



Open Enterprise Server 2018 SP3

Dynamic Storage Technology Administration Guide

July 2021

Legal Notices

© Copyright 2017 - 2021 Micro Focus or one of its affiliates.

The only warranties for products and services of Micro Focus and its affiliates and licensors ("Micro Focus") are as may be set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Micro Focus shall not be liable for technical or editorial errors or omissions contained herein. The information contained herein is subject to change without notice.

Contents

About This Guide	11
1 Overview of Dynamic Storage Technology	13
1.1 Understanding Dynamic Storage Technology	14
1.1.1 Merged View of the File Tree	14
1.1.2 File Systems	16
1.1.3 File Access Protocols	16
1.1.4 Secure File Access	17
1.1.5 Local File Access for Backup and Archive	18
1.2 Benefits of Dynamic Storage Technology	18
1.2.1 Merged View File Access for End Users	18
1.2.2 Policy-Based Migration between Primary and Secondary Storage Areas	18
1.2.3 Faster and Smaller Backups of Important Data	19
1.2.4 Faster Disaster Recovery	19
1.2.5 More Efficient Use of Expensive Storage	19
1.2.6 Fast Storage for Active Data and Slower, Less Expensive Storage for Old Data	20
1.2.7 Moving Files from an Existing Secondary Volume	20
1.2.8 Access to the Secondary Storage Area without the Performance Penalties of HSM Solutions	20
1.3 Shadowing Scenarios	20
1.3.1 Existing Volume as Primary with an Empty Volume as Secondary	21
1.3.2 Empty Volume as Primary with an Existing Volume as Secondary	21
1.4 DST Policy Scenarios	22
1.4.1 Move Files Based on the Last Time Accessed or Modified	22
1.4.2 Move Files Based on File Size	23
1.4.3 Move Files Based on File Extensions	23
1.4.4 Move Selected Files from Primary to Secondary	23
1.5 DST Components	24
1.5.1 NCP Engine	24
1.5.2 Shadow Volume	24
1.5.3 ShadowFS	24
1.5.4 Policy Engine	24
1.6 Management Tools	25
1.7 What's Next	25
2 What's New or Changed in Dynamic Storage Technology	27
2.1 What's New or Changed in Dynamic Storage Technology (OES 2018 SP3)	27
2.2 What's New or Changed in Dynamic Storage Technology (OES 2018 SP1)	27
2.3 What's New or Changed in Dynamic Storage Technology (OES 2018-Update 1-Patch)	27
2.4 What's New or Changed in Dynamic Storage Technology (OES 2018)	27
3 Planning Your Dynamic Storage Technology Server Environment	29
3.1 Open Enterprise Server 2018 SP1	29
3.2 NCP Server and Dynamic Storage Technology	30

3.3	OES Storage Services	30
3.4	Directory Services	30
3.4.1	Active Directory Trustees	30
3.4.2	eDirectory 9.2.1 Trustees	30
3.5	NCP File Access	31
3.6	OES CIFS File Access	31
3.7	Novell Samba File Access with ShadowFS, FUSE, and LUM	31
3.8	Supported Native Linux File Access Protocols with ShadowFS, FUSE, and LUM	32
3.9	OES AFP File Access (Not Supported)	32
3.10	Linux User Management	32
3.11	OES Cluster Services for Linux	32
3.12	OES Remote Manager for Linux	32
3.13	iManager for Linux	33
3.14	SFCB and CIMOM	33
3.15	Other OES Services	34
4	Installing Dynamic Storage Technology	35
4.1	Installing DST on a New OES Server	35
4.2	Installing DST on an Existing OES Server	36
5	Using DST in a Virtual Environment	39
6	Management Tools for DST	41
6.1	Dynamic Storage Technology Plug-In for OES Remote Manager for Linux	41
6.1.1	Accessing OES Remote Manager	41
6.1.2	Starting, Stopping, or Restarting OES Remote Manager on Linux	42
6.1.3	Quick Reference for Dynamic Storage Technology Options	43
6.1.4	Quick Reference for NCP Server Options	45
6.1.5	Quick Reference for DST Global Policy Settings	46
6.1.6	Shadow Volume Inventory and Trustee Reports	46
6.2	NCP Console (NCPCON) Commands	47
6.3	Management Tools for NSS Volumes	47
6.3.1	OES File Access Rights Management (NFARM)	47
6.3.2	Storage Plug-In for iManager 3.2.1	47
6.3.3	Files and Folders Plug-In for iManager 3.2.1	47
6.3.4	NSS Management Utility (NSSMU)	48
6.3.5	Novell Linux Volume Manager (NLVM) Commands	48
6.4	Management Tools for Clustering	48
7	Configuring DST Global Policies	49
7.1	Replicating Branches of the Primary File Tree in the Secondary File Tree	49
7.2	Shifting Files from the Secondary File Tree to the Primary File Tree	50
7.2.1	Understanding Shift Parameters	50
7.2.2	Configuring a Global Policy for Shifting Modified Shadow Files	53
7.2.3	Configuring a Global Policy for Shifting Accessed Shadow Files	54
7.2.4	Configuring a Global Policy for the Days Since Last Access	54
7.2.5	Using the SET Command to Set Global Policies	54
7.3	Resolving Instances of Duplicate Files	55
7.3.1	Understanding Conflict Resolution for Duplicate Files	55

7.3.2	Configuring a Global Policy for Actions to Resolve Duplicate Files Conflicts	58
7.3.3	Enabling or Disabling Broadcast Messages for Duplicate Files Conflicts	58
7.3.4	Resolving Instances of Duplicate Files in the /_DUPLICATE_FILES Directory	59
7.4	Loading ShadowFS	60
7.4.1	Using OES Remote Manager to Set the Autostart of ShadowFS	60
7.4.2	Using the Command Line to Set the Autostart of ShadowFS	60
7.4.3	Manually Starting and Stopping ShadowFS	61
8	Managing Services for DST	63
8.1	Restarting the Novell NCP/NSS IPC (ncp2nss) Daemon	63
8.2	Restarting the eDirectory (ndsd) Daemon	63
8.3	Starting and Stopping ShadowFS	64
9	Planning for DST Shadow Volume Pairs and Policies	65
9.1	Storage Requirements for DST Volume Pairs	65
9.1.1	Storage Devices	66
9.1.2	iSCSI Block Storage Devices	66
9.1.3	File Systems	66
9.2	Data Access Requirements for a DST Shadow Volume Pair	67
9.2.1	Administrator Access	67
9.2.2	User Access and Authorization	68
9.2.3	File Access Protocols	68
9.2.4	ShadowFS and FUSE	71
9.3	Guidelines for Working with DST Shadow Volume Pairs	71
9.3.1	Number of Shadow Volumes per Server	71
9.3.2	Data Volumes	71
9.3.3	Files and Folders	72
9.3.4	File System Trustees and Rights	73
9.3.5	File System Management Utilities	73
9.4	Using NSS Volumes in DST Shadow Volume Pairs	73
9.4.1	DST Support for NSS Volume Attributes	74
9.4.2	DST Support for NSS Features and Actions	76
9.5	Using Trustees, Trustee Rights, and File System Attributes on DST Shadow Volume Pairs	77
9.6	Using NSS Encrypted Volumes in a DST Shadow Volume Pair	78
9.7	Using NSS Quotas on DST Shadow Volume Pairs	78
9.7.1	NSS Volume Quotas	78
9.7.2	NSS Directory Quotas	79
9.7.3	NSS User Quotas	79
9.8	Using Opportunistic File Locking on DST Shadow Volume Pairs	80
9.9	Using OES Cluster Services with DST Shadow Volume Pairs	80
9.10	Using Novell Distributed File Services with DST Shadow Volume Pairs	81
9.11	Using OES Storage Services Auditing Client Logger (VLOG) Utility with DST Shadow Volume Pairs	82
9.12	Using Virus Checking Utilities with DST Shadow Volume Pairs	82
9.13	Using Backup Utilities with DST Shadow Volume Pairs	82
10	Creating and Managing DST Shadow Volumes for NSS Volumes	85
10.1	Understanding DST Volume Pairs	85
10.1.1	Primary Volume	86
10.1.2	Secondary Volume	86

10.1.3	Merged View	86
10.1.4	How Directories Are Created in the Shadow Volume	86
10.1.5	Global DST Policies	86
10.1.6	Shadow Volume Policies	87
10.1.7	File Inventory for the Shadow Volume	87
10.1.8	Moving Specified Files between Volumes	87
10.2	Creating a DST Shadow Volume with NSS Volumes	88
10.2.1	Checking the Availability of an NSS Volume for Primary	88
10.2.2	Checking the Availability of an NSS Volume for Secondary	89
10.2.3	Using an Existing Volume as Secondary	90
10.2.4	Disabling the NCP/NSS Bindings for the Secondary Volume	93
10.2.5	Adding a Shadow to the Primary NSS Volume (Linking the NSS Volumes)	94
10.2.6	Moving Data between the Two Volumes	95
10.3	Giving Users a Merged View of the Shadow Volume	96
10.3.1	NCP	96
10.3.2	OES CIFS	96
10.3.3	Novell Samba with ShadowFS and FUSE	96
10.4	Replicating the Secondary File Tree Structure to the Primary Volume	97
10.5	Configuring the NCP/NSS Bindings for an NSS Volume	97
10.5.1	Disabling the NCP/NSS Bindings for an NSS Volume	98
10.5.2	Enabling the NCP/NSS Bindings for an NSS Volume	99
10.5.3	Enabling or Disabling NCP/NSS Bindings by Editing the /etc/opt/novell/ncp2nss.conf File	100
10.6	Viewing a List of NCP Shares	100
10.7	Mounting and Dismounting DST Shadow Volumes	101
10.8	Viewing the Name and Path Information for a Shadow Volume	101
10.9	Viewing Information about a Shadow Volume	101
10.9.1	Accessing the Volume Information Report	102
10.9.2	Viewing the Shadow Status of a Volume	102
10.9.3	Viewing the Share Information for a Shadow Volume	104
10.10	Auditing File Move Events for the Shadow Volume	106
10.11	Backing Up DST Shadow Volumes	106
10.11.1	Planning Your Backup Solution	106
10.11.2	Planning Your Restore Solution	107
10.11.3	Using the /etc/NCPVolumes XML File for Backup	109
10.11.4	Configuring the Backup Attribute for NSS Volumes	109
10.11.5	Configuring a Backup for Trustee Information on NSS Volumes on Linux	110
10.12	Removing the Shadow Relationship for a Non-Clustered DST Shadow Volume	110
10.12.1	Preparing to Remove a Shadow Volume	110
10.12.2	Removing the Shadow Volume Relationship by Using OES Remote Manager for Linux	111
10.12.3	Removing a Shadow Volume by Editing Configuration Files	112

11 Creating and Managing Policies for Shadow Volumes 113

11.1	Understanding Shadow Volume Policy Options	113
11.1.1	Last Executed	114
11.1.2	Files moved	114
11.1.3	Space moved	114
11.1.4	Total files moved	114
11.1.5	Total space moved	114
11.1.6	Description	114
11.1.7	Start Time	114
11.1.8	End Time	115

11.1.9	Start Day	115
11.1.10	Frequency	115
11.1.11	Command Status	116
11.1.12	Volume Selection	116
11.1.13	Volume Operations	117
11.1.14	Subdirectory Restrictions	117
11.1.15	Search Criteria	119
11.1.16	Stop	121
11.2	Creating a Shadow Volume Policy	121
11.2.1	Prerequisite	121
11.2.2	Guidelines for Shadow Volume Policies	121
11.2.3	Creating a Shadow Volume Policy	121
11.3	Modifying a Shadow Volume Policy	123
11.4	Running a Policy On Demand	124
11.5	Viewing DST Policies and Policy Status	124
11.6	Viewing Information about the Files Moved During a Policy Run	125
11.7	Stopping a Running Policy	126
11.7.1	Stopping All Running Shadow Volume Policies	126
11.7.2	Stopping a Running Individual Shadow Volume Policy	127
11.8	Deleting a Shadow Volume Policy	127

12 Managing Directory and User Space Quotas on DST Volumes 129

12.1	Understanding DST Volume Quotas Management	129
12.2	NPCON Quotas Command	132
12.3	Setting User Space Quotas on DST Volumes	134
12.4	Setting Directory Quotas on DST Volumes	137
12.5	Getting Quotas from an Old Secondary Volume to a New Primary Volume	140
12.6	Viewing User and Directory Quotas for a DST Volume	143
12.6.1	Viewing User Quotas from the Storage Plug-In in iManager	144
12.6.2	Viewing Directory Quotas from the Files and Folders Plug-In in iManager	144
12.6.3	Viewing User and Directory Quotas in NFARM	144

13 Using ShadowFS to Provide a Merged View for Novell Samba Users 145

13.1	Understanding ShadowFS	145
13.2	Prerequisites for Using ShadowFS	146
13.3	Preparing Your System for Using ShadowFS	146
13.4	Installing ShadowFS and FUSE	147
13.5	Setting Rights to ShadowFS Shares	149
13.6	Creating a Samba Share	150
13.7	Adding a User to Samba	150
13.8	Connecting Users to the Share	151
13.9	Testing Shadow Volume Policies	151
13.10	Enabling or Disabling ShadowFS	152
13.10.1	Loading ShadowFS and FUSE	152
13.10.2	Verifying ShadowFS Commands	152
13.11	Starting and Stopping ShadowFS Manually	152
13.11.1	Starting FUSE and ShadowFS	153
13.11.2	Starting FUSE and ShadowFS with novell-shadowfs	154
13.11.3	Stopping Shadowfs	154
13.12	Configuring Trustee Rights for Novell Samba Users	154

14	Generating a File Inventory for DST Shadow Volumes	155
14.1	Understanding the File Inventory for a Shadow Volume	155
14.1.1	Inventory Summary	155
14.1.2	Available Space Trends	156
14.1.3	Graphical Profiles	156
14.1.4	Tabular Profiles	163
14.1.5	Inventory Detail Reports	163
14.1.6	Custom Shadow Volume Options	164
14.2	Creating the Shadow Volume Inventory	166
14.3	Viewing a Saved NCP Volume Report	166
14.4	Viewing Statistics for the Shadow Volume	166
14.5	Using Inventory Detail Reports to Move, Copy, or Delete Files on the Shadow Volume	166
14.6	Generating a Custom Inventory Report	167
15	Troubleshooting for DST	171
15.1	My NCP server information is set to: LOCAL_CODE_PAGE CP437. Why is it not using UTF-8?	171
15.2	A File is listed twice in a directory	171
15.3	Users cannot see some files and directories	172
15.4	Cross-protocol locking stops working	172
15.5	OES Remote Manager connection error when you are working on the DST Options page	172
15.6	Unable to access files on the shadow volume in a DST volume pair	173
15.7	An error message "Shadow Volume Path is a non-NSS Volume." is displayed.	173
16	Security Considerations	175
16.1	Client Access	175
16.2	Linux-Enabled eDirectory Users	176
16.3	Using File System Trustees and Rights	176
16.4	Server-to-Server Access	176
16.5	Hidden Directories and Files	177
16.5.1	Trustee Database	177
16.5.2	Available Space Trends	177
16.6	Shadow Volumes Audit Logs	177
16.7	Shadow File System Audit Logs	177
16.8	NCP Server Auditing and Log Files	177
16.9	Using Secure Remote Connections	178
17	Configuring DST Shadow Volume Pairs with OES Cluster Services	179
17.1	Planning for DST in a Cluster	179
17.1.1	Open Enterprise Server SP3	180
17.1.2	OES Cluster Services	180
17.1.3	NCP Server and Dynamic Storage Technology	180
17.1.4	OES Storage Services File System	180
17.1.5	OES Remote Manager for Linux	180
17.1.6	Merged View Access with NCP	180
17.1.7	Merged View Access with OES CIFS	181
17.1.8	Merged View Access with Novell Samba and ShadowFS	181
17.1.9	No Merged View Access for AFP	182
17.2	Planning for DST Shadow Volume Pairs and Policies in a Cluster	182

17.2.1	DST Pool Cluster Resource	183
17.2.2	Shadow Volume Definition in the /etc/NCPVolumes File	183
17.2.3	Shadow Volume Definition in the ncperv.conf File	183
17.2.4	NCP2NSS Bindings for the Secondary Volume	184
17.2.5	NPCON Mount Command for the Load Script	184
17.2.6	Load Order in the Load Script	184
17.2.7	Unload Order in the Unload Script	185
17.2.8	Monitoring Storage in the Monitor Script	185
17.2.9	Additional Volumes in the Primary Pool	185
17.2.10	Policies for DST Nodes and Volumes in a Cluster	186
17.3	Preparing the Nodes to Support DST in a Cluster Environment.	187
17.4	Configuring the DST Pool Cluster Resource with Two Cluster-Enabled Pools	187
17.4.1	Overview of the Two Pool Cluster Resources	188
17.4.2	Viewing the Scripts for the Two Pool Cluster Resources	189
17.4.3	Adding Commands for the Secondary Clustered Pool and Volume to the Primary Pool Cluster Resource.	191
17.5	Configuring the DST Pool Cluster Resource with a Cluster-Enabled Pool and a Shared Pool	197
17.5.1	Overview of the Pool Cluster Resource and Shared Pool	197
17.5.2	Viewing the Scripts for Pool Cluster Resource	198
17.5.3	Creating a Shared Pool and Volume that Are Not Cluster-Enabled	200
17.5.4	Adding Commands for the Secondary Shared Pool and Volume to the Primary Pool Cluster Resource.	201
17.6	Sample Scripts for a DST Pool Cluster Resource	207
17.6.1	Sample Load Script for a DST Shadow Volume.	208
17.6.2	Sample Unload Script.	208
17.6.3	Sample Monitor Script for a DST Volume	209
17.7	Configuring Shadow Volume Policies for the Clustered DST Volume Pair.	209
17.8	Renaming a Shared Pool in a DST Cluster Resource	209
17.9	Renaming a Shared Volume in a DST Cluster Resource	211
17.10	Removing the Shadow Relationship for a Clustered DST Volume Pair	215
17.10.1	Planning to Remove the Shadow Relationship for a Clustered DST Volume Pair	215
17.10.2	Preparing to Remove a Shadow Relationship.	216
17.10.3	Removing the Shadow Definition and NCP/NSS Bindings Exclusion on All Nodes	217
17.10.4	Preparing the Primary Pool Cluster Resource for Independent Use	218
17.10.5	Preparing the Secondary Pool and Volume for Independent Use	219
17.11	Upgrading a Cluster with DST Resources from OES 2 SP3 to OES 2015 or Later	223

