide

June 2016

Legal Notices

For information about legal notices, trademarks, disclaimers, warranties, export and other use restrictions, U.S. Government rights, patent policy, and FIPS compliance, see <https://www.novell.com/company/legal/>.

Copyright © 2016 Novell, Inc. All Rights Reserved.

Contents

About This Guide	7
1 SMS Overview	9
1.1 Backup Properties and Operations	9
1.2 SMS Components	9
1.2.1 Storage Management Data Requester	10
1.2.2 Target Service Agent (TSA)	10
1.3 Backup Applications	11
2 What's New or Changed in Storage Management Services (SMS)	13
2.1 What's New or Changed in Storage Management Services (SMS) (OES 2015 SP1)	13
2.2 What's New or Changed in Storage Management Services (SMS) (OES 2015)	13
3 Installing and Configuring SMS	15
3.1 System Requirements	15
3.2 Updating Existing Installations	15
3.3 Starting SMS Services	15
3.4 Configuring SMDR	15
3.4.1 Using iManager	16
3.4.2 Using the Command Line	16
3.4.3 Using a Configuration File	16
3.5 Configuring the Target Service Agent for File System	17
3.5.1 Using iManager	17
3.5.2 Using the Command Line	18
3.5.3 Using a Configuration File	19
4 Using SMS	21
4.1 Backing Up Data Using SMS	21
4.1.1 Prerequisites	21
4.1.2 Backing Up the File Systems	22
4.1.3 Backing Up Clusters	24
4.1.4 Additional Backup Features	25
4.2 Restoring Data Using SMS	27
4.2.1 How SMS Restores Data	27
4.2.2 Restoring Trustee or Owner Assignments	28
4.2.3 Limitations for Restoring Data Using SMS	29
4.3 Restoring NetWare Backups to OES 2015 SP1 File Systems	29
4.3.1 Prerequisites	29
4.3.2 Restoring NetWare Backups to NSS	30
4.3.3 Restoring NetWare Backups to Non-NSS File Systems	30
5 Optimizing SMS	31
5.1 Prerequisites	31
5.1.1 Storage Subsystem	31
5.2 Troubleshooting Backup Performance	32

5.2.1	Identify Bottlenecks	33
5.3	Fine-Tuning SMS Performance	33
5.3.1	Basic Configuration	33
5.3.2	Advanced Configuration	36
5.3.3	Additional Parameters	38
6	Running Storage Management Services in a Virtualized Environment	39
7	SMS Security for SMDR	41
7.1	Using SSL in SMDR	41
7.1.1	Using SSL without Certificates	41
7.1.2	Using SSL with Certificates	41
7.1.3	Password-Encrypted Private Key Files	42
7.2	SMDR as a Client and Server	42
7.3	Configuration Options	43
7.3.1	Server Certificate Options	43
7.3.2	Client Certificate Options	43
7.3.3	Miscellaneous Options	44
7.3.4	SSL Option Considerations	44
8	Coexistence and Migration Issues	45
8.1	Coexistence	45
8.1.1	Compatibility	45
8.1.2	Coexistence Issues	46
8.2	Migration	47
9	Troubleshooting SMS	49
9.1	Startup and Connection Issues	49
9.2	Common Backup and Restore Issues	50
9.3	Backup and Restore Issues	52
9.4	Cluster Related Issues	53
A	TSA Features	55
A.1	TSA Options	55
A.1.1	Backup Options	55
A.1.2	Restore Options	58
B	Creating SMS Debug logs	61
B.1	Deciding Which Module to Enable for Debug Logging	61
B.2	Enabling Debug Logging	61
B.2.1	SMDR	62
B.2.2	TSAFS	62
B.3	Location of the Debug Log	62
B.3.1	Debug Log Location	63
B.3.2	Reducing the Debug Log Size	63
C	POSIX File System Support	65
D	SMSLS Utility	67
D.1	Syntax	67

D.2	Options	67
D.3	Examples	68

About This Guide

This guide describes how to use Storage Management Services (SMS) on the Open Enterprise Server (OES) 2015 SP1.

This guide is divided into the following sections:

- ◆ Chapter 1, “SMS Overview,” on page 9
- ◆ Chapter 2, “What’s New or Changed in Storage Management Services (SMS),” on page 13
- ◆ Chapter 3, “Installing and Configuring SMS,” on page 15
- ◆ Chapter 4, “Using SMS,” on page 21
- ◆ Chapter 5, “Optimizing SMS,” on page 31
- ◆ Chapter 6, “Running Storage Management Services in a Virtualized Environment,” on page 39
- ◆ Chapter 7, “SMS Security for SMDR,” on page 41
- ◆ Chapter 8, “Coexistence and Migration Issues,” on page 45
- ◆ Chapter 9, “Troubleshooting SMS,” on page 49
- ◆ Appendix A, “TSA Features,” on page 55
- ◆ Appendix B, “Creating SMS Debug logs,” on page 61
- ◆ Appendix C, “POSIX File System Support,” on page 65
- ◆ Appendix D, “SMSLS Utility,” on page 67

Audience

The guide is intended for network administrators.

Documentation Updates

For the most recent version of the *Storage Management Services Administration Guide*, see the [Novell Open Enterprise Server 2015 documentation Web site \(http://www.novell.com/documentation/oes2015/index.html\)](http://www.novell.com/documentation/oes2015/index.html)

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.

Additional Documentation

- ◆ For documentation on SBCON utility for NetWare, see the [NetWare 6.5 SP8 Documentation Web Site \(http://www.novell.com/documentation/nw65/\)](http://www.novell.com/documentation/nw65/).

1 SMS Overview

Micro Focus 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 Open Enterprise Server 2015 SP1 servers 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 9](#)
- ♦ [Section 1.2, “SMS Components,” on page 9](#)
- ♦ [Section 1.3, “Backup Applications,” on page 11](#)

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.

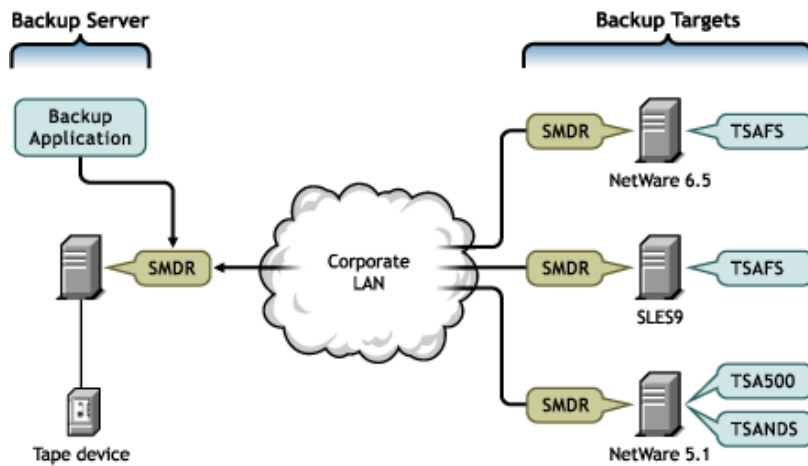
A typical backup must allow for selection, filtering and control of what entities are backed up. This processing granularity provides tremendous benefits during a restore operation where an administrator has the ability to restore specific entities, such as a file or an application-specific object.

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 on an OES server. 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 10](#) 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 10](#) for more information.



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 3.4, "Configuring SMDR,"](#) on page 15.

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
NSS file system	TSAFS
Cluster resources	TSAFS
VFS-compliant file systems	TSAFS
NetIQ eDirectory	TSANDS
GroupWise	TSAFS
	EnableGW
iFolder	TSAIF

- ◆ 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 11](#) to turn on this functionality as required.

File System TSA (TSAFS)

The File System TSA (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 data streams.
- ◆ 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.
- ◆ Ability to handle data across locales by providing data in UTF-8 format.

1.3 Backup Applications

SMS provides the `nbackup` backup application on the OES server. It includes 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.