A Commands and Utilities for Dynamic Storage Technology 225

A.1	Using NPCON Commands for DST	225
A.1.1	Interactive Mode	225
A.1.2	Command Line Mode	226
A.1.3	Scripting Mode	226
A.2	NPCON Commands for Managing DST.	226
A.2.1	Creating a DST Shadow Volume Pair.	227
A.2.2	Removing the Shadow Relationship, or Unlinking the Volumes	227
A.2.3	Listing or Moving Files that Match Search Criteria	228
A.2.4	Listing or Moving a File, or Shifting a File between Volumes	230
A.3	NPCON Commands for DST in a OES Cluster Services Cluster	231
A.3.1	Scenario 1: Primary NSS and Shadow NSS	231
A.3.2	Scenario 2: Primary Non-NSS and Shadow Non-NSS (Not supported).	232
A.3.3	Scenario 3: Primary Non-NSS and Shadow NSS (Not supported).	232
A.3.4	Scenario 4: Primary NSS and Shadow Non-NSS (Not Supported).	232
A.4	Configuring Global DST Policies by Using the SET Command.	233

A.4.1	Understanding DST Parameters for the SET Command	233
A.4.2	Using OES Remote Manager to Configure DST Parameters for the SET Command	234
A.4.3	Using the ncpcon set Command to Configure DST Parameters	236
A.5	DST Commands for /etc/opt/novell/ncpserv.conf	236
A.6	DST Commands for /etc/opt/novell/shadowfs.conf	237
A.7	DST EXCLUDE_VOLUME Command for /etc/opt/novell/ncp2nss.conf	238
A.8	DST Shadow Volume Information in /etc/NCPVolumes	238
A.9	DST ShadowFS Volume Information in /etc/mtab.shadowfs	239
A.10	Verifying and Syncing ACLs (Inherited Rights Filter (IRF) and Trustees) from DST Primary Volume to Shadow Volume	239

About This Guide

This guide describes how to install, configure, and manage the Dynamic Storage Technology for Open Enterprise Server (OES) 2018 SP3. It is divided into the following sections:

- ♦ [Chapter 1, “Overview of Dynamic Storage Technology,” on page 13](#)
- ♦ [Chapter 2, “What’s New or Changed in Dynamic Storage Technology,” on page 27](#)
- ♦ [Chapter 3, “Planning Your Dynamic Storage Technology Server Environment,” on page 29](#)
- ♦ [Chapter 4, “Installing Dynamic Storage Technology,” on page 35](#)
- ♦ [Chapter 5, “Using DST in a Virtual Environment,” on page 39](#)
- ♦ [Chapter 6, “Management Tools for DST,” on page 41](#)
- ♦ [Chapter 7, “Configuring DST Global Policies,” on page 49](#)
- ♦ [Chapter 8, “Managing Services for DST,” on page 63](#)
- ♦ [Chapter 9, “Planning for DST Shadow Volume Pairs and Policies,” on page 65](#)
- ♦ [Chapter 10, “Creating and Managing DST Shadow Volumes for NSS Volumes,” on page 85](#)
- ♦ [Chapter 11, “Creating and Managing Policies for Shadow Volumes,” on page 113](#)
- ♦ [Chapter 12, “Managing Directory and User Space Quotas on DST Volumes,” on page 129](#)
- ♦ [Chapter 13, “Using ShadowFS to Provide a Merged View for Novell Samba Users,” on page 145](#)
- ♦ [Chapter 14, “Generating a File Inventory for DST Shadow Volumes,” on page 155](#)
- ♦ [Chapter 15, “Troubleshooting for DST,” on page 171](#)
- ♦ [Chapter 16, “Security Considerations,” on page 175](#)
- ♦ [Chapter 17, “Configuring DST Shadow Volume Pairs with OES Cluster Services,” on page 179](#)
- ♦ [Appendix A, “Commands and Utilities for Dynamic Storage Technology,” on page 225](#)

Audience

This guide is intended for storage services administrators. Security administrators can find a summary of security information for Dynamic Storage Technology in [Chapter 16, “Security Considerations,” on page 175](#).

It is assumed that the reader has some understanding of the OES Services components that are used with Dynamic Storage Technology, including the OES operating system, the NSS file system, the file access services (NCP (NetWare Core Protocol), CIFS (Common Information File System), Novell Samba (SMB/CIFS), SSH, and OES FTP (Pure-FTPd)), and OES Cluster Services (NCS).

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

Documentation Updates

For the most recent version of the *Dynamic Storage Technology Administration Guide*, see the [Open Enterprise Server 2018 SP3 documentation website](#).

Additional Documentation

For documentation on OES Storage Services (NSS) volumes, see the (<https://www.microfocus.com/documentation/open-enterprise-server/2018-SP3/>).

For documentation on NCP Server and NCP file access, see the *OES 2018 SP3: NCP Server for Linux Administration Guide*.

1 Overview of Dynamic Storage Technology

Dynamic Storage Technology (DST) for Open Enterprise Server (OES) is an information life-cycle management technology. It makes your essential data readily available to the eDirectory and Active Directory users, while tiering files efficiently across a pair of independent OES Storage Services (NSS) volumes, referred to as a *DST shadow volume*. You create policies to control how the files are distributed between the two volumes.

The two NSS volumes reside on different devices on the same server. The primary volume typically contains active or highly critical files, while the secondary volume contains files that are accessed less often. When users connect to a network share on the primary NSS volume, they see a merged view of files on both volumes. Users are not aware of where the files physically reside. Files on both volumes are equally accessible to users. Dynamic Storage Technology pulls data directly to the user from the primary volume or the secondary volume, depending on where the file is located.

NCP (NetWare Core Protocol) client users and CIFS users automatically see a merged view of the files and subdirectories on the shadow volume pair when they access a share on the primary volume. All the actions they take—renaming, deleting, moving, etc.—are synchronized by Dynamic Storage Technology across the two volumes. If you use supported native Linux file access protocols, such as Novell Samba, SSH, or OES FTP (Pure-FTP-d), to access the DST volume, you can enable ShadowFS to provide a merged view location for LUM-enabled users of those protocols. OES CIFS and Novell Samba cannot be installed on the same server.

A Dynamic Storage Technology policy determines what files are moved between the two NSS volumes that form the shadow volume. You can specify which direction to move files (primary to secondary, or secondary to primary). Policy filters allow you to specify one or more conditions to be met, such as frequency of use, file extensions, and file size. Policy enforcement is automated with scheduled and on demand policy runs. You can run multiple policies that start concurrently on a shadow volume. You can also specify a list of files or folders to be moved during a one-time move from the primary volume to the secondary volume.

Dynamic Storage Technology allows you to seamlessly tier storage between high-performance and lower-performance devices. For example, you can establish policies that keep frequently used, mission-critical data on high-performance devices, and move rarely accessed, less-essential data to lower-performance devices.

Backup can be performed separately on the two volumes, which allows for different backup schedules. Backing up essential files takes less time because the seldom used files are stored on the secondary path, where they can be backed up separately and less frequently. This helps narrow the time window needed for backing up critical data.

Dynamic Storage Technology enables you to manage data more efficiently for the enterprise. In doing so, the enterprise can potentially realize significant cost savings in storage management.

All the Active Directory users can now access the data on DST volumes via CIFS. To manage the rights of the Active Directory trustees on the DST volumes, you can use OES File Access Rights Management (NFARM) utility or `rights` utility. For more information, see [Managing the Trustee Rights in the NSS File System](#) and `rights` in the [OES 2018 SP3: NSS File System Administration Guide for Linux](#).

This section provides an overview of Dynamic Storage Technology and its components.

- ♦ [Section 1.1, “Understanding Dynamic Storage Technology,” on page 14](#)
- ♦ [Section 1.2, “Benefits of Dynamic Storage Technology,” on page 18](#)
- ♦ [Section 1.3, “Shadowing Scenarios,” on page 20](#)
- ♦ [Section 1.4, “DST Policy Scenarios,” on page 22](#)
- ♦ [Section 1.5, “DST Components,” on page 24](#)
- ♦ [Section 1.6, “Management Tools,” on page 25](#)
- ♦ [Section 1.7, “What’s Next,” on page 25](#)

1.1 Understanding Dynamic Storage Technology

Dynamic Storage Technology (DST) allows you to specify a shadow relationship between two volumes to form a *shadow volume pair*. The secondary directory tree structure, or *secondary file tree*, shadows the primary file tree. The primary tree and the secondary tree are overlaid to create one virtual volume tree that is transparently presented to the users. Users see a merged view of the files on both volumes.

IMPORTANT: Only OES Storage Services (NSS) volumes are supported for use in a DST shadow volume pair.

Dynamic Storage Technology allows you the flexibility of moving files between the two locations while maintaining a consistent single file tree view of the files for users. If the shadow relationship is removed, the two volumes can once again function independently and normally.

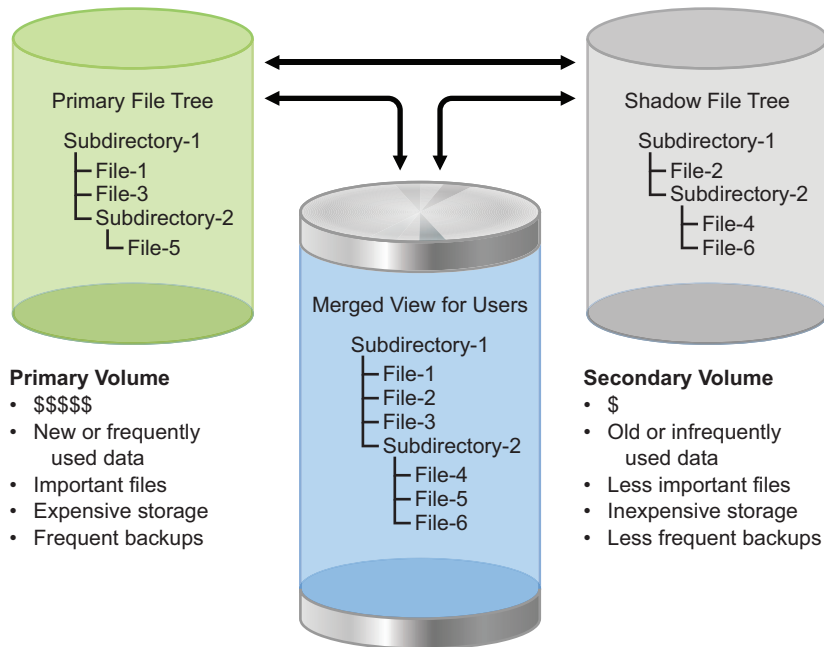
- ♦ [Section 1.1.1, “Merged View of the File Tree,” on page 14](#)
- ♦ [Section 1.1.2, “File Systems,” on page 16](#)
- ♦ [Section 1.1.3, “File Access Protocols,” on page 16](#)
- ♦ [Section 1.1.4, “Secure File Access,” on page 17](#)
- ♦ [Section 1.1.5, “Local File Access for Backup and Archive,” on page 18](#)

1.1.1 Merged View of the File Tree

Dynamic Storage Technology presents the file system directory trees on each volume in a merged view for users. The primary file tree and secondary file tree have the same directory structure, so that each directory has an instance in both locations. Each file has a single instance that resides on only one volume at a time. When accessing files through the merged view, users are not aware of the actual physical location of the files.

An example of the merged view of the shadow volume is shown in [Figure 1-1](#). When an NCP client lists files for `Subdirectory-1`, the user sees `File-1`, `File-2`, and `File-3`. `File-1` and `File-3` are stored in the primary file tree. `File-2` is stored in the shadow file tree.

Figure 1-1 User View of the File System Directory



In general, transactions are executed wherever the file resides. Any file that requires a normal user-level action (copy, delete, and so on) is moved back to the primary for the action to take place, which simplifies the auditing requirements. Some transactions, such as a directory rename, occur in both file trees in order to keep the paths synchronized.

When a client creates new files, the files are automatically stored in the primary file tree. When files in the shadow file tree are modified, a configurable option allows the files to be moved to the primary file tree (default), or left in the shadow file tree. Which method you choose depends upon your storage goals.

For example, if your goal is to place newer files in the primary file tree and to place older files in the shadow file tree, you want an older file in the secondary file tree to move to primary file tree if the file's content is modified. On the other hand, if your goal is to place files of one type (such as `.doc` and `.ppt`) in the primary area and files of a different type (such as `.mp3` and `.jpg`) in the secondary area, you want files to stay where they are whenever they are viewed or modified.

When a new directory is created, it is created in the primary file tree. A configurable option allows the necessary branches of the tree to be created in the shadow file tree in one of two ways:

- The new directory path is created as needed when policies are enforced to move files to the directory in the secondary location. (the default)
- The new directory path is created immediately in the secondary location, and files are moved to the directory as policies are enforced.

Performance is better when the directory branches are created only as needed.

When a directory is deleted, it is deleted in both areas. When a directory is renamed, it is renamed in both areas. The coordination of the file and directory management happens automatically so that the areas remain synchronized and have the same directory structure.

1.1.2 File Systems

The primary volume and secondary volume must use the same type of file system. Only the NSS file system is supported at this time.

Dynamic Storage Technology supports mount points for volumes anywhere in the logical Linux directory tree that is available to the server. This allows you to use DST even for NSS volumes that use non-standard mount points. The primary volume and secondary volume must not use the same mount point.

The primary volume and secondary volume can be located on devices that appear as local devices to the server, including local SCSI devices, Fibre Channel SAN devices, and iSCSI SAN devices, in any combination. The device types and performance can differ between the primary and secondary devices, with the secondary volume typically being on the device with lower performance.

Clustering of a shadow volume pair is supported with OES Cluster Services. For information, see [Chapter 17, “Configuring DST Shadow Volume Pairs with OES Cluster Services,”](#) on page 179.

1.1.3 File Access Protocols

Dynamic Storage Technology works closely with NCP Server to provide the merged view for users. eDirectory users and Active Directory users can see the merged view when they access the primary volume with a supported file access protocol.

- ♦ [“NCP Clients” on page 16](#)
- ♦ [“OES CIFS” on page 16](#)
- ♦ [“Native Linux File Access via ShadowFS” on page 16](#)

NCP Clients

NCP Server supports Dynamic Storage Technology by providing a merged view via an NCP client for eDirectory users. Users connect to the primary volume and automatically see merged view of files on both volumes.

OES CIFS

OES CIFS supports Dynamic Storage Technology by providing a merged view via CIFS for eDirectory users and Active Directory users. OES CIFS leverages the NCP merged view cache to provide the merged view. It is not necessary to enable the users with Linux User Management (LUM). Use OES CIFS with the primary NSS volume in the DST shadow volume pair, just as you would with a regular NSS volume.

Native Linux File Access via ShadowFS

The Dynamic Storage Technology Shadow File System (ShadowFS) provides the merged view for native Linux file access via Novell Samba (SMB/CIFS) and supported Linux file access protocols, such as SSH and OES FTP (Pure-FTPd). ShadowFS uses FUSE (File System in Userspace) to provide a merged view location with a local mount point of `/media/shadowfs/<primary_volume_name>`

for each DST shadow volume. For information, see [Chapter 13, “Using ShadowFS to Provide a Merged View for Novell Samba Users,”](#) on page 145. Novell Samba and OES CIFS cannot be installed on the same server.

When using ShadowFS, the eDirectory users and the native Linux file access services must be enabled with Linux User Management (LUM). For information, see [OES 2018 SP3: Linux User Management Administration Guide](#).

1.1.4 Secure File Access

Dynamic Storage Technology can be accessed by eDirectory users and Active Directory users. Although DST accommodates transactions across both volumes, the file access for a shadow volume pair is ultimately controlled by the same security features that apply for non-shadowed NSS volumes.

Dynamic Storage Technology supports the OES Trustee Model to provide secure file access for eDirectory users. An administrator accesses the shadow volume pair via an NCP client’s merged view, and sets trustees and trustee rights on directories and files in the merged file tree. NCP saves the settings to the trustee database file on the primary NSS volume, and copies the updated trustee database file to the secondary NSS volume. This ensures that the trustee and trustee rights settings are consistent across the two volumes. Trustees and trustee rights are enforced by the same NCP rights management features that apply for non-shadowed NSS volumes.

Authorized users perform directory and file operations (such as read, write, rename, delete, and so on) from a merged view. DST transparently executes directory operations on the directory’s instance on the primary volume and its instance on the secondary volume. DST transparently executes a file transaction on the volume where the file resides. The user is not aware of whether the file is on the primary volume or the secondary volume. As the operation is executed, the directory and file attributes are enforced by the same NSS file system management features used by non-shadowed NSS volumes.

NCP Server and OES CIFS automatically provide a merged view of shadow volume pairs built with NSS volumes. The NCP clients and CIFS users access files normally via this merged view.

A merged view of a shadow volume pair is not available in a native Linux file browser. The primary file tree and the shadow file tree are separate and independent directories. The NSS file system allows the `root` user to view files on these paths, but this access is unrelated to DST. `root` user file access rights are required to back up the two locations as described in [Section 1.1.5, “Local File Access for Backup and Archive,”](#) on page 18.

DST ShadowFS provides a merged view interface where LUM-enabled eDirectory users can access the shadow volume pair via supported native Linux file access protocols. For information, see [“Native Linux File Access via ShadowFS”](#) on page 16. The users access files normally via this merged view interface.

1.1.5 Local File Access for Backup and Archive

Backup administrators and system administrators with `root` user privileges on the server can see the primary file tree and the shadow file tree as separate and independent directories. Thus, backup tools can apply one backup policy to the primary volume's Linux path and apply a different backup policy to the secondary volume's Linux path.

While it is in a shadow relationship, the secondary volume is mounted in Linux, but it is not mounted in NCP. eDirectory users do not access its files directly. The NSS file system allows the `root` user to view files on these paths, but this access is unrelated to DST. The only operations that are intended to take place directly on the secondary volume are backup, or "remove and archive."

1.2 Benefits of Dynamic Storage Technology

Dynamic Storage Technology shadow volume pairs have many benefits:

- ◆ [Section 1.2.1, "Merged View File Access for End Users," on page 18](#)
- ◆ [Section 1.2.2, "Policy-Based Migration between Primary and Secondary Storage Areas," on page 18](#)
- ◆ [Section 1.2.3, "Faster and Smaller Backups of Important Data," on page 19](#)
- ◆ [Section 1.2.4, "Faster Disaster Recovery," on page 19](#)
- ◆ [Section 1.2.5, "More Efficient Use of Expensive Storage," on page 19](#)
- ◆ [Section 1.2.6, "Fast Storage for Active Data and Slower, Less Expensive Storage for Old Data," on page 20](#)
- ◆ [Section 1.2.7, "Moving Files from an Existing Secondary Volume," on page 20](#)
- ◆ [Section 1.2.8, "Access to the Secondary Storage Area without the Performance Penalties of HSM Solutions," on page 20](#)

1.2.1 Merged View File Access for End Users

Dynamic Storage Technology shadow volume pairs present a merged view of the file trees. The end user's files appear to be in the same logical place regardless of their physical location. This allows the administrator to manage the data without disrupting the end user's view of the files.

1.2.2 Policy-Based Migration between Primary and Secondary Storage Areas

Dynamic Storage Technology provides policy-based control of the files to move and the direction that you want to move data between devices. You can set up policies that migrate data by file extension, file size, and the date a file was last accessed or modified. You can also specify a list of files to move in either direction for a one-time move. Policies can be scheduled to run and run on demand. You can set policies so that data stored on the secondary storage volume can be accessed without de-migrating it. For information about using policies, see [Chapter 11, "Creating and Managing Policies for Shadow Volumes," on page 113](#).

1.2.3 Faster and Smaller Backups of Important Data

Backup administrators and system administrators with `root` user privileges on the server can see the primary file tree and the shadow file tree as separate and independent directories. Backup policies can differ for the primary storage volume and the secondary storage volume. For example, the DST administrator can allocate the data between the two volumes in a way that supports separate backup schedules:

- ♦ Important or active data that needs to be maintained on quality storage and backed up frequently
- ♦ Less important or stale data that can be stored on less expensive storage and backed up less frequently

Analyzing the inventory of a volume's data shows that a large portion of its data is seldom used. Having a shadow volume allows the administrator to spend more on the most important data and spend less on the less important data. The important data, which is stored on the primary area, can be backed up nightly. The less important data, which is stored in the secondary area, can be backed up weekly or even monthly.

Getting the less important data out of the way enables the backups of your important data to run more quickly and efficiently. Allocating data in this way can significantly lower the cost of backups by reducing both labor and tape requirements.

For information about backing up data on DST shadow volume pairs, see:

- ♦ [Section 9.13, "Using Backup Utilities with DST Shadow Volume Pairs," on page 82](#)
- ♦ [Section 10.11, "Backing Up DST Shadow Volumes," on page 106](#)

1.2.4 Faster Disaster Recovery

Dynamic Storage Technology allows you to manage your most important files on the primary area. During a disaster recovery, the server administrator can restore the primary area first. This restores the critical files first, and leaves the recovery of the less important secondary area until later. The users can continue working while files they probably do not need immediately are being restored. Also, other fault-tolerant replication solutions like snapshots can be used for the primary area without wasting money on files that do not require the same level of fault tolerance.

1.2.5 More Efficient Use of Expensive Storage

Dynamic Storage Technology policies can help to partition files based on file age, owner, type, size, and so on. You can move the less important files from a higher quality storage array to a lower quality storage, thus reserving the higher-cost storage for your most important files.

For example, you can configure the primary area on block-based SCSI storage devices in a Fibre Channel SAN-based hardware RAID array or storage array, and configure the secondary area on a lower quality storage array using slower devices like SATA. This allows you to get more use out of your Fibre Channel storage solution, and keep it from filling up with unimportant files. You can store more data on your server with a lower overall cost per gigabyte.

1.2.6 Fast Storage for Active Data and Slower, Less Expensive Storage for Old Data

A Dynamic Storage Technology shadow volume pair can use different block storage solutions for its primary volume and secondary volume. Storage costs can be reduced by allowing data that is used infrequently to be stored on lower-cost storage. Locate the primary area on storage drives that are faster and higher quality. Then locate the secondary area on less expensive storage drives. Files that the users are currently working on can be located on the high-performance drives. The files that have not been modified for a long time can be moved to the lower-performance drives to free up space on the high-performance drives. In this way, you can locate a large amount of your data on less expensive, lower-performance storage drives, while your users still get high-quality performance because their active files are located on the high-performance storage drives.

1.2.7 Moving Files from an Existing Secondary Volume

You can start with an empty primary NSS volume, and have the shadow area be an existing volume. The combined view initially presented by the NCP Engine is equivalent to the secondary volume. You can define a policy to move files to the primary as they are modified or accessed. As users access their data through the new primary volume, the files they use are automatically migrated to the new server. This migration-on-demand approach migrates the data gradually, allows users to access the files as usual, and frees the IT department from spending off-hours time migrating the data with the server offline.

1.2.8 Access to the Secondary Storage Area without the Performance Penalties of HSM Solutions

With HSM (hierarchical storage management) solutions, files are migrated from the primary storage to a secondary storage device, and a copy of the file's metadata (stub file) is left behind in the volume's directory tree. If the file is ever accessed again, it needs to be migrated back to the primary storage before it is available.

Dynamic Storage Technology shadow volume pairs can access files directly regardless of which area (primary or secondary) they are in, and without de-migrating them. If a user searches through all the files on a shadow volume, the files are searched without needing to move them to the primary area. Also, shadow volume backups are faster because there are no HSM metadata stub files for the backup software to scan. The backup software does not need to be HSM aware.

1.3 Shadowing Scenarios

A newly defined Dynamic Storage Technology shadow volume pair consists of two new volumes or of one existing volume and one new volume. For this discussion, an existing volume is one that contains data. When you use an existing volume in a pair, the new volume can act as the primary volume or the secondary volume. Shadowing scenarios for new and old volume combinations are described below.

In general, you should not use two unrelated existing volumes in a shadow volume pair. However, if two volumes are already in a shadow relationship, their content is logically and structurally related. You can temporarily remove a shadow relationship for two volumes, and then re-create the pair. For

example, when a shadow volume pair is clustered with OES Cluster Services, the pair is defined when its cluster resource is online, and the pair is undefined when its resource is offline. You also might temporarily remove a shadow relationship to migrate a shadow volume pair between servers as two separate volumes, and then redefine their shadow relationship on the new server. The two volumes should not be independently accessed by users outside their shadow relationship.

The following sections describe shadowing scenarios with one existing volume and one new volume:

- ♦ [Section 1.3.1, “Existing Volume as Primary with an Empty Volume as Secondary,” on page 21](#)
- ♦ [Section 1.3.2, “Empty Volume as Primary with an Existing Volume as Secondary,” on page 21](#)

1.3.1 Existing Volume as Primary with an Empty Volume as Secondary

Your existing NSS volume currently contains many files that are seldom used and rarely change. You want to free space on the high-performance device for the unstructured files that are critical to your growing business. Backing up the volume takes more time than the window reserved for incremental and weekly backups. You want to move the stagnant files to a location where they can be easily accessed but backed up less frequently. This decreases the time it takes to back up or restore the data you use the most.

You set up iSCSI SAN storage to use for the secondary volume. You create a new NSS volume on the iSCSI disk, then define a shadow volume pair for the existing and new NSS volumes.

You configure a policy that governs what files to move to the secondary storage area. Files are returned to the primary area based on a policy of usage or file type. For example, if the user simply views the data in a file, then the data does not move. If the user modifies the file, then the file is moved back to the primary volume. Users are not aware of where the data are physically stored because they see a merged view of both volumes.

The backup administrator configures different backup schedules for the two volumes. The primary volume is backed up incrementally throughout the week, and fully at the weekends. The secondary volume is backed up less frequently. If a restore is needed, the files on the primary can be restored first and more quickly than was previously possible.

1.3.2 Empty Volume as Primary with an Existing Volume as Secondary

Your existing NSS volume resides on an older storage array in the SAN. You want to migrate the data to new high-performance storage arrays. The files are mission critical and the volume cannot be taken offline for a long period while you migrate the data.

You decide to use Dynamic Storage Technology to gradually and transparently migrate files to the new storage. You create a new NSS volume on a storage device in the high-performance storage array. You define a shadow volume pair that uses the empty volume as the primary area, and the existing volume as the secondary area.

You can configure a policy so that the data moves to the primary volume based upon usage. Data gradually flows to the primary volume as it is used. In this way, there is a natural background migration of data from the existing volume to the new volume. The new volume grows, and contains the files used most recently by the users. You can use policies to move all of the files to the new primary volume, or you can leave the secondary volume in place to store the little used and unchanged files.

For example, suppose you have an existing pool that spans multiple LUNs, and contains multiple volumes. The current best practice is to use a separate LUN for each pool, and a single volume per pool. You create a new pool on a new larger LUN (or fewer larger LUNs), then create a single NSS volume in the pool. You might need to rename the old and new NSS volumes if users need to access the data via known paths, because after the shadow volume is created, users access data via the new volume. Repeat this process so that you have one new empty volume for each of the old volumes on the pool. As the new and old volumes are ready, you create a DST shadow volume with the new volume as the primary storage area and an existing volume from the old pool as the secondary storage area.

To begin de-migrating the data, configure the global policies to shift data from the secondary storage area to the primary storage area whenever they are accessed or modified. You can also configure individual shadow volume policies or use inventory reports to shift data on schedule or on-demand based on age, file names, file types, or file size. De-migration occurs with the storage online and accessible to end users; they are not aware of where the data is actually stored. To ensure that the entire existing directory structure (including empty directories) is replicated on the primary volume, an administrator can enumerate the directory from an NCP client that is mapped to the merged view. After all of the data is moved from the old NSS volume to the new one, you can remove the shadow volume relationship, and delete the old NSS volume from the old pool. Users are not aware that the volume is on a new pool. They see only the volume by its name.

You can apply this process to migrate data from the other volumes in the pool. When all data has been migrated and the old volumes are deleted, you can delete the old pool, which frees that storage for other uses.

1.4 DST Policy Scenarios

Dynamic Storage Technology volume policies control how data flows between the primary storage area and the secondary storage area.

- ◆ [Section 1.4.1, “Move Files Based on the Last Time Accessed or Modified,” on page 22](#)
- ◆ [Section 1.4.2, “Move Files Based on File Size,” on page 23](#)
- ◆ [Section 1.4.3, “Move Files Based on File Extensions,” on page 23](#)
- ◆ [Section 1.4.4, “Move Selected Files from Primary to Secondary,” on page 23](#)

1.4.1 Move Files Based on the Last Time Accessed or Modified

You can create a DST volume policy that moves files to the secondary volume that have not been modified or accessed for a period of time, such as 6 months or 1 year. This allows you to keep active files on the primary volume, and seldom-used files on the secondary volume.

You can create an individual volume policy that runs daily, weekly, or monthly to move a file from the secondary storage area to the primary storage area if it has recently been modified. For example, if a backup is performed more frequently on the primary volume, this allows all modified files to be returned to the primary area by running the policy before the primary volume’s scheduled backup window. Typically, the policy run and backup window occur during non-peak hours.

You can also use a DST global policy called **Shift Modified Shadow Files** to ensure that modified files are returned to the primary storage area. This NCP Server parameter is enabled by default. When the parameter is enabled, if a user modifies a file, the file is automatically moved to the primary storage area when the file is closed. This setting applies to all DST shadow volume pairs on the server. The policy does not move files that were modified prior to the setting being enabled.

Another DST global policy called **Shift Accessed Shadow Files** can be used to move files when they are accessed as read-only a second time within a specified time period. This NCP Server parameter is disabled by default. When this parameter is enabled, the accessed file is moved back to the primary area when the file is closed. By default, the period of time is 1 day. Use the **Shift Days Since Last Access** parameter to specify the period of time. Both settings apply to all DST shadow volume pairs on the server.

1.4.2 Move Files Based on File Size

You can create a DST policy that moves files to the secondary volume based on the **File Size Restriction** that includes the options greater than or less than and a specified file size.

To keep this separation of files by size, regardless of whether they are modified, you can disable the DST global policy called **Shift Modified Shadow Files**. This NCP Server parameter is enabled by default. When the parameter is disabled, if a user modifies a file on the secondary volume, the file remains on that volume when it is closed. This setting applies to all DST shadow volume pairs on the server.

To keep this separation of files by size, you should also disable the DST global policy called **Shift Accessed Shadow Files**. This NCP Server parameter is disabled by default.

1.4.3 Move Files Based on File Extensions

You can create a DST policy that moves non-essential types of files based on file extensions, such as *.jpg, *.mp3, *.wma, *.mpeg, *.iso, *.zip, *.cab, and so on.

To keep this separation of files by file extension, regardless of whether they are modified, you can disable the DST global policy called **Shift Modified Shadow Files**. This NCP Server parameter is enabled by default. When the parameter is disabled, if a user modifies a file on the secondary volume, the file remains on that volume when it is closed. This setting applies to all DST shadow volume pairs on the server.

To keep this separation of files by file extension, you should also disable the DST global policy called **Shift Accessed Shadow Files**. This NCP Server parameter is disabled by default.

1.4.4 Move Selected Files from Primary to Secondary

You can use the Shadow Volume Inventory page in OES Remote Manager for Linux to view statistics on files and usage for the DST shadow volume. At the bottom of the page, use the form to move selected files from the primary volume to the secondary volume. This one-time move is not policy based.

1.5 DST Components

There are four main components for Dynamic Storage Technology.

- ♦ [Section 1.5.1, “NCP Engine,” on page 24](#)
- ♦ [Section 1.5.2, “Shadow Volume,” on page 24](#)
- ♦ [Section 1.5.3, “ShadowFS,” on page 24](#)
- ♦ [Section 1.5.4, “Policy Engine,” on page 24](#)

1.5.1 NCP Engine

The NCP Engine provides support for NCP clients and is the main file copy engine for Dynamic Storage Technology. It provides the merged view for NCP users and CIFS users. It supports the ShadowFS access for supported native Linux file access protocols, such as Novell Samba (SMB/CIFS), SSH, and OES FTP (Pure-FTPd).

1.5.2 Shadow Volume

Shadow Volume provides the merged file-tree view for NCP client users. The NCP merged view cache is leveraged to provide a merged view for CIFS users.

Beginning with OES 2015, Active Directory users can also access files residing in shadow volumes.

1.5.3 ShadowFS

The Shadow File System (ShadowFS) works with FUSE (File System in Userspace) to provide a merged view for eDirectory users of native Linux file access protocols, including Novell Samba (SMB/CIFS) and supported Linux file access protocols, such as SSH and OES FTP (Pure-FTPd). ShadowFS requires that eDirectory users and the native Linux service be enabled with Linux User Management (LUM).

IMPORTANT: OES CIFS does not require ShadowFS, FUSE, and LUM. For information, see [“OES CIFS” on page 69](#).

ShadowFS uses FUSE to create a local mount point for each DST shadow volume in `/media/shadowfs/<primary_volume_name>`. FUSE is an open source software package that is installed automatically when you install DST.

1.5.4 Policy Engine

The Dynamic Storage Technology policy engine allows you to create, manage, and enforce policies for a shadow volume pair. You can use the following

- ♦ **Global Policies:** Global policies are a set of NCP Server parameters that govern DST behavior. The settings apply at the server level for every mounted shadow volume pair on the server. They do not affect other NCP volumes. For information, see [Chapter 7, “Configuring DST Global Policies,” on page 49](#).

- ♦ **Volume Policies:** A volume policy applies to one or more specified shadow volume pairs. You can also create a volume policy that automatically applies to all shadow volume pairs. For information, see [Chapter 11, “Creating and Managing Policies for Shadow Volumes,”](#) on page 113.

1.6 Management Tools

The rights of Active Directory (AD) trustees on the Dynamic Storage Technology shadow volumes can be managed by using OES File Access Rights Management (NFARM) utility or `rights` utility. For information about managing rights of AD trustees, see [Section 6.3.1, “OES File Access Rights Management \(NFARM\),”](#) on page 47 and `rights` in the [OES 2018 SP3: NSS File System Administration Guide for Linux](#).

Dynamic Storage Technology shadow volumes, global policies, and shadow volume policies can be managed in OES Remote Manager for Linux. For information about using OES Remote Manager, see [Chapter 6, “Management Tools for DST,”](#) on page 41.

DST shadow volume pairs can be created, mounted, and dismounted with commands by using the NCP Console (NCPCON, `ncpcon(8)`) utility. For information, see [Appendix A, “Commands and Utilities for Dynamic Storage Technology,”](#) on page 225. NCPCON commands are also used in the load, unload, and monitor scripts for clustered DST shadow volume pairs.

1.7 What’s Next

For information about installing NCP Server, Dynamic Storage Technology, and NSS, see [Chapter 4, “Installing Dynamic Storage Technology,”](#) on page 35.

For information about planning your DST solution, see [Chapter 9, “Planning for DST Shadow Volume Pairs and Policies,”](#) on page 65.

2 What's New or Changed in Dynamic Storage Technology

This section describes enhancements and changes in Dynamic Storage Technology, beginning with the initial release of Open Enterprise Server (OES) 2018.