The `nbackup` utility can be run only as an eDirectory user. Active Directory users cannot run this utility.

For information, see the `nbackup(1)` man page.

2 What's New or Changed in Storage Management Services (SMS)

2.1 What's New or Changed in Storage Management Services (SMS) (OES 2015 SP1)

SMS in OES 2015 SP1 has been modified for bug fixes. There are no new features or enhancements in OES 2015 SP1.

2.2 What's New or Changed in Storage Management Services (SMS) (OES 2015)

In addition to bug fixes, OES 2015 provides the following enhancement for SMS:

- ♦ **Support for the NSS 64-bit File System:** SMS in OES 2015 has been modified to support the NSS 64-bit file system.
- ♦ **Support for Active Directory Trustees** SMS now supports backup and restore of Active Directory trustees. You must ensure that the NSS volumes are AD enabled to restore AD trustees.

3 Installing and Configuring SMS

During the installation of Open Enterprise Server (OES) 2015 SP1, Storage Management Services (SMS) is installed by default.

- ♦ [Section 3.1, “System Requirements,” on page 15](#)
- ♦ [Section 3.2, “Updating Existing Installations,” on page 15](#)
- ♦ [Section 3.3, “Starting SMS Services,” on page 15](#)
- ♦ [Section 3.4, “Configuring SMDR,” on page 15](#)
- ♦ [Section 3.5, “Configuring the Target Service Agent for File System,” on page 17](#)

3.1 System Requirements

SMS runs on any system where OES can be installed. The minimum system requirements for SMS are the same as the requirements for an OES server as described in [“Preparing to Install OES 2015 SP1”](#) in the *OES 2015 SP1: Installation Guide*.

3.2 Updating Existing Installations

Existing OES servers can be upgraded to OES 2015 SP1. As a part of the upgrade, SMS is also upgraded.

3.3 Starting SMS Services

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 information, see [Section 3.5.2, “Using the Command Line,” on page 18](#).

NOTE: SMS is not registered with NCS, if NCS is configured after OES configuration. Ensure to either restart SMDR or unload and load TSA to register with NCS.

3.4 Configuring SMDR

- ♦ [Section 3.4.1, “Using iManager,” on page 16](#)
- ♦ [Section 3.4.2, “Using the Command Line,” on page 16](#)
- ♦ [Section 3.4.3, “Using a Configuration File,” on page 16](#)

3.4.1 Using iManager

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

NOTE: iManager standalone version running on Windows cannot be used to manage SMS plugin.

- 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).

On modifying the value for this option, 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.

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

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.

- ♦ **Name Resolution Through HOSTS File:** SMDR can be configured to use a hosts file (`/etc/hosts`) for IP address name resolution. The HOSTS file is automatically configured 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.

3.4.2 Using the Command Line

SMDRD is located at `/opt/novell/sms/bin` folder. SMDR can be configured using the command line options:

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

For information, see the `smdrd(8)` man page in OES server.

3.4.3 Using a Configuration File

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

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

3.5 Configuring the Target Service Agent for File System

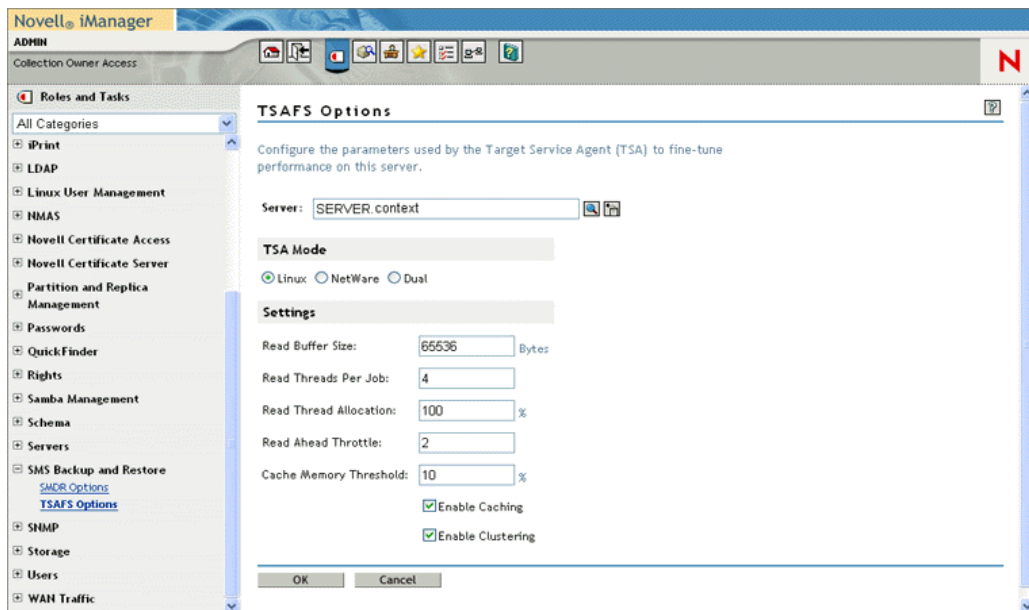
TSAFS provides configurable parameters to help tune its performance. It also provides configurable parameters to control certain specific features that are supported by it.

- ◆ Section 3.5.1, “Using iManager,” on page 17
- ◆ Section 3.5.2, “Using the Command Line,” on page 18
- ◆ Section 3.5.3, “Using a Configuration File,” on page 19

3.5.1 Using iManager

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

- 1 In iManager, click **SMS Backup and Restore**, then click **TSAFS Options**.



- 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 2015 SP1.

If the Dual Mode is selected, both NetWare and Linux semantics are simultaneously exposed and can be used independently of each other. By default, the TSA Mode is set to Linux.

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

- ◆ **Read Buffer Size:** 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.

The modified value takes effect when you reload the `tsafs` module.

- ◆ **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.

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

- ◆ **Read Thread Allocation:** The percentage of Read Threads Per Job that processes a data set.

Read Thread Allocation can be used to control the maximum number of read threads that may be allocated to process a single data set. By default, the Read Thread Allocation is set to 100 (%). It can have a value in the range 10 (%) to 100 (%). It is recommended that this value be set to 100 (%) if the backup application requests data sets serially.

- ◆ **Read Ahead Throttle:** The maximum number of data sets that the TSA processes simultaneously.

Read Ahead Throttle can be used to limit the number of simultaneous data sets that are being cached. In certain runtime scenarios, it helps in overriding the Read Thread Allocation in order to complete processing of large data sets. By default, the Read Ahead Throttle is set to 2. It can have a value in the range 1 to 32.

- ◆ **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.

The modified value takes effect when you reload the `tsafs` module.

- ◆ **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.

The modified value takes effect when you reload the `tsafs` module.

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

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

The modified value takes effect when you reload the `tsafs` module.

3.5.2 Using the Command Line

The Storage Management Data Requestor (SMDR) is a daemon process 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.

The `/opt/novell/sms/bin/smsconfig` command can be used to load or unload the (TSA) with the SMDR daemon:

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

For example:

<code>smsconfig -r, --refresh</code>	Specify this option in Cluster Unload scripts to inform SMDRD regarding the cluster resource failover. This option must be passed with <code>-p <pool_servername></code> . This option unregisters the pool and is supported with TSAFS.
<code>smsconfig -l tsafs</code>	Registers the TSA, <code>tsafs</code> with SMDR
<code>smsconfig -u tsafs</code>	De-registers the TSA, <code>tsafs</code> 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. This is the default.
	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 independently of each other.

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.

```
smsconfig -r tsafs -p ClusterPool_Server
```

On a cluster resource failover, the specified pool is unregistered.

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

For advanced configuration, see [Section 5.3, “Fine-Tuning SMS Performance,” on page 33](#).

To enable additional backup features, see [Section 4.1.4, “Additional Backup Features,” on page 25](#).

3.5.3 Using a Configuration File

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

For information, see the `tsafs.conf(5)` man page.

4 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, see the vendor's documentation pertaining to the specific application.