2.1 What's New or Changed in Dynamic Storage Technology (OES 2018 SP3)

DST in OES 2018 SP3 has been modified for bug fixes, DST provides the following enhancements and changes in OES 2018 SP3

Replicating Branches of the Primary File Tree in the Secondary File Tree

A configurable option called Replicate Primary Tree to Shadow is by default set to 1, so when the existing NSS volumes are configured as primary volume for DST, the primary volume directory tree structure will get replicated to the shadow volume directory tree. See [Section 7.1, "Replicating Branches of the Primary File Tree in the Secondary File Tree,"](#) on page 49

2.2 What's New or Changed in Dynamic Storage Technology (OES 2018 SP1)

DST in OES 2018 SP1 has been modified for bug fixes. There are no new features or enhancements in OES 2018 SP1.

2.3 What's New or Changed in Dynamic Storage Technology (OES 2018-Update 1-Patch)

The `dst_verify_resync_trustee.py` utility is newly introduced in this patch to verify and sync the ACLs from DST primary volume to shadow volume. For more information on `dst_verify_resync_trustee.py` utility, see [Verifying and Syncing ACLs \(Inherited Rights Filter \(IRF\) and Trustees\) from DST Primary Volume to Shadow Volume.](#)

2.4 What's New or Changed in Dynamic Storage Technology (OES 2018)

DST in OES 2018 has been modified for bug fixes. There are no new features or enhancements in OES 2018.

3 Planning Your Dynamic Storage Technology Server Environment

This section describes the software requirements and configuration guidelines for installing and using Dynamic Storage Technology (DST) on your Open Enterprise Server (OES) servers.

IMPORTANT: For information about planning a Dynamic Storage Technology solution for your OES Storage Services (NSS) volumes, see [Chapter 9, “Planning for DST Shadow Volume Pairs and Policies,”](#) on page 65.

- ◆ [Section 3.1, “Open Enterprise Server 2018 SP1,”](#) on page 29
- ◆ [Section 3.2, “NCP Server and Dynamic Storage Technology,”](#) on page 30
- ◆ [Section 3.3, “OES Storage Services,”](#) on page 30
- ◆ [Section 3.4, “Directory Services,”](#) on page 30
- ◆ [Section 3.5, “NCP File Access,”](#) on page 31
- ◆ [Section 3.6, “OES CIFS File Access,”](#) on page 31
- ◆ [Section 3.7, “Novell Samba File Access with ShadowFS, FUSE, and LUM,”](#) on page 31
- ◆ [Section 3.8, “Supported Native Linux File Access Protocols with ShadowFS, FUSE, and LUM,”](#) on page 32
- ◆ [Section 3.9, “OES AFP File Access \(Not Supported\),”](#) on page 32
- ◆ [Section 3.10, “Linux User Management,”](#) on page 32
- ◆ [Section 3.11, “OES Cluster Services for Linux,”](#) on page 32
- ◆ [Section 3.12, “OES Remote Manager for Linux,”](#) on page 32
- ◆ [Section 3.13, “iManager for Linux,”](#) on page 33
- ◆ [Section 3.14, “SFCB and CIMOM,”](#) on page 33
- ◆ [Section 3.15, “Other OES Services,”](#) on page 34

3.1 Open Enterprise Server 2018 SP1

Dynamic Storage Technology runs on OES servers with 64-bit processors. For information about installing and configuring OES, see the [OES 2018 SP3: Installation Guide](#).

3.2 NCP Server and Dynamic Storage Technology

The **Novell NCP Server / Dynamic Storage Technology** pattern on the OES installation disk includes the NCP Server software, the Dynamic Storage Technology software, and the plug-in components for OES Remote Manager that allow you to create and manage DST shadow volume pairs.

NCP Server must be installed and running in order for DST to work. NCP Server provides the NCP file access services and merged view for a DST shadow volume pair. The NCP Console (NCPCON) commands can be used to define a DST shadow volume pair for two shared NSS volumes. For information about managing NCP Server for Linux, see the [OES 2018 SP3: NCP Server for Linux Administration Guide](#).

DST is automatically enabled when NCP Server is running and enabled, even if there are no shadow volume pairs currently defined. There is no way to separately turn DST on or off in OES Remote Manager for Linux or in the YaST 2 Runlevel Editor.

Some NCP Server parameters are used to control DST global policies. These parameters apply only for shadow volume pairs, and not to NCP volumes in general. For information, see [Chapter 7, “Configuring DST Global Policies,”](#) on page 49.

3.3 OES Storage Services

Dynamic Storage Technology supports only OES Storage Services (NSS) volumes in shadow volume pairs. You must install OES Storage Services and any other OES Services that it requires. For information, see [“Installing and Configuring OES Storage Services”](#) in the [OES 2018 SP3: NSS File System Administration Guide for Linux](#).

IMPORTANT: Some NSS file system features require special handling when using NSS volumes in a DST shadow volume pair. For information, see [Section 9.4, “Using NSS Volumes in DST Shadow Volume Pairs,”](#) on page 73.

3.4 Directory Services

- ♦ [Section 3.4.1, “Active Directory Trustees,”](#) on page 30
- ♦ [Section 3.4.2, “eDirectory 9.2.1 Trustees,”](#) on page 30

3.4.1 Active Directory Trustees

The Active Directory users can now access the data on DST volumes via CIFS. To manage the rights of the Active Directory trustees on the DST volumes, you can use OES File Access Rights Management (NFARM) utility or `rights` utility. For more information, see [Managing the Trustee Rights in the NSS File System](#) and `rights` in the [OES 2018 SP3: NSS File System Administration Guide for Linux](#).

3.4.2 eDirectory 9.2.1 Trustees

Users whose User objects are defined in eDirectory 9.2.1 also have access to data. For information about configuring eDirectory and users, see the [About this Book and the Library](#).

IMPORTANT: All users of data on the shadow volume pair must be eDirectory users. The server's `root` user is the only local user who should access data without authenticating in eDirectory.

eDirectory users who access files via native Linux file access protocols must be enabled for Linux with Linux User Management (LUM), that is, they are *LUM-enabled*. For information, see [Section 3.10, “Linux User Management,” on page 32](#).

3.5 NCP File Access

NCP (NetWare Core Protocol) client users and applications automatically see a merged view of the files and subdirectories on the shadow volume pair when they access a share on the primary volume.

For information about managing NCP Server for Linux, see the [OES 2018 SP3: NCP Server for Linux Administration Guide](#).

3.6 OES CIFS File Access

OES CIFS is supported to give CIFS users access to the data on DST shadow volume pairs that are built with NSS volumes. CIFS users automatically see a merged view of the data by accessing a CIFS share on the primary volume. For information about configuring and managing OES CIFS, see the [OES 2018 SP3: OES CIFS for Linux Administration Guide](#). For planning information, see “OES CIFS” on page 69.

OES CIFS is supported as an alternative to the Novell Samba solution. OES CIFS does not require ShadowFS, FUSE, or LUM.

3.7 Novell Samba File Access with ShadowFS, FUSE, and LUM

ShadowFS and FUSE can be used to provide a merged view for Novell Samba users. Novell Samba and OES CIFS are mutually exclusive; they cannot be installed on the same server. The merged view is available only when an instance of ShadowFS is running. For information, see [Chapter 13, “Using ShadowFS to Provide a Merged View for Novell Samba Users,” on page 145](#).

The SMB/CIFS users must be eDirectory users who are enabled with Linux User Management (LUM). Samba must also be enabled for LUM. For information about using LUM, see [OES 2018 SP3: Linux User Management Administration Guide](#).

3.8 Supported Native Linux File Access Protocols with ShadowFS, FUSE, and LUM

In addition to Novell Samba, ShadowFS and FUSE can be used to provide a merged view for the native Linux file access protocols SSH and OES FTP (Pure-FTPd). The merged view is available only when an instance of ShadowFS is running. For information about using ShadowFS, see [Chapter 13, “Using ShadowFS to Provide a Merged View for Novell Samba Users,”](#) on page 145.

The SSH and Pure-FTPd users must be eDirectory users who are enabled with Linux User Management (LUM). The native Linux file access services must also be enabled with LUM. For information about using LUM, see [OES 2018 SP3: Linux User Management Administration Guide](#).

The merged view access is not supported for other native Linux file access protocols, including HTTP and NFS.

3.9 OES AFP File Access (Not Supported)

OES AFP does not support a merged view of DST shadow volume pairs. The AFP users are able to see only the data that is on the primary volume. Primary or secondary volumes that are used in a DST shadow volume should not be exposed through AFP.

3.10 Linux User Management

Linux User Management is selected and installed automatically when you install NCP Server and Dynamic Storage Technology. LUM is required if you use Novell Samba or supported native Linux file access protocols with DST shadow volume pairs. The eDirectory users and the native Linux service must both be LUM-enabled in order to see a merged view of files in a DST shadow volume pair, and an instance of ShadowFS must be running. For information about how to configure users and services for LUM, see the [OES 2018 SP3: Linux User Management Administration Guide](#).

3.11 OES Cluster Services for Linux

NCP Server and Dynamic Storage Technology support cluster-enabling DST shadow volume pairs with OES Cluster Services. NCP Server and DST are not cluster-aware. They must be installed and configured on each node in the cluster where you plan to fail over DST shadow volume pairs. In addition, some DST volume and policy information must be manually copied to all nodes. For information, see [Section 17.2.10, “Policies for DST Nodes and Volumes in a Cluster,”](#) on page 186.

3.12 OES Remote Manager for Linux

OES Remote Manager for Linux is required for managing NCP Server services and Dynamic Storage Technology. It is selected and installed by default when you install NCP Server and Dynamic Storage Technology.

For information about management options for DST, see [Section 6.1, “Dynamic Storage Technology Plug-In for OES Remote Manager for Linux,”](#) on page 41.

For information about managing OES Remote Manager and using its other features, see the [OES 2018 SP3: OES Remote Manager Administration Guide](#).

3.13 iManager for Linux

iManager for Linux is required for managing eDirectory users. It can also be used to manage the OES Storage Services pools and volumes, OES Cluster Services for Linux clusters and cluster resources, OES CIFS, Novell Samba, and OES Linux User Management.

It is not necessary to install iManager on every server, but it must be installed somewhere in the same eDirectory tree. For information about installing and using iManager, see the [NetIQ iManager Administration Guide](#).

You use the Storage plug-in for iManager to share devices and to create and manage NSS pools and volumes. For information, see “[Storage Plug-In Quick Reference](#)” in the [OES 2018 SP3: NSS File System Administration Guide for Linux](#). You can alternatively use NSSMU to manage NSS pools and volumes.

You use the Clusters plug-in for iManager to manage the DST pool cluster resource, and to modify the resource load script, unload script, and monitor script. For information, see “[Configuring and Managing Cluster Resources](#)” in the [OES 2018 SP3: OES Cluster Services for Linux Administration Guide](#).

3.14 SFCB and CIMOM

The Small Footprint CIM Broker (SFCB) service replaces OpenWBEM for CIMOM activities in OES 11 and later. For information, see “[Small Footprint CIM Broker \(SFCB\)](#)” in the [OES 2018 SP3: Planning and Implementation Guide](#).

IMPORTANT: SFCB service must be running and working properly whenever you modify the settings for a cluster or cluster resources.

Port 5989 is the default setting for secure HTTP (HTTPS) communications. If you are using a firewall, the port must be opened for CIMOM communications.

The storage-related plug-ins for iManager require CIMOM connections for tasks that transmit sensitive information (such as a user name and password) between iManager and the `_admin` volume on the OES 2015 or later server that you are managing. Typically, CIMOM is running, so this should be the normal condition when using the server. CIMOM connections use Secure HTTP (HTTPS) for transferring data, and this ensures that sensitive data is not exposed.

If CIMOM is not currently running when you click **OK** or **Finish** for the task that sends the sensitive information, you get an error message explaining that the connection is not secure and that CIMOM must be running before you can perform the task.

IMPORTANT: If you receive file protocol errors, it might be because SFCB service is not running.

You can use the `rcsblim-sfcb` command to help resolve CIMOM and SFCB service issues:

To perform this task	At a terminal console prompt, enter as the <code>root</code> user
To start SFCB service	<code>rcsblim-sfcb start</code>
To stop SFCB service	<code>rcsblim-sfcb stop</code>
To check SFCB service status	<code>rcsblim-sfcb status</code>
To restart SFCB service	<code>rcsblim-sfcb restart</code>

3.15 Other OES Services

Ensure that you install and configure additional OES services that might be required by each of the other services mentioned in this section. Refer to the individual guides for those services for information about how to install and manage them.

4 Installing Dynamic Storage Technology

This section describes how to install Dynamic Storage Technology on a Open Enterprise Server (OES) server. You can install DST during or after an OES installation. Before you install DST, ensure that you understand the requirements and guidelines for using DST as described in [Chapter 3, “Planning Your Dynamic Storage Technology Server Environment,”](#) on page 29.

IMPORTANT: After you install DST, you must configure the DST global policies as described in [Chapter 7, “Configuring DST Global Policies,”](#) on page 49.

- ♦ [Section 4.1, “Installing DST on a New OES Server,”](#) on page 35
- ♦ [Section 4.2, “Installing DST on an Existing OES Server,”](#) on page 36

4.1 Installing DST on a New OES Server

You can install the **Novell NCP Server / Dynamic Storage Technology** pattern during the OES installation. For general installation instructions for OES, see the [OES 2018 SP3: Installation Guide](#).

- 1 From the boot menu, select **Installation** and press enter, then continue with the installation as desired until you get to the Installation Settings page.
- 2 On the **Installation Settings** page, click **Software** to open the Software Selection and System Tasks page.
- 3 Under **Open Enterprise Server**, select **Novell NCP Server / Dynamic Storage Technology**.
When you select **Novell NCP Server / Dynamic Storage Technology**, the following additional **Open Enterprise Server** patterns are automatically selected:
 - ♦ **Novell Backup / Storage Management Services**
 - ♦ **NetIQ eDirectory**
 - ♦ **OES Linux User Management**
 - ♦ **OES Remote Manager (NRM)**
- 4 Select **OES Storage Services** from the **Open Enterprise Server** patterns.
NSS is required because DST shadow volume pairs are supported only for NSS volumes.
- 5 (Optional) Select **Novell iManager** from the **Open Enterprise Server** patterns.
You must install Novell iManager somewhere in your network, but it is not necessary to install it on every server.
- 6 (Optional) If you plan to configure shared DST shadow volumes in a cluster, select **OES Cluster Services (NCS)** from the **Open Enterprise Server** patterns.

For detailed information about configuring OES Cluster Services after the installation is complete, see [“Configuring OES Cluster Services”](#) in the [OES 2018 SP3: OES Cluster Services for Linux Administration Guide](#).

- 7 (Optional) If you plan to provide access to DST shadow volumes for CIFS users, select one of the following from the **Open Enterprise Server** patterns:
 - ◆ **OES CIFS**
- 8 Click **OK** to accept the software choices and return to the Installation Settings page.
- 9 Continue with the installation.

When the installation is complete, NCP Server and Dynamic Storage Technology run automatically on system restart.
- 10 After the installation is completed, continue with [Chapter 7, “Configuring DST Global Policies,”](#) on page 49.

4.2 Installing DST on an Existing OES Server

You can install the **Novell NCP Server / Dynamic Storage Technology** pattern at any time after the initial OES installation by using **YaST > Open Enterprise Server > OES Install and Configuration**. For general instructions for installing OES services on an existing OES server, see “[Installing or Configuring OES Services on an Existing OES 2018 SP3 Server](#)” in the *OES 2018 SP3: Installation Guide*.

- 1 Log in to the server as the `root` user, then launch YaST 2.
- 2 In YaST, select **Open Enterprise Server > OES Install and Configuration**.

Wait for the Package Manager to load the Software Selection page. The Open Enterprise Server patterns appear at the end of the list.
- 3 On the Software Selection page under **Open Enterprise Server**, select **Novell NCP Server / Dynamic Storage Technology** and any other compatible OES patterns that you want to install.

IMPORTANT: Services that are already installed are indicated by a check mark in a blue status check box next to the service. If a service is already installed, do not select it again.

When you select **Novell NCP Server / Dynamic Storage Technology**, the following additional **Open Enterprise Server** patterns are automatically selected if they are not already installed:

- ◆ **Novell Backup / Storage Management Services**
 - ◆ **NetIQ eDirectory**
 - ◆ **OES Linux User Management**
 - ◆ **OES Remote Manager (NRM)**
- 4 If it is not already installed, select **OES Storage Services** from the **Open Enterprise Server** patterns.

NSS is required because DST shadow volume pairs are supported only for NSS volumes.
 - 5 (Optional) If it is not already installed, select **Novell iManager** from the **Open Enterprise Server** patterns.

You must install Novell iManager somewhere in your network, but it is not necessary to install it on every server.
 - 6 (Optional) If it is not already installed, and if you plan to configure shared DST shadow volume pairs, select **OES Cluster Services (NCS)** from the **Open Enterprise Server** patterns.

For information about configuring the server for clustering, see [“Configuring OES Cluster Services”](#) in the *OES 2018 SP3: OES Cluster Services for Linux Administration Guide*.

- 7 (Optional) If you plan to provide access to DST shadow volumes for CIFS users, select one (not both) of the following protocol solutions from the **Open Enterprise Server** patterns if it is not already installed:
 - ◆ **OES CIFS**
- 8 Click **Accept** to install the selected patterns.

Wait while the patterns’ packages are installed and the server is configured for the added software.
- 9 On the Open Enterprise Server Configuration page, click **Next**.
- 10 After the installation is completed, click **Finish**, then continue with [Chapter 7, “Configuring DST Global Policies,”](#) on page 49.

5 Using DST in a Virtual Environment

Dynamic Storage Technology (DST) for Open Enterprise Server (OES) works similarly on a native hardware environment and on a virtual machine, with the following caveats:

- ◆ DST supports up to 16 shadow volumes and up to 16 ShadowFS volumes in a virtualized guest server environment.
- ◆ DST is not supported for use in the virtualization host server environment.

Limits for the number of devices assigned to a virtual machine:

- ◆ **Para-virtualized:** 16 devices
- ◆ **Fully virtualized:** 4 devices

To get started with Xen virtualization and KVM virtualization, see the [Virtualization Guide \(https://www.suse.com/documentation/sles-12/book_virt/data/book_virt.html\)](https://www.suse.com/documentation/sles-12/book_virt/data/book_virt.html).

For information on setting up virtualized OES, see [Installing, Upgrading, or Updating OES on a VM](#) in the [OES 2018 SP3: Installation Guide](#).

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 you are using.

6 Management Tools for DST

This section provides an overview of the management tools for Dynamic Storage Technology (DST) in Open Enterprise Server (OES).

- ♦ [Section 6.1, “Dynamic Storage Technology Plug-In for OES Remote Manager for Linux,” on page 41](#)
- ♦ [Section 6.2, “NCP Console \(NCPCON\) Commands,” on page 47](#)
- ♦ [Section 6.3, “Management Tools for NSS Volumes,” on page 47](#)
- ♦ [Section 6.4, “Management Tools for Clustering,” on page 48](#)

6.1 Dynamic Storage Technology Plug-In for OES Remote Manager for Linux

The Dynamic Storage Technology Options plug-in to OES Remote Manager for Linux allows you to create, manage, and remove shadow volume pairs built with OES Storage Services (NSS) volumes. The plug-in is automatically installed in OES Remote Manager when you install NCP Server and Dynamic Storage Technology on your OES server.

- ♦ [Section 6.1.1, “Accessing OES Remote Manager,” on page 41](#)
- ♦ [Section 6.1.2, “Starting, Stopping, or Restarting OES Remote Manager on Linux,” on page 42](#)
- ♦ [Section 6.1.3, “Quick Reference for Dynamic Storage Technology Options,” on page 43](#)
- ♦ [Section 6.1.4, “Quick Reference for NCP Server Options,” on page 45](#)
- ♦ [Section 6.1.5, “Quick Reference for DST Global Policy Settings,” on page 46](#)
- ♦ [Section 6.1.6, “Shadow Volume Inventory and Trustee Reports,” on page 46](#)

6.1.1 Accessing OES Remote Manager

You must log in as the server’s root user to manage Dynamic Storage Technology shadow volume pairs and NCP Server parameters for DST global policies.

- 1 In a web browser, go to the IP address of the server that you want to manage. Use the default ports 8008 or 8009, or you can use any special port you configured for OES Remote Manager.

```
http://server_IP_address:8008  
https://server_IP_address:8009
```

For example:

```
http://192.168.123.11:8008  
https://192.168.123.11:8009
```

- 2 Log in to OES Remote Manager as the `root` user of the server.

6.1.2 Starting, Stopping, or Restarting OES Remote Manager on Linux

OES Remote Manager on Linux is installed and runs by default. If it hangs, you can use the `rcnovell-httpstkd status/start/stop/restart` or `systemctl status/start/stop/restart novell-httpstkd.service` to get status, start, to stop, or restart `httpstkd`. For the latest information about `httpstkd`, see “Starting or Stopping HTTPSTKD” in the *OES 2018 SP3: OES Remote Manager Administration Guide*.

- 1 Open a terminal console, then log in as the `root` user.
- 2 At the terminal console prompt, enter the command for the task you need to perform:

Task	Command
Status	<code>rcnovell-httpstkd status</code>
	or
	<code>systemctl status novell-httpstkd.service</code>
Start	<code>rcnovell-httpstkd start</code>
	or
	<code>systemctl start novell-httpstkd.service</code>
Stop	<code>rcnovell-httpstkd stop</code>
	or
	<code>systemctl stop novell-httpstkd.service</code>
Restart	<code>rcnovell-httpstkd restart</code>
	or
	<code>systemctl restart novell-httpstkd.service</code>

6.1.3 Quick Reference for Dynamic Storage Technology Options

The Dynamic Storage Technology Options plug-in (shown in [Figure 6-1](#)) in OES Remote Manager for Linux is the primary tool for configuring global policies for all DST shadow volumes, creating and managing DST shadow volumes, and configuring shadow volume policies. [Table 6-1](#) describes the management tasks available on this page.

Figure 6-1 View File System > Dynamic Storage Technology Options




Table 6-1 View File System > Dynamic Storage Technology Options

Management Tasks	Description
Volume Information	View a list of NCP volumes and NSS volumes on the server and their shadow status.
Add Shadow link	<p>Click the link to go to the NCP volume's Share Information page, then scroll down to the Volume Tasks area to find the Add Shadow Volume task.</p> <p>The Share Information page and Add Shadow Volume page do not distinguish or validate whether the volumes you choose are actually supported file systems and available combinations.</p> <p>IMPORTANT: NSS volumes must already exist when you create the shadow volume. The Create if not present option is available for future support of NCP volumes on Linux file systems. Do not use this option for NSS volumes.</p>

Management Tasks	Description
Inventory link	View statistics and graphical trend displays for the volume's files and directories. For a DST shadow volume pair, the report includes information for both the primary NSS volume and the secondary NSS volume (shadow).
View Log link	View the audit log for the selected volume.
Information (i) icon	<p>View the NCP volume's share information, such as the Linux file system path for the volume, file system type, NCP volume ID, status, capacity, and cache statistics.</p> <p>Open File Information option allows you to view a list of open NCP connections for the selected volume.</p> <p>Add a shadow volume for the selected NSS volume.</p> <p>For unmounted DST shadow volumes, click the Info icon to access the dialog box to remove the shadow volume relationship. This removes the entry in the <code>/etc/opt/novell/ncpserv.conf</code> file, but does not delete the volume itself.</p> <p>To unmount a shadow volume, click Manage NCP Services > Manage Shares, then click the Unmount option next to the shadow volume.</p>
Dynamic Storage Technology Policies	<p>View a list of existing policies.</p> <p>Click the Policy Name link to modify or delete the policy.</p>
Create a new policy	Click this option to define a new volume policy. For information about policy options, see Section 11.1, "Understanding Shadow Volume Policy Options," on page 113.
Stop all running policies	Click this option to stop all volume policies that are currently running. For information, see Section 11.7, "Stopping a Running Policy," on page 126.
Duplicate File Resolution Options	<p>Select (default) or deselect the option to broadcast messages about duplicate files to users.</p> <p>Select the action to be taken when duplicate files are present on the secondary volume (shadow area), then click Submit.</p> <ul style="list-style-type: none"> ◆ Show duplicate shadow files (the default) ◆ Hide duplicate shadow files ◆ Rename duplicate shadow files ◆ Delete duplicate files from shadow area ◆ Move duplicate shadow files to <code>_DUPLICATE_FILES</code>

Management Tasks	Description
ShadowFS Configuration	<p>Select or deselect (default) Load ShadowFS, then click Submit. This option starts and instance of ShadowFS and configures it to start automatically on system restart.</p> <p>Use ShadowFS only if you need to provide merged view access to users via supported native Linux file access protocols, such as Novell Samba (SMB/CIFS).</p> <p>IMPORTANT: OES CIFS does not require ShadowFS.</p>

6.1.4 Quick Reference for NCP Server Options

Table 6-2 describes the DST tasks available for the **Manage NCP Services > Manage Shares** task in OES Remote Manager for Linux. For a complete list of NCP Server management tasks, see “[Quick Reference for the NCP Server Plug-In for OES Remote Manager for Linux](#)” in the *OES 2018 SP3: NCP Server for Linux Administration Guide*.

Table 6-2 *Manage NCP Services > Manage Shares*

Subtasks	Management Tasks
NCP/NSS Bindings	In the Configuration area, click NCP/NSS Bindings to view a list of NSS volumes on the server. Set the NCP Accessible setting to No for NSS volumes that you want to use as secondary storage locations for DST shadow volumes.
Mount/Unmount	Mount or unmount the primary volume for a shadow volume. The primary volume must be unmounted in order to access the Remove Shadow Volumes task.
Info > Remove Shadow Volume	After you unmount a DST shadow volumes, click the Info icon to access the dialog box to remove the shadow volume relationship. This removes the entry in the <code>ncpserv.conf</code> file, but does not delete the two volumes and their data.

6.1.5 Quick Reference for DST Global Policy Settings

Table 6-3 describes the DST parameters available for the **Manage NCP Services > Manage Server** task in OES Remote Manager for Linux. For descriptions of the parameters, see [Section A.4.1, “Understanding DST Parameters for the SET Command,”](#) on page 233.

Table 6-3 *Manage NCP Services > Manage Server > Server Parameter Information*

Parameter Name	Default Value	Valid Values
SHIFT_MODIFIED_SHADOW_FILES	1	0 - Disable 1 - Allow
SHIFT_ACCESSED_SHADOW_FILES	0	0 - Disable 1 - Allow
SHIFT_DAYS_SINCE_LAST_ACCESS	1	0 - Disable 1 to 365 (in days)
DUPLICATE_SHADOW_FILE_ACTION	0	0 - Show duplicate shadow files (default) 1 - Hide duplicate shadow files 2 - Rename duplicate shadow files 3 - Delete duplicate files from shadow area 4 - Move duplicate shadow files to / . _DUPLICATE_FILES
DUPLICATE_SHADOW_FILE_BROADCAST	1	0 - Disable 1 - Allow
REPLICATE_PRIMARY_TREE_TO_SHADOW	1	0 - Disable 1 - Allow

6.1.6 Shadow Volume Inventory and Trustee Reports

In OES Remote Manager, the Volume Inventory feature detects shadow volumes and displays information from the primary and secondary volumes. The complete inventory profile displays three categories of information: combined areas, primary area, and shadow area. With OES Remote Manager’s shadow volume inventory, you can also select files that meet specific criteria (such as files that have not been accessed for two years, files that have not been modified in a year, all .mp3 files, and so on). Use the inventory information to profile each area’s files and move them as needed.

For general information about the volume inventory feature, see [“Generating Inventories for Directories or NCP Volumes”](#) in the *OES 2018 SP3: OES Remote Manager Administration Guide*.

OES Remote Manager also allows you to generate a trustee report for the shadow volume. For information, see [“Generating and Viewing NCP Trustee Reports for NSS Volumes”](#) in the *OES 2018 SP3: OES Remote Manager Administration Guide*.

6.2 NCP Console (NCPCON) Commands

You can optionally use the NCP Console (NCPCON, `ncpcon(8)` command) to manage Dynamic Storage Technology pairs from a terminal console. For information, see [Section A.1, “Using NCPCON Commands for DST,”](#) on page 225.

6.3 Management Tools for NSS Volumes

- ◆ [Section 6.3.1, “OES File Access Rights Management \(NFARM\),”](#) on page 47
- ◆ [Section 6.3.2, “Storage Plug-In for iManager 3.2.1,”](#) on page 47
- ◆ [Section 6.3.3, “Files and Folders Plug-In for iManager 3.2.1,”](#) on page 47
- ◆ [Section 6.3.4, “NSS Management Utility \(NSSMU\),”](#) on page 48
- ◆ [Section 6.3.5, “Novell Linux Volume Manager \(NLVM\) Commands,”](#) on page 48

Run the following commands before upgrading the volumes to OES 2018:

- 1 `ncpcon set REPLICATE_PRIMARY_TREE_TO_SHADOW=1.`
- 2 `ncpcon set SYNC_TRUSTEES_TO_NSS_AT_VOLUME_MOUNT=1.`
- 3 Remount the volumes.
 - 3a `ncpcon dismount <vol name>.`
 - 3b `ncpcon mount <vol name>.`

6.3.1 OES File Access Rights Management (NFARM)

You can use the NFARM utility or `rights` utility to manage the rights of Active Directory trustees on the NSS volumes that you use as DST shadow volumes. For more information, see [“Managing the Trustee Rights in the NSS File System”](#) and [“rights”](#) in the *OES 2018 SP3: NSS File System Administration Guide for Linux*.

6.3.2 Storage Plug-In for iManager 3.2.1

Use the Storage plug-in for iManager to create and manage OES Storage Services (NSS) volumes that you use as DST shadow volumes. You can enable the User Quotas attribute and to set user quotas on the primary volume and secondary volume. You can enable the Directory Quotas attribute for the primary volume. For information, see [“iManager and Storage-Related Plug-Ins”](#) in the *OES 2018 SP3: NSS File System Administration Guide for Linux*.

6.3.3 Files and Folders Plug-In for iManager 3.2.1

Use the Files and Folders plug-in for iManager to manage file system trustees, trustee rights, and inherited rights filters for files and directories on NSS volumes that you use as DST shadow volumes. You can also set file ownership, directory quotas, and file system attributes. For information, see [“Files and Folders Plug-In Quick Reference”](#) in the *OES 2018 SP3: NSS File System Administration Guide for Linux*.

6.3.4 NSS Management Utility (NSSMU)

You can also use the NSS Management Utility (NSSMU, `nssmu(8)`) to create and manage NSS volumes that you use in DST shadow volumes. For information, see “[NSS Management Utility \(NSSMU\) Quick Reference](#)” in the *OES 2018 SP3: NSS File System Administration Guide for Linux*.

6.3.5 Novell Linux Volume Manager (NLVM) Commands

You can also use the Novell Linux Volume Manager (NLVM, `nlvm(8)`) commands to create and manage NSS volumes that you use in DST shadow volumes. For information, see “[NLVM Commands](#)” in the *OES 2018 SP3: NLVM Reference*.

6.4 Management Tools for Clustering

Use the Clustering plug-in for iManager to create and manage the cluster resources, load scripts, and unload scripts for clustered NSS pools that contain the NSS volumes you use as DST shadow volumes. For information, see “[Creating Cluster Resources](#)” in the *OES 2018 SP3: OES Cluster Services for Linux Administration Guide*.

7 Configuring DST Global Policies

Dynamic Storage Technology allows you to configure global policies that govern DST behavior for all shadow volume pairs on a server. Global policies are special NCP Server parameters that apply only to shadow volume pairs, and not to NCP volumes in general. Before you configure DST shadow volume pairs, ensure that you understand the DST global policies and how their settings affect DST behavior. You can modify the settings as needed.

In a cluster, all nodes must have the same DST global policy settings. When you modify DST's NCP Server parameters settings on one node, ensure that you make the same changes on all nodes before the resource is cluster migrated or failed over to a different node.

- ♦ [Section 7.1, “Replicating Branches of the Primary File Tree in the Secondary File Tree,” on page 49](#)
- ♦ [Section 7.2, “Shifting Files from the Secondary File Tree to the Primary File Tree,” on page 50](#)
- ♦ [Section 7.3, “Resolving Instances of Duplicate Files,” on page 55](#)
- ♦ [Section 7.4, “Loading ShadowFS,” on page 60](#)

7.1 Replicating Branches of the Primary File Tree in the Secondary File Tree

You can create a global policy to control when branches in the primary file tree are replicated to the secondary file tree.

The NSS filesystem is DST aware for new volumes, so the primary volume directory tree structure will always be replicated on the shadow volume directory tree structure. A configurable option called Replicate Primary Tree to Shadow is by default set to 1, so when the existing NSS volumes are configured as primary volume for DST, the primary volume directory tree structure will get replicated to the shadow volume directory tree.

IMPORTANT: If you use shadow volumes in a cluster, ensure that you set the same global policies on each OES node in the cluster.

Valid settings for the Replicate Primary Tree to Shadow are:

- ♦ **Disabled (0):** Branches in the existing primary volume are not replicated to the new secondary volume when the existing volume is configured as a primary volume of the DST. It is replicated as needed when data is moved from the primary storage area to the secondary storage area.
- ♦ **Enabled (1, default):** Branches in the existing primary volume are replicated to the new secondary volume when the existing volume is configured as a primary volume of the DST.

To configure the Replicate Primary Tree to Shadow parameter:

- 1 Log in as the `root` user to OES Remote Manager.
- 2 Select **Manage NCP Services > Manage Server** to view the **Server Parameter Information**.

- 3 Click the link for the `REPLICATE_PRIMARY_TREE_TO_SHADOW` setting.
- 4 In **New Value**, do one of the following (applicable when the existing volume is configured as primary volume of the DST):
 - ♦ **Disable Immediate Path Replication:** Type 0 to stop replicating the paths in the existing primary volume directory structure to the new secondary volume directory structure. Paths are replicated as they are needed when the data is actually moved to the secondary storage area.
 - ♦ **Allow Immediate Path Replication:** Type 1 to replicate all paths in the existing primary volume directory structure to the new secondary volume directory structure.
- 5 Click **Change**.
- 6 On the **Server Parameter Information** page, verify that the new setting is displayed for the `REPLICATE_PRIMARY_TREE_TO_SHADOW` parameter.

For information about using the `SET` command to modify this global policy, see [Section A.4, “Configuring Global DST Policies by Using the SET Command,”](#) on page 233.

7.2 Shifting Files from the Secondary File Tree to the Primary File Tree

You can configure global policies for how files in the secondary file tree are automatically moved back to the primary volume. By default, files are moved back to the primary if they are modified, but not if they are accessed.

- ♦ [Section 7.2.1, “Understanding Shift Parameters,”](#) on page 50
- ♦ [Section 7.2.2, “Configuring a Global Policy for Shifting Modified Shadow Files,”](#) on page 53
- ♦ [Section 7.2.3, “Configuring a Global Policy for Shifting Accessed Shadow Files,”](#) on page 54
- ♦ [Section 7.2.4, “Configuring a Global Policy for the Days Since Last Access,”](#) on page 54
- ♦ [Section 7.2.5, “Using the SET Command to Set Global Policies,”](#) on page 54

7.2.1 Understanding Shift Parameters

You can control how files are automatically moved from the secondary storage area to the primary storage area by configuring three parameters:

- ♦ Shift Modified Shadow Files
- ♦ Shift Accessed Shadow Files
- ♦ Shift Days Since Last Access

IMPORTANT: If you use shadow volumes in a cluster, ensure that you set the same global policies on each OES node in the cluster.

This section describes the parameters, and recommends combinations of the policies to achieve different goals.

- ♦ [“Shift Modified Shadow Files”](#) on page 51
- ♦ [“Shift Accessed Shadow Files”](#) on page 51

- ♦ [“Shift Days Since Last Access” on page 52](#)
- ♦ [“Use Cases for Shifting Shadow Files” on page 52](#)

Shift Modified Shadow Files

When files in the secondary file tree are modified, a configurable global policy called *Shift Modified Shadow Files* allows the files to be moved to the primary file tree (default), or kept in the secondary file tree. If a file is modified when this parameter is enabled, the file is automatically moved back to the primary storage area when the file is closed. This global policy applies to all DST shadow volumes on a given server. The policy does not move files that were modified prior to the setting being enabled.