The following topics are discussed in this section:

- ♦ [Section 4.1, "Backing Up Data Using SMS," on page 21](#)
- ♦ [Section 4.2, "Restoring Data Using SMS," on page 27](#)
- ♦ [Section 4.3, "Restoring NetWare Backups to OES 2015 SP1 File Systems," on page 29](#)

4.1 Backing Up Data Using SMS

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

- ♦ [Section 4.1.1, "Prerequisites," on page 21](#)
- ♦ [Section 4.1.2, "Backing Up the File Systems," on page 22](#)
- ♦ [Section 4.1.3, "Backing Up Clusters," on page 24](#)
- ♦ [Section 4.1.4, "Additional Backup Features," on page 25](#)

4.1.1 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, see Appendix A, File-Level Snapshot Commands in the *OES 2015 SP1: NSS File System Administration Guide for Linux*. 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 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:

- Users performing backup need to be LUM-enabled.
- Load the controller and storage device drivers on the backup server.
- Load the SMDR and TSAs on the backup and target server.

See [Section 3.3, “Starting SMS Services,” on page 15](#) for information on how to start SMS services.

4.1.2 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 22](#)).

TSAFS supports backing up:

- ◆ File system metadata such as name spaces, extended attributes, trustee rights, and data streams on OES servers.
- ◆ All POSIX* compliant file systems on ReiserFS, Ext2, Ext3, and XFS file systems OES servers, see [Appendix C, “POSIX File System Support,” on page 65](#).
- ◆ NSS file system and associated metadata on OES which are not available through POSIX interfaces.

TSAFS uses the ECMA SIDF standard format to store the file system data. For information, see 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 23](#)
- ♦ [“Backing Up Links” on page 23](#)
- ♦ [“Backing Up NCP Volumes” on page 23](#)
- ♦ [“Backing Up Dynamic Storage Technology Volumes” on page 24](#)

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.

NOTE: SMS now supports backup of Active Directory trustees on an AD-enabled NSS volumes.

For additional information about object ID and trustee issues, see [“Restoring Trustee or Owner Assignments” on page 28](#).

Backing Up Links

TSAFS supports backup of hard and soft links when backing up POSIX compliant file systems on OES servers. It supports backup of hard links on the NSS file system on OES servers.

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

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

TSAFS supports backup and restore of additional metadata for files and directories under NCP volumes. When backing up an local NCP volume file or directory, trustee assignments and inherited rights filters for the data set is additionally backed up using the Novell client libraries. This is not supported for clustered NCP volumes.

For information on setting up NCP volumes and the NCP metadata on them, see [“Managing NCP Volumes”](#) and [“Managing File System Trustees, Trustee Rights, and Attributes on NCP Volumes”](#) in the *OES 2015 SP1 Beta: NCP Server for Linux Administration Guide*.

Backing Up Dynamic Storage Technology Volumes

NCP server for Linux allows administrators to create shadow volume pairs by using two NSS volumes as the primary volume and the secondary volume. The metadata on each volume is the same as for independent NSS volumes. When the DST volume is mounted, the secondary volume is mounted in Linux but is not mounted in NCP. This allows the secondary volume to be visible to backup software even though users do not directly access the files.

You must separately back up the primary volume and the secondary volume to back up the files and directories on them. Backup tools can apply one backup policy to the primary file tree and a different backup policy to the secondary file tree. The only operations that take place on the secondary volume are backup, or remove and archive.

For information about backing up the primary NSS volume and secondary NSS volume in a DST volume, see [“Using Backup Utilities with DST Shadow Volume Pairs”](#) and [“Backing Up DST Shadow Volumes”](#) in the *OES 2015 SP1: Dynamic Storage Technology Administration Guide*.

4.1.3 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 cluster migration of individually managed cluster resources. For more information, see the *OES 2015 SP1: Novell Cluster Services for Linux Administration Guide*.

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, see [Section 3.5, “Configuring the Target Service Agent for File System,” on page 17](#). This will enable listing of all cluster volumes as part of the cluster node instead of virtual server resource.

Backing Up Mixed Node Clusters

During a rolling cluster conversion from NetWare to OES, it is possible to have mixed node clusters, where different nodes in the cluster run NetWare or OES. For more information regarding mixed node clusters, see the *OES 2015 SP1: Novell Cluster Services NetWare to Linux Conversion Guide*.

TSAFS supports mixed node cluster backup. Path names are represented differently on NetWare and OES servers. In order to achieve a consistent backup, TSAFS supports a NetWare emulation mode on an OES 2015 SP1 server. This mode is used to make TSAFS behave as a NetWare target on OES 2015 SP1 server. This resolves any path name conflicts that might arise because of mixed node clusters in the setup. To enable this feature, see the [“NetWare Emulation Mode” on page 26](#).

4.1.4 Additional Backup Features

This section describes additional features supported by TSAFS:

- ♦ [“Apparmor Profile for SMDR daemon” on page 25](#)
- ♦ [“GroupWise Backup” on page 25](#)
- ♦ [“Non-Caching Mode of Operation” on page 25](#)
- ♦ [“Code Page Support” on page 26](#)
- ♦ [“NetWare Emulation Mode” on page 26](#)
- ♦ [“Autoloading the TSAFS Settings” on page 26](#)

Apparmor Profile for SMDR daemon

The default apparmor profile `opt.novell.sms.bin.smdrd` is available in `/etc/apparmor/profiles/extras/` folder. The profile contains all permissions to the paths and libraries and permissions that `smdr` requires during its execution. The profile contains the `rw` permissions to the file system `root (/)` to enable the backup of any path on the file system. You can modify the profile as per your security requirements. On modifying the profile, reload Apparmor with the `rcapparmor` command.

GroupWise Backup

TSAFS supports backing up GroupWise database files. TSAFS is integrated with GroupWise 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:

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

To autoload the TSAFS settings, perform the steps mentioned in the section [“Autoloading the TSAFS Settings” on page 26](#).

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:

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

To autoload the TSAFS settings, perform the steps mentioned in the section [“Autoloading the TSAFS Settings” on page 26](#).

Code Page Support

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 information on codesets, see the `tsafs(1)` man page.

To autoload the TSAFS settings, perform the steps mentioned in the section [“Autoloading the TSAFS Settings” on page 26](#).

NetWare Emulation Mode

TSAFS by default exposes the 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 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 targets. In `netware` mode, the TSA displays only NetWare File System targets. 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.

To autoload the TSAFS settings on OES Linux, perform the steps mentioned in the section [“Autoloading the TSAFS Settings” on page 26](#).

Autoloading the TSAFS Settings

To make the TSAFS settings persistent, perform the following steps:

- 1 Go to the file `/etc/opt/novell/sms/smdrd.conf`.
- 2 Modify the line `autoload: tsafs` to read as `autoload: tsafs --filename`
For example, Modify the line `autoload: tsafs` to `autoload: tsafs --EnableGW --noCachingMode`.
- 3 Restart `novell-smdrd` using `rcnovell-smdrd restart` command. This command shuts down and restarts the `smdrd` daemon.
- 4 Wait for few seconds and then issue `smsconfig -t` command. The following is displayed:
The loaded TSAs are, `tsafs --filename`. For example, `autoload: tsafs --EnableGW --noCachingMode`.