Valid settings for Shift Modified Shadow Files are:

- ♦ **Disabled (0):** When a file that resides on the secondary storage area is modified, it remains on the secondary storage area.

IMPORTANT: Applications are not aware that DST stores files in two locations. Depending on how an application works, a file might reside on the secondary storage when it is opened, and reside on the primary storage after it is modified.

For example, when you open a file to modify it, Microsoft Word creates a new temporary file and copies the content to it. It saves any changes in the new file, and deletes the old one. Because DST creates all new files on the primary location, the temporary file is created and saved on the primary storage, and the old file is deleted on the secondary location.

This behavior is not unique to Microsoft applications; other word processors and applications behave in the same fashion. When you plan your solution, you must be aware of how the applications you use actually work. If an application’s behavior overrides your intended data locations in the shadow volume, you can use policies to achieve the desired separation.

-
- ♦ **Enabled (1, default):** If a file that resides on the secondary storage area is modified, it is automatically shifted to the primary storage area after it is closed. The file remains on the primary storage area until a policy is enforced that shifts it to the secondary storage area.

For example, if your policy is to place newer files in the primary file tree and to place older files in the secondary file tree, you want an older file in the secondary file tree to move to primary file tree if the file’s content is modified. The Shift Modified Shadow Files parameter is enabled by default, so this is the default behavior.

On the other hand, if you are placing files of one type (such as `.doc` and `.ppt`) in the primary area and files of a different type (such as `.mp3` and `.jpg`) in the secondary area, you want files to stay where they are whenever they are modified. In this case, you should disable the Shift Modified Shadow Files parameter.

Shift Accessed Shadow Files

When files in the secondary file tree are accessed (but not changed), a configurable global policy called *Shift Accessed Shadow Files* allows the files to be left in the secondary file tree (default), or to be moved to the primary file tree. When this parameter is enabled, a file is shifted if it is accessed as read-only a second time during a specified period of time. The file is automatically moved back to

the primary area when the file is closed. By default, the period of time is 1 day. Use the Shift Days Since Last Access parameter to specify the period of time. This global policy applies to all DST shadow volumes on a given server.

Valid settings for the Shift Accessed Shadow Files are:

- ◆ **Disabled (0, default):** When a file that resides on the secondary storage area is accessed twice in the specified period, it remains on the secondary storage area.
- ◆ **Enabled (1):** If a file that resides on the secondary storage area is accessed twice in the specified period, it is automatically shifted to the primary storage area after it is closed. The file remains on the primary storage area until a policy is enforced that shifts it to the secondary storage area.

For example, if you are placing files that are changing in the primary area and files that are not changing in the secondary area, you want files to stay where they are whenever they are accessed but not changed. The Shift Accessed Shadow Files parameter is disabled by default, so this is the default behavior.

On the other hand, if your policy is to place in-use files in the primary file tree and to place unused files in the secondary file tree, you want an in-use file in the secondary file tree to move to primary file tree if the file is accessed, whether it changes or not. In this case, you should enable the Shift Accessed Shadow Files parameter.

Shift Days Since Last Access

The Shift Days Since Last Access parameter specifies the number of days to use when determining if a file should be moved back to the primary storage area. When it is used with `SHIFT_ACCESSED_SHADOW_FILES`, the parameter sets the time when files are migrated back to the primary storage area after the second access within the specified elapsed time.

Valid settings for the Shift Accessed Shadow Files are:

- ◆ **Disabled (0):** Files are not shifted on access.
- ◆ **Number of Days (1 to 365):** If a file that resides on the secondary storage area is accessed twice in the specified period, it is automatically shifted to the primary storage area after it is closed. The default is 1 day.

Use Cases for Shifting Shadow Files

[Table 7-1](#) describes use cases for shifting files based on the global policies.

Table 7-1 Shift Behaviors for Files in the Secondary File Tree

	Don't Shift Modified Shadow File to Primary	Shift Modified Shadow File to Primary (Default)
Don't Shift Accessed Shadow File to Primary (Default)	<p>Files can be modified or accessed without being shifted to the primary file tree.</p> <p>For example, you can separate files by file type, with the less important files in the secondary area. Thereafter, the files remain where you moved them. You can periodically apply volume-level policies that move file types from the primary to the secondary.</p> <p>Back up the primary area more frequently because it contains the most important file types.</p>	<p>Modified files are shifted to the primary file tree, but accessed files are not. This is the default combination.</p> <p>Separate files so that recently modified files are located in the primary area. Older files remain in the secondary area.</p> <p>Back up the primary area more frequently because it contains all of the recently changed files.</p>
Shift Accessed Shadow File to Primary	<p>Files are shifted when they are accessed twice in a specified period, but not when they are modified.</p> <p>No use case exists for this combination.</p>	<p>Files are shifted when they are modified, or if they are accessed twice in a specified period.</p> <p>This is desirable for migration-on-demand solutions that move data gradually from an old volume to a new, higher-performance location.</p> <p>Unchanged, seldom-used files are available to users, but do not require frequent backups.</p>

7.2.2 Configuring a Global Policy for Shifting Modified Shadow Files

To configure the Shift Modified Shadow Files parameter:

- 1 Log in as the `root` user to OES Remote Manager.
- 2 Select **Manage NCP Services > Manage Server** to view the Server Parameter Information page.
- 3 Click the link for the `SHIFT_MODIFIED_SHADOW_FILES` setting.
- 4 In **New Value**, do one of the following:
 - ◆ **Disable Modified Files from Shifting to Primary:** Type 0 to keep files on the secondary storage area when they are modified.
 - ◆ **Allow Modified Files to Shift to Primary:** Type 1 to shift files on the secondary storage area to the primary storage area when they are modified. This is the default.
- 5 Click **Change**.
- 6 On the Server Parameter Information page, verify that the new setting is displayed for the `SHIFT_MODIFIED_SHADOW_FILES` parameter.

7.2.3 Configuring a Global Policy for Shifting Accessed Shadow Files

To configure the Shift Accessed Shadow Files parameter:

- 1 Log in as the `root` user to OES Remote Manager.
- 2 Select **Manage NCP Services > Manage Server** to view the Server Parameter Information page.
- 3 Click the link for the `SHIFT_ACCESSED_SHADOW_FILES` setting.
- 4 In **New Value**, do one of the following:
 - ♦ **Disable Accessed Files from Shifting to Primary:** Type 0 to keep files on the secondary storage area when they are accessed. This is the default.
 - ♦ **Allow Accessed Files to Shift to Primary:** Type 1 to shift files on the secondary storage area to the primary storage area when they are accessed twice during a specified period.
- 5 Click **Change**.
- 6 On the Server Parameter Information page, verify that the new setting is displayed for the `SHIFT_ACCESSED_SHADOW_FILES` parameter.

7.2.4 Configuring a Global Policy for the Days Since Last Access

To configure the Shift Days Since Last Access parameter:

- 1 Log in as the `root` user to OES Remote Manager.
- 2 Select **Manage NCP Services > Manage Server** to view the Server Parameter Information page.
- 3 Click the link for the `SHIFT_DAYS_SINCE_LAST_ACCESS` setting.
- 4 In **New Value**, do one of the following:
 - ♦ **Disable:** Type 0 to disable this parameter.
 - ♦ **Number of Days:** Type an integer value from 1 to 365 (in days) that specifies the number of days to wait for a second access of a shadow file. If the second access occurs during this period, the file can be moved if the `SHIFT_ACCESSED_SHADOW_FILES` parameter is also enabled.
- 5 Click **Change**.
- 6 On the Server Parameter Information page, verify that the new setting is displayed for the `SHIFT_DAYS_SINCE_LAST_ACCESS` parameter.

7.2.5 Using the SET Command to Set Global Policies

For information about using the `SET` command to modify these global policies, see [Section A.4, “Configuring Global DST Policies by Using the SET Command,”](#) on page 233.

7.3 Resolving Instances of Duplicate Files

You might want to change the default policies for how duplicate files are resolved for DST shadow volumes.

- ◆ [Section 7.3.1, “Understanding Conflict Resolution for Duplicate Files,” on page 55](#)
- ◆ [Section 7.3.2, “Configuring a Global Policy for Actions to Resolve Duplicate Files Conflicts,” on page 58](#)
- ◆ [Section 7.3.3, “Enabling or Disabling Broadcast Messages for Duplicate Files Conflicts,” on page 58](#)
- ◆ [Section 7.3.4, “Resolving Instances of Duplicate Files in the /._DUPLICATE_FILES Directory,” on page 59](#)

7.3.1 Understanding Conflict Resolution for Duplicate Files

The Duplicate File Resolution policies are designed to handle the case where files with the same name are located in matching directories in both the primary storage location and the secondary storage location. Duplicate files typically are caused by restoring instances of the same file to both the primary storage location and the secondary storage location. If you back up the primary volume more frequently than the secondary volume, the instance of the file that is restored on the primary storage area should be the most current of the two files.

IMPORTANT: If you use shadow volumes in a cluster, ensure that you set the same global policies on each OES node in the cluster.

The following global policies can be set to govern handling of duplicate files for all shadow volumes on the server:

- ◆ [“Handling Instances of Duplicate Files” on page 55](#)
- ◆ [“Broadcasting Conflict Messages to NCP Users” on page 57](#)
- ◆ [“Recommended Policy Settings for Duplicate Files Conflict Resolution” on page 57](#)

Handling Instances of Duplicate Files

[Table 7-3](#) describes the options for handling duplicate instances of files. For information about configuring the **Actions to be taken** parameter, see [Section 7.3.2, “Configuring a Global Policy for Actions to Resolve Duplicate Files Conflicts,” on page 58](#).

Table 7-2 Actions for Duplicate File Resolution

Parameter Options	User View	Resolution
Show duplicate shadow files (default)	The file name appears twice in directory listings.	The administrator or user manually renames one of the files so the system can tell them apart. The user should then determine whether or not to delete one of the instances, and which instance to delete.
Hide duplicate shadow files	Only one instance of the file name is displayed in the directory listings. Client file operations are directed to the instance located on the primary area. If the client deletes the file, the instance in the primary area is deleted, and the instance in the secondary area is then visible.	The users are not aware that a conflict exists. However, the user might see files randomly reappear after they delete a file.
Rename duplicate shadow files	Automatically renames the duplicate file located on the secondary area by adding a unique extension to the name.	Both instances of the file (the file on the primary area and the renamed file on the secondary area) appear in directory listings. The user needs to be informed that such instances might occur so the user can determine which file instance to keep.
Delete duplicate files from the shadow area	Automatically deletes duplicate files located on the secondary storage area.	The users are not aware that a conflict exists. Because duplicate files are typically caused by restoring instances of the same file to both the primary and secondary areas, the instance located on the primary area should be the most current of the two.
Move duplicate shadow files to / ._DUPLICATE_FILES	Causes the duplicate file located on the secondary storage area to be moved to the / ._DUPLICATE_FILES directory at the root of the secondary volume. If there is a file name conflict in the destination directory, then a unique extension is also added to the file name.	The users are not aware that a conflict exists. This option is less risky than automatically deleting duplicate files. It might require occasional cleanup work to be performed in the /._DUPLICATE_FILES directory.

Broadcasting Conflict Messages to NCP Users

DST leverages the broadcast message capability of NCP Server for Linux. You can disable the broadcast messages option in DST if you choose not to broadcast messages when duplicate files are discovered. If the option is enabled, the messages are received only by client versions that support broadcast messages, and only if the client itself has broadcast messages enabled.

If the option is enabled, a message is broadcast by default to NCP users of the file, whenever duplicate file conflicts occur.

There are two prerequisites for using broadcast messages:

- ♦ **NCP Server:** NCP Server must be configured to support broadcast messages by setting the Disable Broadcast parameter for the `SET` command to 0 (disabled).
- ♦ **Client for Open Enterprise Server:** The Client for Open Enterprise Server version being used by the NCP users must be capable of receiving broadcast messages, and the client must be configured to receive broadcast messages.

The broadcast message capability is called Send Message in the Client for Open Enterprise Server.

For information about configuring the Broadcast Conflict Messages to Users parameter, see [Section 7.3.3, “Enabling or Disabling Broadcast Messages for Duplicate Files Conflicts,”](#) on page 58.

Recommended Policy Settings for Duplicate Files Conflict Resolution

The settings for broadcasting messages and handling files are configured separately. [Table 7-3](#) summarizes recommendations for combining the two. However, if users are SMB/CIFS users who cannot receive broadcast messages, or if the version of the Client for Open Enterprise Server that is in use does not support receiving broadcast messages, you should simply disable broadcast, and select an action that makes sense in your environment.

For information, see [“Handling Instances of Duplicate Files”](#) on page 55.

Table 7-3 Recommended Global Policies for Duplicate Files Resolution

Action to be Taken	Broadcast Conflict Messages to Users
Show duplicate shadow files (default)	Enable broadcast (default)
Hide duplicate shadow files	Disable broadcast
Rename duplicate shadow files	Optionally enable broadcast
Delete duplicate files from the shadow area	Disable broadcast
Move duplicate shadow files to / ._DUPLICATE_FILES	Disable broadcast

7.3.2 Configuring a Global Policy for Actions to Resolve Duplicate Files Conflicts

You can set a global policy for the actions to be taken to resolve duplicate file conflicts.

By default, the **Actions to be taken** parameter is set to show duplicate shadow files to the user. For information about the other options, see [“Handling Instances of Duplicate Files”](#) on page 55.

For information about using the `SET` command to modify this global policy, see [Section A.4, “Configuring Global DST Policies by Using the SET Command,”](#) on page 233.

- 1 In OES Remote Manager for Linux, select **View File System**, then select **Dynamic Storage Technology Options**.
- 2 In the **Duplicate File Resolution Options** area, view the current setting for **Actions to be taken**.
- 3 From the **Actions to be taken** drop-down list, select one of the following options:
 - ◆ **Show duplicate shadow files** (default)
 - ◆ **Hide duplicate shadow files**
 - ◆ **Rename duplicate shadow files**
 - ◆ **Delete duplicate files from shadow area**
 - ◆ **Move duplicate shadow files to** / `._DUPLICATE_FILES`
- 4 In the **Duplicate File Resolution Options** area, click **Submit** to save and apply the change.

7.3.3 Enabling or Disabling Broadcast Messages for Duplicate Files Conflicts

You can set a global policy that enables or disables broadcast messages to be sent to NCP clients when duplicate file conflicts are detected.

When **Broadcast Conflict Messages to Users** is enabled (the default setting), a message is broadcast to NCP users of the file when duplicate instances of the file occur on the primary storage location and secondary storage location. For information, see [“Broadcasting Conflict Messages to NCP Users”](#) on page 57 and [“Recommended Policy Settings for Duplicate Files Conflict Resolution”](#) on page 57.

IMPORTANT: In order for users to be able to receive the duplicate-file-conflict messages, NCP Server must be configured to support broadcast messages and the Clients for Open Enterprise Server must be configured to receive broadcast messages. For instructions, see [“Enabling or Disabling Broadcast Message Support”](#) in the *OES 2018 SP3: NCP Server for Linux Administration Guide*.

For information about using the `SET` command to modify this global policy, see [Section A.4, “Configuring Global DST Policies by Using the SET Command,”](#) on page 233.

- 1 In OES Remote Manager for Linux, select **View File System**, then select **Dynamic Storage Technology Options**.
- 2 In the **Duplicate File Resolution Options** area, enable or disable **Broadcast Conflict Messages to User** by selecting or deselecting the check box next to it. It is enabled by default.

Duplicate File Resolution Options	
Broadcast conflict message to user:	<input checked="" type="checkbox"/>
Action to be taken:	Show duplicate shadow files
<input type="button" value="Submit"/>	

- 3 In the **Duplicate File Resolution Options** area, click **Submit** to save and apply the change.
- 4 If you enabled Broadcast Conflict Messages, ensure that NCP Server is configured to support broadcast messages by verifying that the Disable Broadcast (DISABLE_BROADCAST) parameter for the SET command is disabled.
 - 4a In OES Remote Manager for Linux, select **Manage NCP Services**, then select **Manage Server**.
 - 4b In the **Set Parameter Information** table, locate the DISABLE_BROADCAST parameter, then view the current value of the parameter. By default, the parameter is disabled (set to 0), which means that NCP Server supports broadcast messages.

DISABLE_BROADCAST	0
-------------------	-------------------

- 4c If the DISABLE_BROADCAST parameter is enabled (set to 1), click the link for the value in the **Parameter Value** column to open a page where you can change the value.

DISABLE_BROADCAST	1
-------------------	-------------------

- 4d In **New Value**, type 0, then click **Change** to save and apply the settings that disable the DISABLE_BROADCAST parameter, which enables broadcasting for NCP Server.

IMPORTANT: Messages are received only by logged-in users who are using Client for Open Enterprise Server versions that are capable of receiving broadcast messages, and that are configured to receive them.

DISABLE_BROADCAST	
Current Value	New Value
1	<input type="text" value="0"/> <input type="button" value="Change"/>
<input type="button" value="Back"/>	

7.3.4 Resolving Instances of Duplicate Files in the /._DUPLICATE_FILES Directory

If you enable **Move duplicate shadow files to /._DUPLICATE_FILES** as the action to be taken when duplicate file conflicts occur, it might require occasional cleanup work to be performed in the /._DUPLICATE_FILES directory.

7.4 Loading ShadowFS

ShadowFS is required to provide a merged view access to DST volumes for users via SMB/CIFS or native Linux file access. You can start it manually, or set a policy to automatically load ShadowFS at boot time. For information about using and managing ShadowFS, see [Chapter 13, “Using ShadowFS to Provide a Merged View for Novell Samba Users,”](#) on page 145.

By default, ShadowFS and FUSE are not started automatically at boot time. You can set the **ShadowFS Configuration > Load ShadowFS** option that starts them at boot time. It also starts them when you first enable the global policy if they are not already running. A single instance of ShadowFS should be running on a server. The setting applies for all DST volumes on the server.