4.2 Restoring Data Using SMS

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

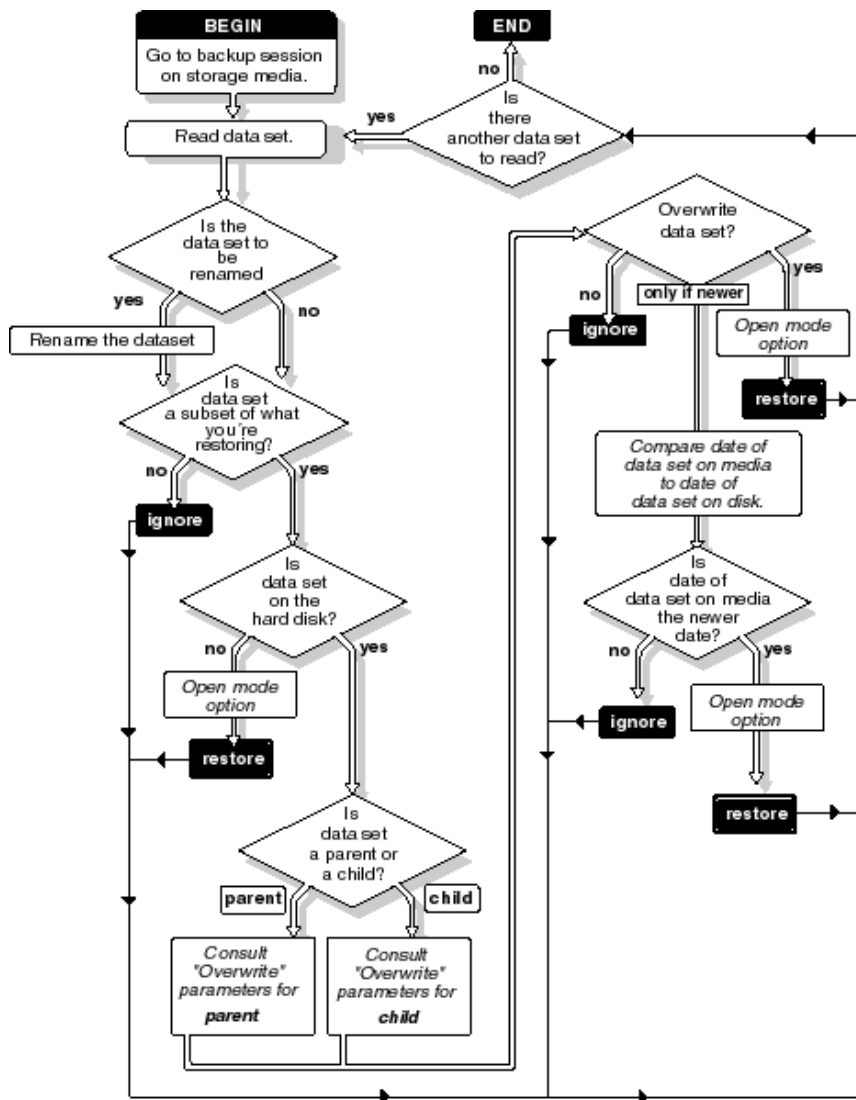
- ◆ [Section 4.2.1, “How SMS Restores Data,” on page 27](#)
- ◆ [Section 4.2.2, “Restoring Trustee or Owner Assignments,” on page 28](#)
- ◆ [Section 4.2.3, “Limitations for Restoring Data Using SMS,” on page 29](#)

4.2.1 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?

NOTE: When machine is running, system libraries cannot be restored because smdr uses dynamically loaded libraries from `/lib` folder for restoration.



4.2.2 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.

SMS now restores Active Directory trustees that were backed up. Ensure that the NSS volumes are AD enabled to restore the AD trustees.

If the volumes are not AD-enabled, then the files are successfully restored but the owner, user quota, and trustee information of the Active Directory user is not restored.

4.2.3 Limitations for Restoring Data Using SMS

- ♦ [“Restoring an NCP POSIX Volume” on page 29](#)
- ♦ [“Open Files” on page 29](#)

Restoring an NCP POSIX Volume

On restoring an NCP POSIX volume, the NSS user must be LUM-enabled to preserve the user’s ID.

Open Files

Restoring an open file is not supported.

Restoring a file fails when a file with same name and extension already exists on the target and is used by some process or is in open state.

4.3 Restoring NetWare Backups to OES 2015 SP1 File Systems

This section describes steps to move data transparently between NetWare and OES 2015 SP1 file systems:

- ♦ [Section 4.3.1, “Prerequisites,” on page 29](#)
- ♦ [Section 4.3.2, “Restoring NetWare Backups to NSS,” on page 30](#)
- ♦ [Section 4.3.3, “Restoring NetWare Backups to Non-NSS File Systems,” on page 30](#)

4.3.1 Prerequisites

- ♦ [“Linux Mode” on page 29](#)
- ♦ [“NetWare Emulation Mode” on page 29](#)
- ♦ [“NFS Name Space Support on OES” on page 29](#)

Linux Mode

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

NetWare Emulation Mode

If the NetWare emulation mode (see [“NetWare Emulation Mode” on page 26](#)) is used, then TSAFS on OES can be treated like a NetWare backup or restore.

NFS Name Space Support on OES

OES 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.

4.3.2 Restoring NetWare Backups to NSS

When restoring data that was backed up from a NetWare source to the NSS file system on OES server, 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 server.

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

4.3.3 Restoring NetWare Backups to Non-NSS File Systems

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, see *Novell Client for Linux Installation and Administration Guide* (http://www.novell.com/documentation/linux_client/index.html)

5 Optimizing SMS

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

- ◆ [Section 5.1, “Prerequisites,” on page 31](#)
- ◆ [Section 5.2, “Troubleshooting Backup Performance,” on page 32](#)
- ◆ [Section 5.3, “Fine-Tuning SMS Performance,” on page 33](#)

5.1 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.

5.1.1 Storage Subsystem

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

If backup is critical and a non-time 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, see the respective hardware reference guides.

File System (NSS) and Applications

- ◆ The file system performance tuning and networking parameters should be configured for improved performance. See 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 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.

See the “Managing Software RAID Devices” section in the [OES 2015 SP1: NSS File System Administration Guide for Linux](#) for a detailed discussion on NSS tuning parameters.

SMS

SMS can be configured to optimally exploit the underlying subsystem capability by fine tuning its working parameters. For more information, see [Section 5.3, “Fine-Tuning SMS Performance,” on page 33](#).

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, see 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, see the appropriate vendor documentation.

5.2 Troubleshooting Backup Performance

This section provides troubleshooting information that you can use to optimize the Backup Performance.

5.2.1 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 information on improving storage subsystem performance, see [“Storage Subsystem” on page 31](#).
- ♦ 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, see [Section 5.1, “Prerequisites,” on page 31](#).

Monitor the disk throughput to determine the efficiency of the disk system to provide data for backup. As backup is an disk I/O bound problem, this 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, see [Section 5.3, “Fine-Tuning SMS Performance,” on page 33](#).

For more information on TSATEST, see the [TSATEST Readme for Linux \(https://www.novell.com/documentation/developer/samplecode/smscomp_linux_sample/tsatest/tsatest_linux.html\)](https://www.novell.com/documentation/developer/samplecode/smscomp_linux_sample/tsatest/tsatest_linux.html).

5.3 Fine-Tuning SMS Performance

TSATEST is used to determine bottlenecks in the backup systems. For more information, see [Section 5.2, “Troubleshooting Backup Performance,” on page 32](#). Using this information, the following tunable switches can be used to identify sweet spots that helps improve throughput of your backup systems.

5.3.1 Basic Configuration

Configure the following basic tunable parameters to enhance the SMS performance. For more information about configuring the switches, see [Section 3.5, “Configuring the Target Service Agent for File System,” on page 17](#).

Table 5-1 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	<code>TSAFS / readthreadsperjob=value</code>

Task	Purpose	Field Name in the iManager Interface	Command
Set the read buffer size	<p>This is the number of data bytes read from the file system by a single read operation. This switch is based on the buffer sizes requested by the engine. For example, if the engine requests 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	<pre>TSAFS / readbuffersize= value</pre>
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	<pre>TSAFS / cachememorythreshold= value</pre>

Task	Purpose	Field Name in the iManager Interface	Command
Enable or disable caching based on engine usage and workload being backed up	<p>This option is used to specify if TSA should do predictive caching during backups. Caching improves backup performance, on certain workloads, by prefetching files in memory.</p> <p>The default value is <code>cachingMode</code>.</p> <p>If the datasets are not requested in the order in which they were prefetched, backup performance may degrade for some engines and for certain workloads.</p> <p>To determine if caching will improve the backup performance, enable caching and load TSA with the following TSA debug options: <code>smsdebug=800003c</code> and <code>smsdebug2=fffff100</code></p> <p>The TSA debug log file displays the number of datasets opened by the engine and the TSA. If the difference in both the values is significant (>50%), then you are recommended to disable caching for optimal performance.</p> <p>For information on enabling debug logging, see Appendix B, "Creating SMS Debug logs," on page 61</p>	Enable Caching	<code>TSAFS /CachingMode noCachingMode</code>

5.3.2 Advanced Configuration

Configure the following advanced tunable parameters to enhance the SMS performance. For more information about configuring the switches, see [Section 3.5, "Configuring the Target Service Agent for File System,"](#) on [page 17](#).