IMPORTANT: If you use shadow volumes in a cluster, ensure that you set the same global policies on each OES 2018 node in the cluster.

- ◆ [Section 7.4.1, “Using OES Remote Manager to Set the Autostart of ShadowFS,”](#) on page 60
- ◆ [Section 7.4.2, “Using the Command Line to Set the Autostart of ShadowFS,”](#) on page 60
- ◆ [Section 7.4.3, “Manually Starting and Stopping ShadowFS,”](#) on page 61

7.4.1 Using OES Remote Manager to Set the Autostart of ShadowFS

The **Load ShadowFS** option in OES Remote Manager provides a GUI interface for setting up ShadowFS to start at boot time. It also starts ShadowFS when the option is first enabled if ShadowFS is not already running.

- 1 In OES Remote Manager for Linux, select **View File System**, then select **Dynamic Storage Technology Options**.
- 2 In the **ShadowFS Configuration** area, view the current setting for **Load ShadowFS**.
- 3 Enable or disable **Load ShadowFS** by selecting or deselecting the check box.
- 4 In the **ShadowFS Configuration** area, click **Submit** to save and apply the change.

7.4.2 Using the Command Line to Set the Autostart of ShadowFS

You can set the service to autostart upon future reboots at the command line instead of using OES Remote Manager:

- 1 Log in as the `root` user, then open a terminal console.
- 2 Do one of the following:
 - ◆ **Enable Autostart:** Enter the following at a command prompt to enable the autostart of `novell-shadowfs.service`:

```
systemctl enable novell-shadowfs.service
```
 - ◆ **Disable Autostart:** Enter the following at a command prompt to disable the autostart of `novell-shadowfs.service`:

```
systemctl disable novell-shadowfs.service
```

7.4.3 Manually Starting and Stopping ShadowFS

Only one instance of ShadowFS should be loaded at a time. Before you attempt to manually start ShadowFS, ensure that you have stopped any running instances of it.

- 1 Log in as the `root` user, then open a terminal console.
- 2 Do one of the following:
 - ♦ **Start:** Enter the following at a command prompt to start ShadowFS:

```
systemctl start novell-shadowfs.service
```
 - ♦ **Stop:** Enter the following at a command prompt to stop ShadowFS:

```
systemctl stop novell-shadowfs.service
```


8

Managing Services for DST

The health of Dynamic Storage Technology depends on other services that are running on the Open Enterprise Server (OES) server. This section identifies those dependencies and provides instructions for how to get them started again if they are not running.

- ♦ [Section 8.1, “Restarting the Novell NCP/NSS IPC \(ncp2nss\) Daemon,” on page 63](#)
- ♦ [Section 8.2, “Restarting the eDirectory \(ndsd\) Daemon,” on page 63](#)
- ♦ [Section 8.3, “Starting and Stopping ShadowFS,” on page 64](#)

8.1 Restarting the Novell NCP/NSS IPC (ncp2nss) Daemon

If NSS is installed, NCP Server runs the Novell NCP/NSS IPC daemon in order to synchronize its settings with NSS. You must restart `ncp2nss` whenever you modify NCP Server settings or the Dynamic Storage Technology global parameter settings for NCP Server. by using any of the following methods:

- ♦ OES Remote Manager for Linux
- ♦ `ncpcon set` commands
- ♦ Editing the values in the `/etc/opt/novell/ncpserv.conf` file

The changes do not take effect until you have restarted `ncp2nss` and eDirectory (see [Section 8.2, “Restarting the eDirectory \(ndsd\) Daemon,” on page 63](#)).

To restart the `ncp2nss` daemon:

- 1 Log in as the `root` user, then open a terminal console.
- 2 At the terminal console prompt, enter

```
systemctl restart ncp2nss.service
```

8.2 Restarting the eDirectory (ndsd) Daemon

You must restart the eDirectory (`ndsd`) daemon if you modify NCP Server settings or the Dynamic Storage Technology global parameter settings for NCP Server. You can modify the settings in OES Remote Manager, or by editing the `/etc/opt/novell/ncpserv.conf` file. The changes do not take effect until you restart the eDirectory daemon.

IMPORTANT: Restarting or stopping `ndsd` automatically warns the users and terminates their connections. Users can reconnect to the mapped NCP share after the `ndsd` service restarts.

Use the following steps to stop and start `ndsd` when a single instance is running:

- 1 Log in as the `root` user, then open a terminal console.
- 2 Enter one of the following commands to stop `ndsd`:

```
rcndsd stop
```

or

```
systemctl stop ndsd.service
```

- 3 Enter one of the following commands to start `ndsd`:

```
rcndsd start
```

or

```
systemctl start ndsd.service
```

8.3 Starting and Stopping ShadowFS

ShadowFS must be running in order to provide a merged view to SMB/CIFS users if you are using Novell Samba to provide file access to a DST shadow volume. You might also use ShadowFS to provide a merged view for supported native file access protocols, such as SSH and OES FTP.

To enable ShadowFS to start up automatically on system restart, see [Section 7.4, “Loading ShadowFS,”](#) on page 60.

Only one instance of ShadowFS should be loaded at a time. Before you attempt to manually start ShadowFS, ensure that you have stopped any running instances of it.

- 1 Log in as the `root` user, then open a terminal console.
- 2 Enter the following command to stop ShadowFS:

```
systemctl stop novell-shadowfs.service
```

- 3 Enter the following command to start ShadowFS:

```
systemctl start novell-shadowfs.service
```


9 Planning for DST Shadow Volume Pairs and Policies

This section describes guidelines for using Dynamic Storage Technology (DST) shadow volume pairs and policies on Open Enterprise Server (OES) servers. Use this section to understand the expected behavior of shadow volume pairs and policies with commonly used services and storage features.

For installation requirements, see [Chapter 3, “Planning Your Dynamic Storage Technology Server Environment,”](#) on page 29.

- ♦ [Section 9.1, “Storage Requirements for DST Volume Pairs,”](#) on page 65
- ♦ [Section 9.2, “Data Access Requirements for a DST Shadow Volume Pair,”](#) on page 67
- ♦ [Section 9.3, “Guidelines for Working with DST Shadow Volume Pairs,”](#) on page 71
- ♦ [Section 9.4, “Using NSS Volumes in DST Shadow Volume Pairs,”](#) on page 73
- ♦ [Section 9.5, “Using Trustees, Trustee Rights, and File System Attributes on DST Shadow Volume Pairs,”](#) on page 77
- ♦ [Section 9.6, “Using NSS Encrypted Volumes in a DST Shadow Volume Pair,”](#) on page 78
- ♦ [Section 9.7, “Using NSS Quotas on DST Shadow Volume Pairs,”](#) on page 78
- ♦ [Section 9.8, “Using Opportunistic File Locking on DST Shadow Volume Pairs,”](#) on page 80
- ♦ [Section 9.9, “Using OES Cluster Services with DST Shadow Volume Pairs,”](#) on page 80
- ♦ [Section 9.10, “Using Novell Distributed File Services with DST Shadow Volume Pairs,”](#) on page 81
- ♦ [Section 9.11, “Using OES Storage Services Auditing Client Logger \(VLOG\) Utility with DST Shadow Volume Pairs,”](#) on page 82
- ♦ [Section 9.12, “Using Virus Checking Utilities with DST Shadow Volume Pairs,”](#) on page 82
- ♦ [Section 9.13, “Using Backup Utilities with DST Shadow Volume Pairs,”](#) on page 82

9.1 Storage Requirements for DST Volume Pairs

- ♦ [Section 9.1.1, “Storage Devices,”](#) on page 66
- ♦ [Section 9.1.2, “iSCSI Block Storage Devices,”](#) on page 66
- ♦ [Section 9.1.3, “File Systems,”](#) on page 66

9.1.1 Storage Devices

Volumes in a shadow volume pair can reside on any device that is seen as local to the DST server, such as direct-attached storage and Fibre Channel SAN devices. For information about using iSCSI SAN devices, see [Section 9.1.2, “iSCSI Block Storage Devices,” on page 66](#).

Block storage devices used for the primary and secondary storage can have different performance characteristics. Typically, the secondary storage area is slower and less expensive.

9.1.2 iSCSI Block Storage Devices

Dynamic Storage Technology supports using target iSCSI block storage devices to store the primary and secondary volumes in a shadow volume pair. Any iSCSI block storage device should work in a shadow volume pair, if it is compatible with the Linux iSCSI initiator software running on the OES server where you create and manage the shadow volume pair. However, only iSCSI targets running on the following OES servers (or later versions) have been tested and are supported for OES 11 SP1 and later:

- ◆ OES 11 or later
- ◆ OES 2 SP3

IMPORTANT: Third-party iSCSI solutions have not been tested, so they are not supported.

For information about setting up iSCSI target devices on OES servers, see [“Setting Up an iSCSI LIO Target Server”](#) (https://www.suse.com/documentation/sles-12/stor_admin/data/sec_iscsi_target.html) in the *SLES 12 Storage Administration Guide* (https://www.suse.com/documentation/sles-12/stor_admin/data/stor_admin.html).

The iSCSI targets are connected by using the Linux iSCSI initiator software on the OES server where you are creating DST shadow volumes. For information about configuring iSCSI initiators and discovering iSCSI target devices, see [“Configuring iSCSI Initiator”](#) (https://www.suse.com/documentation/sles-12/stor_admin/data/sec_iscsi_initiator.html) in the *SLES 12 Storage Administration Guide* (https://www.suse.com/documentation/sles-12/stor_admin/data/stor_admin.html).

IMPORTANT: OES 2015 or later does not support running iSCSI target software and initiator software on the same server.

9.1.3 File Systems

Dynamic Storage Technology supports using two OES Storage Services (NSS) volumes in a shadow volume pair.

IMPORTANT: Mixing file system types for the primary storage area and secondary storage area in a given DST shadow volume pair is not supported.

For information about using NSS features and file system attributes while the NSS volumes are in a shadow volume pair, see [Section 9.4, “Using NSS Volumes in DST Shadow Volume Pairs,” on page 73](#).

9.2 Data Access Requirements for a DST Shadow Volume Pair

All user and administration access to data is intended to be through the merged view of DST shadow volume pair. Users should never have direct access to the secondary volume. Consider the guidelines in this section when planning how to provide access to the merged view of data in the DST shadow volume.

- ♦ [Section 9.2.1, “Administrator Access,” on page 67](#)
- ♦ [Section 9.2.2, “User Access and Authorization,” on page 68](#)
- ♦ [Section 9.2.3, “File Access Protocols,” on page 68](#)
- ♦ [Section 9.2.4, “ShadowFS and FUSE,” on page 71](#)

9.2.1 Administrator Access

Administrators should manage files, directories, and authorization via the merged view of the DST shadow volume pair. Ensure that you use tools that support the merged view, such as the NCP client and tools described in [Chapter 6, “Management Tools for DST,” on page 41](#).

File system access rights are based on the OES Trustee Model just as they are for a single NSS volume. You must connect to the share on the primary volume by using the Client for Open Enterprise Server or similar NCP client tools, and then set permissions while working in the merged view of the data. The permissions for data on both locations are saved on the primary volume. Dynamic Storage Technology uses those permissions to control access to data stored on the secondary volume. The XML file that contains trustees and trustee rights settings is copied to the secondary volume.

All users (except the `root` user) of the DST volume pair must have User objects defined in eDirectory. NSS volumes by design are not visible to any locally defined user except the `root` user when accessing files natively from the server. The `root` user of the server is the only local user who has direct access to the Linux file structure view of the two volumes.

All the Active Directory users can now access the data on DST volumes via CIFS. To manage the rights of the Active Directory trustees on the DST volumes, you can use OES File Access Rights Management (NFARM) utility or `rights` utility. For more information, see [Managing the Trustee Rights in the NSS File System](#) and `rights` in the [OES 2018 SP3: NSS File System Administration Guide for Linux](#).

Some management tasks are performed as the `root` user; they require direct access to the secondary volume:

- ♦ Volume backup and restore

For information, see [Section 9.13, “Using Backup Utilities with DST Shadow Volume Pairs,” on page 82](#) and [Section 10.11, “Backing Up DST Shadow Volumes,” on page 106](#).

- ♦ Duplicate files management

For information, see [Section 7.3, “Resolving Instances of Duplicate Files,” on page 55](#).

9.2.2 User Access and Authorization

All user file access to data on the DST volume pair is done via the merged view. Users connect to a share on the primary volume to see the merged view.

File system access rights are based on the OES Trustee Model just as they are for a single NSS volume. Trustees and trustee rights are set from the merged view of the DST shadow volume pair. DST enforces those settings whether the file resides on the primary volume or secondary volume.

All the Active Directory users can now access the data on DST volumes via CIFS. To manage the rights of the Active Directory trustees on the DST volumes, you can use OES File Access Rights Management (NFARM) utility or `rights` utility.

All eDirectory users (except the `root` user) of the DST volume pair must have User objects defined in eDirectory.

9.2.3 File Access Protocols

A merged user view of the file system is available for both NCP and CIFS users. The CIFS access can be set up by using OES CIFS or using Novell Samba, but OES does not support using both of these CIFS user access solutions on the same server. In this guide, Novell Samba access is referred to as *SMB/CIFS*. NCP Server is required to be installed and running even if all users access data via OES CIFS or Novell Samba.

This section describes the requirements and guidelines for file access protocols:

- ◆ [“Cross-Protocol File Locking” on page 68](#)
- ◆ [“NCP” on page 68](#)
- ◆ [“OES CIFS” on page 69](#)
- ◆ [“Novell Samba with ShadowFS and FUSE” on page 70](#)
- ◆ [“OES AFP \(Not Supported\)” on page 70](#)
- ◆ [“Other Linux Protocols \(Supported and Not Supported\)” on page 70](#)

Cross-Protocol File Locking

When users access the files via multiple protocols, you should enable the NCP Server Cross-Protocol File Locks option to protect against data corruption. For information, see [“Configuring Cross-Protocol File Locks for NCP Server”](#) in the *OES 2018 SP3: NCP Server for Linux Administration Guide*.

NCP

The DST Shadow Volumes engine supports file access for NCP users. Users access data via an NCP share on the primary storage location, by using the Client for Open Enterprise Server or other NCP clients. For information about configuring NCP Server for the OES server, see the *OES 2018 SP3: NCP Server for Linux Administration Guide*.

Client for Open Enterprise Server: See the following resources for the latest release of the Client for Open Enterprise Server, which provides NCP access for users on Windows clients:

- ◆ [Client for Open Enterprise Server \(https://www.novell.com/documentation/windows_client/\)](https://www.novell.com/documentation/windows_client/)

Only NCP client versions that are configured to receive broadcast messages are eligible to receive the duplicate file conflict messages. For information, see [Section 7.3.3, “Enabling or Disabling Broadcast Messages for Duplicate Files Conflicts,”](#) on page 58.

NetStorage: NetStorage for Linux has limited use for accessing files on shadow volumes. NetStorage presents a merged view of the data, and the user can see, read, and write files. However, certain management functions (such as getting file properties, setting trustees, and salvaging files) work only if the files are on the primary volume. The user will find that the commands work for some files but not others because they are not aware of where the file is physically stored. For information about using NetStorage, see [OES 2018 SP3: NetStorage Administration Guide for Linux](#).

OES CIFS

OES CIFS supports use of NSS volumes in a DST shadow volume configuration. It supports the following DST features:

- ♦ **Merged View:** OES CIFS works with NCP Server to provide a merged view of the two NSS volumes. CIFS users can access data on both volumes via a OES CIFS share on the primary NSS volume.
- ♦ **Duplicate Files:** CIFS can handle duplicate files, but it does not support the broadcast message notification via NCP. It shows the instance of the file on the primary volume to users. The administrator or user can rename the file so that the secondary instance of the file is again visible. The user can then determine which instance to delete.

If the global policy is set to hide duplicate files, CIFS moves the files on the secondary volume to the `/._DUPLICATE_FILES` folder where the administrator can access them for recovery, if necessary.

- ♦ **Global DST Policies:** When users access or modify files, CIFS honors the global DST policies for moving files from the secondary volume to the primary volume.

Setting up OES CIFS for use with a DST shadow volume is similar to setting up CIFS for an NSS volume. You create the CIFS share on the primary volume, but not on the secondary volume. Enable the cross-protocol file locking parameter for NCP Server on the DST server.

OES CIFS features should work as expected, including cross-protocol file locking. The key difference is that users access the merged view of data in both volumes via the CIFS share on the primary NSS volume. The users do not know that the data is stored on two different volumes.

For information, see the following:

- ♦ [OES 2018 SP3: OES CIFS for Linux Administration Guide](#)
- ♦ “[Configuring Cross-Protocol File Locks for NCP Server](#)” in the [OES 2018 SP3: NCP Server for Linux Administration Guide](#)

Consider the following requirements when configuring OES CIFS for use with DST volumes:

- ♦ OES CIFS supports only NSS volumes on Linux. Thus, OES CIFS can be used only with DST volumes built on NSS volumes.
- ♦ CIFS users access a merged view of the DST shadow volume by using a OES CIFS share on the primary volume. Create a CIFS share on the Primary volume only; delete the share on the secondary (or do not give users rights to access the secondary share).

- ♦ CIFS users don't see broadcast messages if the Broadcast Messages for Duplicate Files Conflicts feature is enabled for Duplicate Files errors. This DST option works only for Client for Open Enterprise Server users as described in [Section 7.3.3, "Enabling or Disabling Broadcast Messages for Duplicate Files Conflicts,"](#) on page 58. If you have only CIFS users and no NCP users, you might as well disable the broadcasting option.
- ♦ If you use OES CIFS with a DST volume in a cluster, you need to add the OES CIFS lines in the load/unload scripts for the DST cluster resource. The differences are described in [Chapter 17, "Configuring DST Shadow Volume Pairs with OES Cluster Services,"](#) on page 179.
- ♦ You cannot configure OES CIFS and the SMB/CIFS setup on the same server. This limitation is derived from the requirement that Novell Samba and OES CIFS cannot be installed on the same server, and is unrelated to DST.