Table 5-2 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

5.3.3 Additional Parameters

Configure the following additional tunable parameters to enhance the SMS performance.

Table 5-3 Additional Tunable Parameters to Enhance SMS Performance

Task	Purpose	Parameter	Command
Changing the I/O scheduler to deadline scheduler	<p>The Deadline scheduler sets a cap on per request latency and ensures good disk throughput.</p> <p>Service queues are prioritized by deadline expiration, making this a good choice for real-time applications, databases and other disk-intensive applications.</p> <p>In a multi-path environment change the I/O scheduler on the multipath DM object.</p> <p>The I/O scheduler setting has to be modified per device and per server.</p>	deadline	<pre>echo deadline > /sys/block/{DEVICE-NAME}/queue/scheduler</pre> <p>This not persistent, rebooting the server sets the value to default.</p>
Modifying the slice_idle parameters of the CFQ scheduler	<p>Completely Fair Queuing (CFQ) I/O scheduler provides a good compromise between throughput and latency by treating all competing processes. Each process is given a separate request queue and a dedicated time slice of disk access.</p> <p>When a task has no more I/O to submit in its time slice, the I/O scheduler waits for a while before scheduling the next thread to improve locality of I/O.</p> <p>In a multi-path environment, disable idling on the CFQ queues on the multipath DM object.</p> <p>The slice_idle setting has to be modified per device and per server.</p>	slice_idle	<pre>echo 0 > /sys/block/device/queue/iosched/slice_idle</pre> <p>This not persistent, rebooting the server sets the value to default.</p> <p>The default value is eight millisecs, set it to zero to improve the backup performance.</p>

6 Running Storage Management Services in a Virtualized Environment

SMS runs in a virtualized environment just as it does on a physical server running Open Enterprise Server (OES) 2015 SP1, and requires no special configuration or other changes.

To get started with Xen virtualization, see [Virtualization with Xen documentation](#).

To get started with KVM virtualization, see [Virtualisation with KVM virtualisation](#).

To get started with third-party virtualization platforms, such as Hyper-V from Microsoft and the different VMware product offerings, refer to the documentation for the product that you are using.

For information on setting up virtualized OES 2015 SP1, see “[Installing, Upgrading, or Updating OES on a VM](#)” in the *OES 2015 SP1: Installation Guide*.

7 SMS Security for SMDR

SMS provides remote backup and restore services using the Storage Management Data Requester (SMDR). This gives you the flexibility to use SMS from a single server to back up the entire network. This section details the security features available in SMS to secure your backup networks.

The SSL protocol is used to secure remote backup connections that are created by SMDR. For more information on SSL/TLS, see [RFC 2246 \(http://www.ietf.org/rfc/rfc2246.txt\)](http://www.ietf.org/rfc/rfc2246.txt).

The following topics are discussed in this section:

- ♦ [Section 7.1, “Using SSL in SMDR,” on page 41](#)
- ♦ [Section 7.2, “SMDR as a Client and Server,” on page 42](#)
- ♦ [Section 7.3, “Configuration Options,” on page 43](#)

7.1 Using SSL in SMDR

SMDR is enabled to use SSL and uses Novell TLS (NTLS) library for SSL operations.

When SMDR is configured to use SSL, it encrypts the channel by securing data that is exchanged between servers. For configuration options in SMDR, see [Section 7.3, “Configuration Options,” on page 43](#).

SMDR provides you the option of using SSL with or without certificates.

- ♦ [Section 7.1.1, “Using SSL without Certificates,” on page 41](#)
- ♦ [Section 7.1.2, “Using SSL with Certificates,” on page 41](#)

7.1.1 Using SSL without Certificates

When using SSL without certificates, SMDR uses cipher suites based on the anonymous Diffie-Hellman protocol to exchange session keys. This mechanism provides session security because the data is encrypted across the connections. However, this does not provide server identity authentication because certificates are not used to validate server identity.

7.1.2 Using SSL with Certificates

When SMDR is configured to use SSL with certificates, it has the ability to authenticate the server identity and secure data on the network. The server’s certificate that is exchanged during the SSL channel establishment provides server authentication.

Certificate Types

SMDR supports PEM (Privacy Enhanced Mail) encoded or DER (Distinguished Encoding Rules) encoded certificates. Certificates in other formats must be converted to either PEM or DER format in order to be used with SMDR.

PEM and DER are used by openssl to represent public and private keys and signatures for X.509 compliant certificates. The DER format is a block of base64 encoded data for a digital certificate. The PEM format is the DER format encoded with additional header and footer lines.

7.1.3 Password-Encrypted Private Key Files

Private keys stored on servers are typically encrypted using passwords. SMDR supports SSL private key files that are encrypted using this method.

If the private key file is password-encrypted, SMDR loads and displays a screen to accept the password. Enter the password at the prompt to continue loading SMDR.


7.2 SMDR as a Client and Server

SMDR is capable of connecting and backing up remote systems, which makes it behave both as a client and as a server. You need to select appropriate options to configure SMDRs for the required SSL behavior.

When SMDR connects to remote servers and communicates with loaded TSAs on those servers, it acts as a client. In turn, the SMDR module running on those remote servers act as server, servicing the client SMDR requests.

In [Figure 7-1](#), SMDR running on Server A behaves as a client that services requests for backup application and for TSAFS running on the remote Server B. SMDR on Server B behaves as a server to SMDR running on Server A.

Figure 7-1 SMDR as a Client and Server



7.3 Configuration Options

To configure SMDR to use SSL with certificates, update the SMDR configuration file with the SSL options in this section.

- ◆ [Section 7.3.1, “Server Certificate Options,” on page 43](#)
- ◆ [Section 7.3.2, “Client Certificate Options,” on page 43](#)
- ◆ [Section 7.3.3, “Miscellaneous Options,” on page 44](#)
- ◆ [Section 7.3.4, “SSL Option Considerations,” on page 44](#)

7.3.1 Server Certificate Options

The options specified below, enables you to configure server-side SMDR to use certificate-based SSL protocol. See [Section 7.2, “SMDR as a Client and Server,” on page 42](#) for details on SMDR’s behavior as a server.

Table 7-1 Server Certificate Options

Options	Description
PublicKey	Path of the server’s public key certificate file. By default, this is disabled.
PublicKeyType	The format of the server’s <code>publickey</code> certificate file. This option accepts either PEM or DER. By default, the value is PEM.
PrivateKey	Path of the server’s private key certificate file. By default, this is disabled.
PrivateKeyType	The format of the server’s <code>privatekey</code> certificate file. This option accepts either PEM or DER. By default, the value is PEM.

7.3.2 Client Certificate Options

The following options are used to enable client-side SMDR to use the certificate-based SSL protocol. See [Section 7.2, “SMDR as a Client and Server,” on page 42](#) for details on SMDR behavior as a client

Table 7-2 Client Certificate Options

Options	Description
TrustedCertificate	Path of the trusted CA certificate. By default, this is disabled.
TrustedCertificateType	The format of the trusted CA certificate. This option accepts either PEM or DER. By default, the value is PEM.

7.3.3 Miscellaneous Options

Table 7-3 Miscellaneous Options

Options	Description
LegacyConnections	Specifies if connections can be established with older SMDRs that do not support SSL. This is a Boolean switch and can be configured as enable or disable. The default value is enable. This is applicable only when SMDR behaves as a client. See Section 7.2, “SMDR as a Client and Server,” on page 42 for more information.
DataEncryption	Specifies if the data needs to be encrypted or not. This can be configured as optional or mandatory. If the host server and target server are configured as optional, the data is not encrypted and only authentication information is encrypted. If either the client or the server is configured as mandatory then the data is also encrypted.

NOTE: Because performance is critical during backup, you can optionally configure SMDR to use SSL to only encrypt sensitive authentication information instead of all communications between SMDRs. To do this, disable the DataEncryption option.

7.3.4 SSL Option Considerations

When configured with some SSL options, SMDR can result in connection failures to other SMDRs on the network. The following information details how different options impact the connection behavior in SMDR.

[Table 7-4](#) lists client SMDR configuration options that force client SMDR to establish communication to only SSL-enabled SMDRs on the network. Attempts to connect to SMDRs that are not configured to use SSL result in connection failures.

Table 7-4 SSL Interoperability between Client and Server SMDRs

Client SMDR Options	Server SMDR Configuration	Connection Status
TrustedRootCertificate: <path> LegacyConnections: disable	PublicKey and PrivateKey	Pass
TrustedRootCertificate: <path> LegacyConnections: disable	PublicKey and PrivateKey configured	Fail
LegacyConnections: disable	SSL enabled SMDR	Pass
LegacyConnections: disable	Legacy SMDR	Fail

8 Coexistence and Migration Issues

One of the top priorities in designing Open Enterprise Server (OES) 2015 SP1 was to ensure that new OES components, 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 2015 SP1 platform.

This section discusses the issues involved in the coexistence and migration of SMS in OES. It is divided into the following sections:

- ♦ [Section 8.1, “Coexistence,” on page 45](#)
- ♦ [Section 8.2, “Migration,” on page 47](#)

For a general discussion of coexistence and migration issues in OES, see the [Migration and Coexistence Web site \(http://www.novell.com/documentation/oes2/migrate-consolidate-coexist.html#migrate-consolidate-coexist\)](http://www.novell.com/documentation/oes2/migrate-consolidate-coexist.html#migrate-consolidate-coexist).

8.1 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](#)

8.1.1 Compatibility

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

SMDR

- ♦ [“Wire Compatibility of the SMDR Protocol” on page 45](#)
- ♦ [“Discovery Protocols Used in SMDR” on page 46](#)

Wire Compatibility of the SMDR Protocol

The SMDR (see [Section 1.2.1, “Storage Management Data Requester,” on page 10](#)) 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. 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. SAP is not supported for SMDR on OES 2015 SP1. See [“OES 2015 SP1 on IPX-Based Networks” on page 47](#) for more information on how compatibility for SAP can be achieved.

Using iManager

You can use the latest SMS plug-in to iManager running on an OES 2015 SP1 server to configure SMS services.

TSAFS

- ◆ [“Data Stream Compatibility” on page 46](#)

Data Stream Compatibility

TSAFS (see [“File System TSA \(TSAFS\)” on page 11](#)) 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.

TSAFS provides full data stream compatibility between NetWare and OES 2015 SP1. 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 2015 SP1 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 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 traditional file system data or NSS file system data to non-NSS file systems on OES 2015 SP1:

- ◆ 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

8.1.2 Coexistence Issues

- ◆ [Backup Application Support for OES File System Backup](#)
- ◆ [OES 2015 SP1 on IPX-Based Networks](#)

Backup Application Support for OES 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 2015 SP1. For more information, refer to the commercial vendor's backup application documentation.

SMS also supports a NetWare emulation mode (see [“NetWare Emulation Mode” on page 26](#)) 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 the future (After all backup applications have moved to backing up OES 2 or later natively), data backed up using this option would be recoverable by all future TSAs. nbackup (see [Section 1.3, “Backup Applications,” on page 11](#)) supports backing up of OES 2015 SP1. However, these backup applications are technology demonstrators and are not positioned as enterprise backup applications.

OES 2015 SP1 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 2015 SP1, so in a pure IPX environment, SMDR on OES 2015 SP1 cannot discover or resolve SMDRs on NetWare and vice versa. For SMS services on OES to work independently of platforms, discovery and name resolution protocols that are supported by SMDR must be common across all the platforms.

For more information on protocols supported by SMDR, see [Section 3.4, “Configuring SMDR,” on page 15](#).

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 during a rolling cluster conversion. For more information regarding mixed node clusters, see the [OES 2015 SP1: Novell Cluster Services NetWare to Linux Conversion Guide](#).

TSAFS supports backing up of NSS file system resources on nodes in a mixed node cluster with failover/failback support. To back up cluster resources in a mixed node environment, use the TSAFS on OES 2015 SP1 in the emulation mode of operation. See [“NetWare Emulation Mode” on page 26](#) for more information on how to use the emulation mode.

8.2 Migration

As SMS does not store any data on the server apart from its configuration files. You 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 environment. Any configuration updates must be done using iManager. For information, see [Section 3.4, “Configuring SMDR,” on page 15](#) and [Section 3.5, “Configuring the Target Service Agent for File System,” on page 17](#).

9 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 9.1, “Startup and Connection Issues,” on page 49](#)
- ♦ [Section 9.2, “Common Backup and Restore Issues,” on page 50](#)
- ♦ [Section 9.3, “Backup and Restore Issues,” on page 52](#)
- ♦ [Section 9.4, “Cluster Related Issues,” on page 53](#)

9.1 Startup and Connection Issues

- ♦ [“novell-smdrd fails to start with the default AppArmor profile” on page 49](#)
- ♦ [“Unable to connect to target server or service” on page 49](#)
- ♦ [“Unable to modify, SMDR listener IP address using iManager” on page 49](#)

novell-smdrd fails to start with the default AppArmor profile

Possible Cause: Novell AppArmor restricts the access to novell-smdrd.

Action: Using YaST, open Novell AppArmor > Update Profile Wizard and update novell-smdrd profile to grant the required access.

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.

9.2 Common Backup and Restore Issues

- ♦ [“Backup does not include the modifications” on page 50](#)
- ♦ [“Backup is slow” on page 50](#)
- ♦ [“Files were restored but the error file contains a message specifying the name space formats not restored” on page 50](#)
- ♦ [“Restore is slow” on page 51](#)
- ♦ [“Unable to backup open files in NSS Volumes” on page 51](#)
- ♦ [“Unable to restore compressed files” on page 51](#)
- ♦ [“Unable to set the data set name” on page 51](#)
- ♦ [“Out of disk space” on page 51](#)
- ♦ [“DOS namespace based restores fail for certain files” on page 52](#)
- ♦ [“Backed up filenames display characters of the form \[xxxx\]” on page 52](#)

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.

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 21](#).

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

Action: See [“Backing Up Migrated Files” on page 22](#).

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 21](#)) 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 21](#).

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 21](#).

Unable to restore compressed files

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

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.

9.3 Backup and Restore Issues

- ♦ [“Restore Fails When TSAFS is Set to Netware Mode on a OES Server” on page 52](#)
- ♦ [“During full system restore, smdr crashes and restore fails” on page 52](#)
- ♦ [“Certain files do not get backed up on non-NSS file systems” on page 52](#)
- ♦ [“Backup or restore hangs on submitting a request” on page 53](#)

Restore Fails When TSAFS is Set to Netware Mode on a OES Server

Action: To resolve the issue, set the parameter `tsamode` to `Linux` in the `/etc/opt/novell/sms/tsafs.conf` file.

During full system restore, smdr crashes and restore fails

Possible Cause: SMDR uses dynamically loaded libraries from `/lib` folder for restoration. On a running machine, during full restore the entire file system is reloaded this causes SMDR to crash and fails to restore system libraries.

Action: During full system restore, you must restore the system libraries to a non-default path and not to the `/lib` folder. At a later time, move the restored libraries to `/lib` folder when file system is not in use.

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 page 26](#).

Backup or restore hangs on submitting a request

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.

9.4 Cluster Related Issues

- ♦ [“smdr.novell Is Not Registered with SLP for a New Cluster Resource” on page 53](#)
- ♦ [“On the failover of a cluster resource, smdr daemon terminates” on page 53](#)
- ♦ [“Backup fails in a mixed node cluster environment” on page 53](#)
- ♦ [“Clustered volumes are not backed up during full server backups” on page 54](#)
- ♦ [“Cluster pools are not listed on Linux” on page 54](#)
- ♦ [“Reconnect to TSAFS fails when a cluster resource migrates to another cluster node in a mixed mode cluster.” on page 54](#)

smdr.novell Is Not Registered with SLP for a New Cluster Resource

Explanation: You might get an error after creating a cluster resource indicating that smdr.novell is not registered with SLP for cluster resources, but the smdr.novell service for the node is registered.

```
Error: "cluster--<212>: Read ResVol error -603"
```

Action: The first time a cluster resource is created, smdrd cannot figure it out. Restart smdrd. Thereafter, smdrd is aware of the cluster resource. and advertise it correctly.

- 1 Log in to the server as the root user, open a terminal console, then enter

On the failover of a cluster resource, smdr daemon terminates

Possible Cause: Cluster failover scripts sends SIGTERM command to smdr to release the cluster resources.

Action: SMDR needs time to clear the existing connections and release the cluster volumes. Increase the sleep time in the cluster failover script to facilitate the release of the cluster volume.

```
novell-smdrd restart
```

Backup fails in a mixed node cluster environment

Explanation: In a mixed node cluster environment, backup fails when node fail-over and fall-back happens between NetWare and Linux servers.

Possible Cause: Volume names are in lowercase.

Action: On backing up data in a mixed node cluster environment, the volume name must be passed in uppercase. Linux filenames are case sensitive.

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 3.5, "Configuring the Target Service Agent for File System,"](#) on page 17.

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 3.5, "Configuring the Target Service Agent for File System,"](#) on page 17.

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 49.

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.

A 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 55](#)

A.1 TSA Options

- ♦ [Section A.1.1, “Backup Options,” on page 55](#)
- ♦ [Section A.1.2, “Restore Options,” on page 58](#)

A.1.1 Backup Options

All backup types contain advanced options to allow you to customize your backup. 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. You can specify how to scan what you are backing up. These options allow you to perform the following tasks.

- ♦ [“Choosing Subsets of Data to Back Up” on page 55](#)
- ♦ [“Scanning Data Sets” on page 57](#)

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 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. A list of existing mount points is maintained by SMS, by default these mount points are excluded.

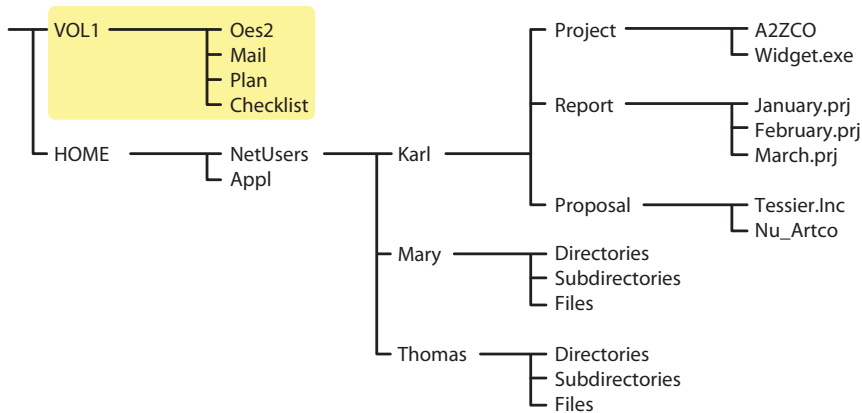
Include

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.

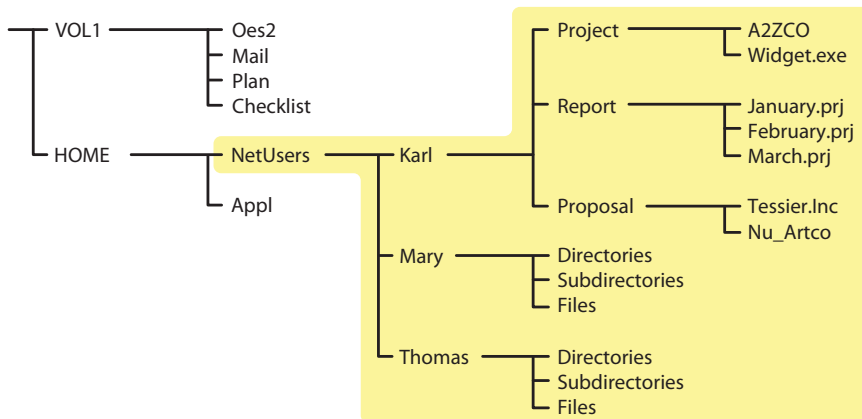
In the figure shown below, volume VOL1 is selected as an **Include** option. All other areas of the file system structure are excluded from the backup. You can exclude some subdirectories or files beneath your selection if necessary.

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.

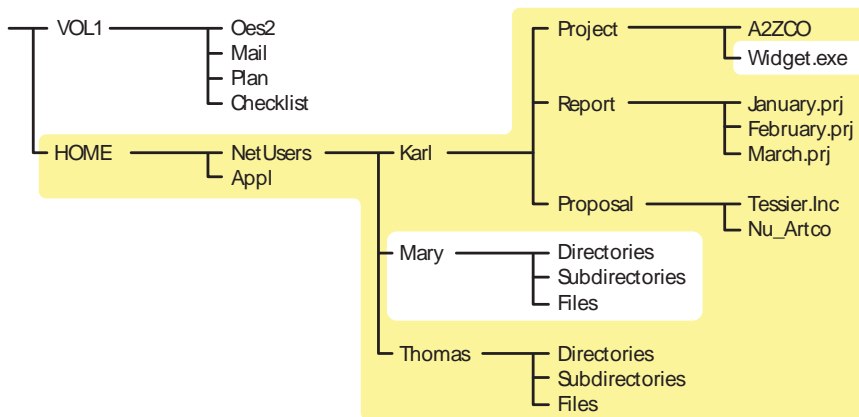
For example, the following command sequence results in volume home being included in the backup with the exception of the Mary directory and the widget.exe file.

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

```
Exclude directories (full path): /media/nss/HOME/netusers/mary
```

```
Exclude path/files /media/nss/HOME/netusers/karl/project/widget.exe
```


Figure A-3 Example combining nbackup Include and Exclude options



Scanning Data Sets

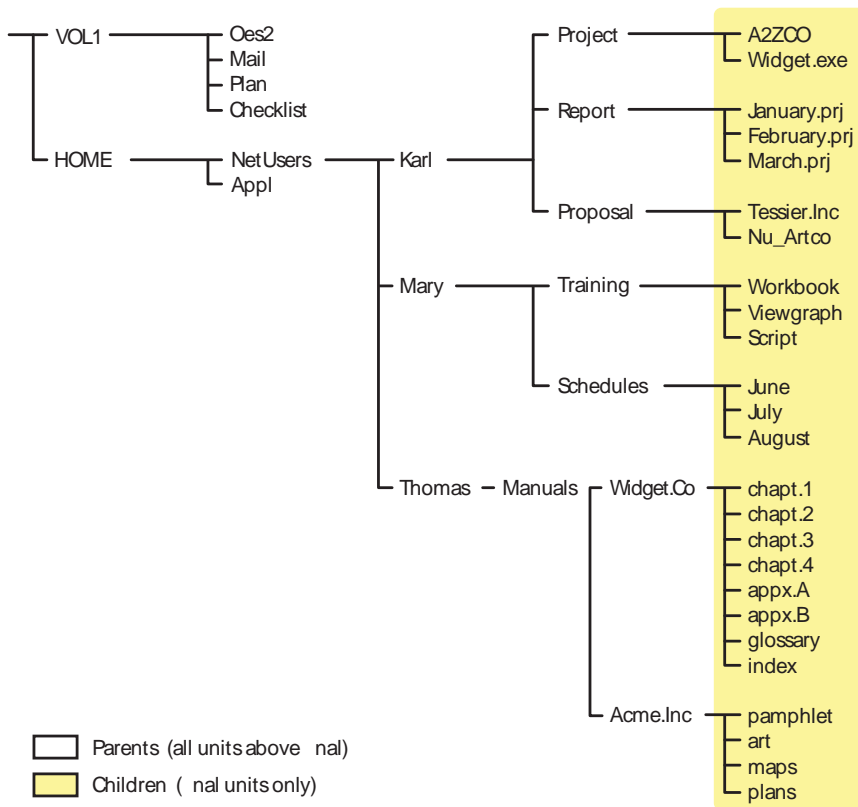
You can specify a different type of data set to be scanned.

A data set is a group of data that can be manipulated by nbackup. 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 nbackup, 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

nbackup 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 55](#).

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.

B 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.

SMS supports creation of debug logs for `smārd` and `tsafs`.

- ◆ [Section B.1, “Deciding Which Module to Enable for Debug Logging,” on page 61](#)
- ◆ [Section B.2, “Enabling Debug Logging,” on page 61](#)
- ◆ [Section B.3, “Location of the Debug Log,” on page 62](#)

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.

Use the following options to enable or disable debug logs:

- ◆ [Section B.2.1, “SMDR,” on page 62](#)
- ◆ [Section B.2.2, “TSAFS,” on page 62](#)

B.2.1 SMDR

- ♦ [“Enabling Debug Logging” on page 62](#)
- ♦ [“Disabling Debug Logging” on page 62](#)

Enabling Debug Logging

- 1 Kill the SMDR process

```
pkill smdrd
```

- 2 Enable debug and restart SMDR

```
/opt/novell/sms/bin/smdrd --smsdebug ffffffff --smsdebug2 ffffffff
```

- 3 Perform the tests to reproduce the problem.
- 4 The error messages will be logged in the debug log file.

Disabling Debug Logging

```
pkill smdrd
```

B.2.2 TSAFS

- ♦ [“Enabling Debug Logging” on page 62](#)
- ♦ [“Disabling Debug Logging” on page 62](#)

Enabling Debug Logging

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

Disabling Debug Logging

- ♦ Unload TSA

```
smsconfig -u tsafs
```

- ♦ Reload TSA

```
smsconfig -l tsafs
```

B.3 Location of the Debug Log

- ♦ [Section B.3.1, “Debug Log Location,” on page 63](#)
- ♦ [Section B.3.2, “Reducing the Debug Log Size,” on page 63](#)

B.3.1 Debug Log Location

By default, the debug files are created in `/var/opt/novell/log/sms/`. The debug file name for `smdrd` is `smdrd_debug_xxxx.log` where `xxxx` is the process id for `smdrd`. The debug file name for `tsafs` is `tsafs_debug_xxxx.log` where `xxxx` is the process id for `smdrd`.

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/`, enter the following command:

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

The above command creates a file `smdrd_xxxx.log` in `/home/testuser/` directory.

B.3.2 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.

C POSIX File System Support

This section provides information regarding the TSAFS support of POSIX-compliant file systems like BtrFS, ReiserFS, Ext2, Ext3, and XFS file systems on the Open Enterprise Server (OES) 2015 SP1. 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).

Backup of Linux POSIX file systems requires that ACLS and POSIX permissions be set on the Linux path for the LUM-enabled user performing the backup. The root user has all permissions needed to perform backup of any Linux path. You can use the Linux `chmod(8)` and `chown(8)` commands to give the backup user the Linux POSIX permissions to the directory being backed up.

The following table lists metadata that is backed up or restored to non-NSS file systems on OES 2015 SP1. 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
<code>st_mode</code>	Mode of the file, including File types and File access permissions
<code>st_nlink</code>	Number of hard links to the file
<code>st_uid</code>	User ID of the file
<code>st_gid</code>	Group ID of the file
<code>off_t</code>	Size of the file
<code>st_atime</code>	Time of last access
<code>st_mtime</code>	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`.

POSIX Compliant File Systems Limitations

- ◆ Backing up and restoring Extended Attributes is supported within the same file system, but restoration is not supported across different file systems.
- ◆ During restoration, a non-`root` user cannot overwrite the read-only files to the POSIX-complaint file systems, because write access is required for updating the files.

D SMSLS Utility

The smsls utility generates a list of files and directories based on the specified filter options. This utility stores the last executed time and EFL epoch in a user-specified configuration file and generates the file paths in the specified output file.

- ♦ [Section D.1, “Syntax,” on page 67](#)
- ♦ [Section D.2, “Options,” on page 67](#)
- ♦ [Section D.3, “Examples,” on page 68](#)

D.1 Syntax

```
smsls [options]
```

D.2 Options

Table D-1 SMSLS Options

Linux Options	Description
-a --list-archived	Generates a list of files whose archive bit is set.
-m--list-meta-archived	Generates a list of files whose metadata archive bit is set.
-o=FILE_PATH --output=FILE_PATH	Specifies the path to the file that stores the deleted file list. A new FILE_PATH.deleted file is created to store the deleted file list.
--nodirs	Includes only the files in the modified file list, not the directories. By default, directories are included in the modified file list.
-U --user	Specifies the username for connecting to the Target Service Agent (TSA). This is a mandatory option.
-P --password	Specifies the password to connect to the TSA. Specify the password by using the environment variable SMSLS_USER_PASSWORD. If no password is specified, smsls prompts for a password.
--path	Specifies the volume name or the location to generate the modified file list and deleted file list. This is a mandatory option.
--conf-file=FILE_PATH	Specifies the location of the file for storing the configuration details. The smsls utility stores the last run time and the epoch details in this configuration file. This is a mandatory option.
--init	Initializes the smsls utility. This stores the EFL epoch and the current date and time in the specified configuration file.

Linux Options	Description
--list-epochs	Lists all the active epochs for the specified volume.
-h	Displays the help information for the smsls utility.
--help	

D.3 Examples

On performing a full backup, you must initialize the smsls tool. During initialization, the smsls utility stores the EFL epoch and the current date and time in the configuration file. For the subsequent executions, the same configuration file needs to be provided.

This file list is used by engines that take a list of paths as input for more granular incremental backups. You can use the smsls utility to view information about files and epochs..

- ♦ To initialize the smsls utility for the volume `VOL1` and to store the current configuration in the file `smsls_vol1.conf`, enter

```
smsls --init --conf-file=/home/smsls_vol1.conf --path=/media/nss/VOL1 -U admin
```

- ♦ To generate the list of modified files and deleted files by using the configuration file `sys:backup-data/smsls_vol1.conf`, enter

```
smsls -o=/home/fileslist_vol1 --conf-file=/home/smsls_vol1.conf --path=/media/nss/VOL1 -U admin
```

The modified files are listed in the `/home/filelist_vol1` file.

The deleted files are listed in the `/home/filelist_vol1.deleted` file.

- ♦ To generate the list of files within the directory for a specified volume, enter

```
smsls -o=/home/fileslist_vol1 --conf-file=/home/smsls_vol1.conf --path=/media/nss/VOL1/dir1/dir2 -U admin
```

- ♦ To use the archive bit and metadata archive bits to generate the modified file list, enter

```
smsls -am -o=/home/fileslist_vol1 --conf-file=/home/smsls_vol1.conf --path=/media/nss/VOL1 -U admin
```

- ♦ To list the current active epochs, enter

```
smsls --list-epochs --conf-file=/home/smsls_vol1.conf --path=/media/nss/VOL1 -U admin
```

IMPORTANT:

- ♦ You must store the output of a file in the directory.
 - ♦ If a filtering option is specified for the archive bit and metadata archive bit, the modified time option is not used for displaying the file list.
-