Novell Samba with ShadowFS and FUSE

Novell Samba is supported for providing SMB/CIFS user access to shadow volumes. This Samba version is the standard Linux Samba that has been integrated with eDirectory. For information, see the [OES 2018: Novell Samba Administration Guide](#).

In order for SMB/CIFS users to see a merged view of the shadow volume, you must also set up ShadowFS (Shadow File System) and FUSE (File System in Userspace). See [Chapter 13, "Using ShadowFS to Provide a Merged View for Novell Samba Users,"](#) on page 145 for installation and configuration requirements.

SMB/CIFS users must be Linux-enabled users of the OES server. The Linux Samba service must also be LUM enabled. For information, see the [OES 2018 SP3: Linux User Management Administration Guide](#).

Enable the cross-protocol file locking parameter for NCP Server. For information, see "[Configuring Cross-Protocol File Locks for NCP Server](#)" in the [OES 2018 SP3: NCP Server for Linux Administration Guide](#).

OES AFP (Not Supported)

OES AFP (Apple Filing Protocol) does not support DST shadow volumes. AFP users are able to see only the data that is on the primary volume. Do not create AFP shares on the primary or secondary volumes that are used in a DST shadow volume.

Other Linux Protocols (Supported and Not Supported)

Supported Linux file access protocols include SSH and OES FTP (Pure-FTPd). You must use ShadowFS and FUSE on the system to provide a merged view for SSH and FTP users. In addition, the eDirectory users and the native Linux service must be enabled with Linux User Management (LUM). ShadowFS must be running in order for the users to access the data.

IMPORTANT: This configuration works best when you use only NCP for normal user access, or if you must allow CIFS access, you use Novell Samba instead of OES CIFS.

User access to shadow volumes via all other native Linux protocols (such as HTTP, NFS, and others) is not supported.

9.2.4 ShadowFS and FUSE

The Shadow File System (ShadowFS) works with FUSE (File System in Userspace) to provide a merged view of the shadow volume tree for eDirectory users of native Linux file access protocols, including Novell Samba (SMB/CIFS) and supported Linux file access protocols, such as SSH and OES FTP (Pure-FTPd). ShadowFS requires that eDirectory users and the native Linux service be enabled with Linux User Management (LUM). ShadowFS must be running in order for the SMB/CIFS users to access the data.

When ShadowFS is running, it automatically creates a shadow file system directory for each of the shadow volumes, not just the ones where you plan to allow SMB/CIFS access. SMB/CIFS users see only those volumes where they have file system trustee rights.

ShadowFS requires FUSE (File System in Userspace) to be installed and running. For information, see [Chapter 13, “Using ShadowFS to Provide a Merged View for Novell Samba Users,”](#) on page 145.

9.3 Guidelines for Working with DST Shadow Volume Pairs

Consider the guidelines in this section when working with DST shadow volumes.

- ◆ [Section 9.3.1, “Number of Shadow Volumes per Server,”](#) on page 71
- ◆ [Section 9.3.2, “Data Volumes,”](#) on page 71
- ◆ [Section 9.3.3, “Files and Folders,”](#) on page 72
- ◆ [Section 9.3.4, “File System Trustees and Rights,”](#) on page 73
- ◆ [Section 9.3.5, “File System Management Utilities,”](#) on page 73

9.3.1 Number of Shadow Volumes per Server

DST supports the following number of DST shadow volume pairs per DST server:

- ◆ **Physical server:** 32
- ◆ **Virtual server:** 16

IMPORTANT: This constraint is imposed because of a known defect in FUSE (File System in Userspace).

9.3.2 Data Volumes

DST shadow volumes are intended for use with data volumes that contain unstructured data. Consider the following guidelines for choosing which volumes to use with DST:

- ◆ Do not put system files and application files on DST shadow volumes.
- ◆ Do not create a DST shadow volume for the `_ADMIN` volume.

- ◆ If the volume contains database files, rebuild situations might occur because of the additional latency related to the DST handling, or if the secondary storage area becomes unavailable for any reason.

Policies should exclude directories that contain databases such as those for Novell GroupWise and MySQL. You can alternatively create policies in such a way that they do not affect database files.

9.3.3 Files and Folders

- ◆ New files are created via the merged view and are saved on the primary volume.
- ◆ New folders are created via merged view and are saved on the primary volume and secondary volume immediately.
- ◆ While a file is moved from one area to the other, the file is locked so that clients cannot access it during relocation.
- ◆ A policy cannot move a file between the areas if the file is open.

When DST enforces policies or moves files, the relocation request fails if a user has the file open; only files that are not in use can be moved.

- ◆ Always use the merged view when you perform actions on a file, such as open, close, create, delete, rename, modify, set trustees, set trustee rights, or set inherited rights.

Do not perform these actions by directly accessing files on the secondary location, except when the `root` user backs up files, restores files, or resolves instances of duplicate files.

- ◆ Always use the merged view when you perform actions on a folder, such as open, close, create, delete, rename, modify, set trustees, set trustee rights, or set inherited rights.

Do not perform these actions by directly accessing folders on the secondary location, except when the `root` user backs up files or restores files.

When you rename a folder via the merged view, its name is changed on the instance of the folder in the primary location and the secondary location.

If you rename a folder by directly accessing it on the secondary location, the instance of the folder on the primary location is not renamed or deleted. The renamed folder contains the files that were stored in it when it was renamed. The files appear to have disappeared from the primary instance of the folder.

If you delete a folder by directly accessing it on the secondary location, the instance of the folder on the primary location is not deleted.

- ◆ If you define a DST shadow volume pair with an existing volume as secondary and a new volume as primary, the directory structure is replicated gradually to the primary volume as files are used or moved there with a policy. Empty folders are not automatically replicated on the new primary volume. To ensure that the entire existing directory structure (including empty directories) is replicated on the primary volume, an administrator can enumerate the directory from an NCP client that is mapped to the merged view.

9.3.4 File System Trustees and Rights

When the NCP protocol is used in conjunction with the NSS file system, all native NCP functionality (security, rights, trustees) is preserved in a DST environment. No functionality is lost, and no management patterns are changed.

When you use OES CIFS or Novell Samba, all native CIFS functionality for security, rights, and so on is preserved in a DST environment. The conversion of CIFS ACLs (access control lists) to NSS ACLs based on the POSIX definitions is based on code resident in Samba and is not supported for modification by Novell.

IMPORTANT: The CIFS support of ACLs is offered as-is, and is not modified to take advantage of the expanded management features of NSS file systems.

9.3.5 File System Management Utilities

You can continue using existing file management utilities that currently execute successfully against the designated file systems. DST is transparent to this operation. All file management operations currently available to NSS users through iManager, NSSMU, and OES Remote Manager for Linux function transparently for shadow volumes. File operations and the location of the file are transparent to the NCP and CIFS clients.

9.4 Using NSS Volumes in DST Shadow Volume Pairs

Dynamic Storage Technology supports shadow volumes created with OES Storage Services volumes. NSS volume attributes, features, and actions can behave differently while the volume is in a shadow relationship. Consider the guidelines and caveats in this section when planning your shadow volume solution.

- ♦ [Section 9.4.1, “DST Support for NSS Volume Attributes,” on page 74](#)
- ♦ [Section 9.4.2, “DST Support for NSS Features and Actions,” on page 76](#)

9.4.1 DST Support for NSS Volume Attributes

Ensure that you enable the same NSS volume attributes on both volumes in the shadow relationship to ensure a consistent user experience. For example, if Salvage is enabled for the primary volume but not for the secondary volume, files that are deleted when they reside on the secondary volume are purged immediately, and are not available for salvage.

[Table 9-1](#) describes which NSS volume attributes are supported for use with Dynamic Storage Technology, and any caveats to consider when using them. For information about the volume attributes, see “[Volume Attributes](#)” in the *OES 2018 SP3: NSS File System Administration Guide for Linux*.

Table 9-1 DST Support for NSS Volume Attributes

NSS Volume Attribute	Supported	Caveats
AD Enabled	Yes	For AD users to access an NSS32 or NSS64 volume, <ul style="list-style-type: none">the volume must be part of a pool that is AD media upgradedthe volume must be AD Enabled For more information, see Volume AD-enabling and Understanding Volume Properties in the <i>OES 2018 SP3: NSS File System Administration Guide for Linux</i> .
Allow mount point to be renamed	No	DST does not track the renaming of NSS volumes or their mount points. Before you rename or modify the mount point for an NSS volume, you must remove the shadow volume definition. Afterwards, you can re-create the shadow volume.
Backup	Yes	The Linux file system sees both volumes, so you back up each volume separately.
Compression	Yes	You can set compression on one or both volumes. Compressed files are uncompressed when they are moved from the primary volume to secondary volume, and vice versa. In order for the move to occur, there must be sufficient space on the source volume to allow both the uncompressed and compressed copies of the file to coexist until the move is completed. There must also be sufficient space on the destination volume for the uncompressed file to be stored. The file is re-compressed according to the compression schedule and settings in the destination volume.
Data Shredding	Yes	For security compliance reasons, you should set this attribute on both volumes if you use it.
Directory Quotas	Yes	Set a directory quota for a directory only on the primary volume. For more information, see Section 9.7, “Using NSS Quotas on DST Shadow Volume Pairs,” on page 78.
File-level Snapshot	Yes	No known issues.
Flush Files Immediately	Yes	No known issues.

NSS Volume Attribute	Supported	Caveats
Lookup Namespace	Yes	The default Lookup Namespace for NSS on Linux is Long, which treats file names as case insensitive. In prior versions, the default name space is UNIX. Using the Long name space helps improve performance because NetWare and Windows treat file names as case insensitive. This is especially important when files are accessed through the SMB/CIFS protocol.
Migration (to near-line HSM storage)	No	DST should not be used in combination with HSM storage solutions.
Modified File List (Use Event File List APIs instead.)	No	<p>By default, modified files are moved back to the primary location. If you disable the Shift Modified Files parameter, modified files might also be located on the secondary location.</p> <p>Modified File List is rarely used. It has been replaced by the Event File List APIs that provide more information than the Modified File List. For information, see “Using the Event File List to Refine the Backup” in the <i>OES 2018 SP3: NSS File System Administration Guide for Linux</i>.</p>
Salvage	Yes	<p>Deleted files on a NSS volume that are salvageable remain salvageable after that volume is used in a shadow volume pair.</p> <p>The Salvage attribute is set separately on each of the NSS volumes in the pair. Because a file can exist on either the primary volume or the secondary volume, a single instance of a deleted file is saved to the volume where it resides when it is deleted. A salvaged file is restored to the volume where it resided when it was deleted.</p> <p>The primary volume and the secondary volume each have an instance of a folder in their file trees. When you delete a folder, both instances are deleted, and each folder’s content is also deleted. Each of the deleted folder instances contain different deleted files. DST does not present a merged view of the deleted folders and files. Duplicate deleted folders are presented when you use the Salvage (undelete) and Purge options for folders. In a Salvage list or Purge list, the folder instance that resided on the primary volume when the folder was deleted has a valid Deleter ID. The folder instance that resided on the secondary volume reports that the Deleter ID is [Supervisor]. You must salvage both deleted folders to access the deleted files in each.</p> <p>NetStorage does not see the deleted files that are available for salvage on the secondary volume.</p> <p>In NFARM utility, when you use salvage or purge options, the salvage or purge list displays only one instance of the deleted folder with the correct Deleter ID. You can salvage the deleted folder to their original location and access the deleted files. The deleted files are restored at their appropriate primary and secondary volumes.</p>
User Space Quotas	Yes	Set up the user space quotas separately on each of the volumes. For more information, see Section 9.7, “Using NSS Quotas on DST Shadow Volume Pairs,” on page 78.
User-level Transaction Model	No	NSS does not support the NetWare Transaction Tracking System for NSS volumes on Linux.

9.4.2 DST Support for NSS Features and Actions

Table 9-2 describes caveats for using the NSS volume features and actions when working with DST shadow volumes.

Table 9-2 Caveats for NSS Features and Actions

NSS Feature	Supported	Caveats
Novell Distributed File Services	Yes	Some limitations apply. For information, see Section 9.10, “Using Novell Distributed File Services with DST Shadow Volume Pairs,” on page 81.
Encryption	Yes	Using encrypted NSS volumes is supported for DST shadow volume pairs. For information, see Section 9.6, “Using NSS Encrypted Volumes in a DST Shadow Volume Pair,” on page 78.
Hard links	No	DST does not support hard links on NSS volumes used in a shadow volume. if a file is a hard link, and the hard-linked file is moved between the primary and the secondary area, the move is really a copy and has the effect of breaking the hard link and creating an additional version of the file that is not linked to the other ones.
Media format for enhanced hard links	Yes	The media format that supports enhanced hard links is supported, but the hard links themselves are not.
Multiple Server Activation Prevention	Yes	Each pool enforces this separately.
Pool low-space warnings and watermarks	Yes	You must set the pool-level watermarks for low-space warnings separately for the primary pool and the secondary pool. IMPORTANT: NSS does not have a volume-level low-space-warning feature. However, you can take advantage of the NCP Server global parameters for managing low-space warnings for NCP volumes on NSS, Ext3, and Reiser file systems. For information, see “NCP Volumes Low-Space Warning” in the <i>OES 2018 SP3: NCP Server for Linux Administration Guide</i> .
Pool snapshots	Yes	Take pool snapshots separately for the primary and secondary pools. IMPORTANT: For NSS on Linux, pool snapshots are not supported for clustered pools. If the primary volume is configured to contain the most frequently used data, pool snapshots of the primary pool have a higher percentage of changed data than does the secondary pool.
Renaming a volume’s mount point	No	Renaming a volume’s mount point breaks the shadow volume. If you need to rename a volume’s mount point, remove the shadow, rename the volume’s mount point, then create the shadow volume again.

NSS Feature	Supported	Caveats
Renaming a volume	No	Renaming a volume breaks the shadow volume. If you need to rename a volume, remove the shadow, rename the volume, then create the shadow volume again.
Resizing (growing) a pool	Yes	No known issues.
Sharing a pool and its volumes in a cluster	Yes	When using NSS volumes in shared pools, you must manage both pools' resources in the primary pool resource load and unload scripts. For information, see Section 9.9, "Using OES Cluster Services with DST Shadow Volume Pairs," on page 80.
Volume quotas	Yes	Set the volume quotas separately for each volume. For more information, see Section 9.7, "Using NSS Quotas on DST Shadow Volume Pairs," on page 78.

9.5 Using Trustees, Trustee Rights, and File System Attributes on DST Shadow Volume Pairs

Authentication and file access is controlled by the file system trustees and rights that you set from the merged view. Users do not have direct access to the secondary volume. You can set trustees, trustee rights, and file system attributes by using the NCP client or the Files and Folders plug-in to Novell iManager and accessing the primary volume. When you access the volume with NCP tools via the primary volume, NCP automatically shows the merged view.

To manage the rights of Active Directory trustees on DST volumes, you can use OES File Access Rights Management (NFARM) utility or `rights` utility. For more information, see [Managing the Trustee Rights in the NSS File System](#) and `rights` in the [OES 2018 SP3: NSS File System Administration Guide for Linux](#).

IMPORTANT: You set trustees, trustee rights, and file system attributes exactly as you would on any NSS volume. DST applies the settings appropriately. You do not need to be aware of where the data physically resides between the two volumes. For information about setting file system attributes, see ["Configuring File System Trustees, Trustee Rights, Inherited Rights Filters, and Attributes"](#) in the [OES 2018 SP3: NSS File System Administration Guide for Linux](#).

Trustees and rights are tied very closely to NCP and NSS, and are coordinated between them via an event system. Therefore, if you use the NCP client or Files and Folders plug-in from the merged view to add, modify, or remove trustees and trustee rights, those changes are automatically synchronized from NCP to NSS, and the event is also automatically synchronized from the primary volume to the secondary volume.

Explicit trustee settings for files and folders are stored in both volumes.

Inherited trustee rights are calculated and enforced based on the trustee settings for the folders on the primary volume. The primary folder tree contains instances of the folders on the secondary volume in order to support this function.

File system attributes are enforced by the NSS file system after they are set.

9.6 Using NSS Encrypted Volumes in a DST Shadow Volume Pair

If encrypted NSS volumes are used, both the primary and secondary volumes should be encrypted in order to provide the same level of security on both volumes.

The first time the volumes are mounted after a server reboot, the encrypted volumes must be mounted manually by using NSSMU or NSS commands in order to provide the encryption passwords. Mount the secondary volume first so that it is available to DST when you mount the primary volume. You can use the standard mount procedure for encrypted NSS volumes as described in [“Mounting an Encrypted NSS Volume with NSSMU”](#) and [“Mounting Encrypted NSS Volumes with NSS Commands”](#) in the *OES 2018 SP3: NSS File System Administration Guide for Linux*.

9.7 Using NSS Quotas on DST Shadow Volume Pairs

DST supports using volume, directory, and user quotas features of NSS volumes. However, DST does not have a unified quota system for the two volumes that manages quotas for the combined primary and secondary volumes in a shadow volume pair.

- [Section 9.7.1, “NSS Volume Quotas,” on page 78](#)
- [Section 9.7.2, “NSS Directory Quotas,” on page 79](#)
- [Section 9.7.3, “NSS User Quotas,” on page 79](#)

9.7.1 NSS Volume Quotas

Volume quotas specify how big a volume can grow within a NSS pool. You can set a volume quota on the primary volume, secondary volume, or both volumes in the shadow volume pair. Each quota is enforced independently of the other. For information about setting the volume quota by using NSS tools, see [“Managing NSS Volume Quotas”](#) in the *OES 2018 SP3: NSS File System Administration Guide for Linux*.

Users of the shadow volume pair can map drives only to the primary volume. They are not aware of the existence of the secondary volume. Users see only the volume quota status for the primary volume. The volume quota information is not presented with a total space usage across both volumes. Users can actually store up to the quota amount set on each of the volumes, where each limit is enforced separately.

When the user has data stored on both the primary and secondary volume, the user sees the amount of space used only on the primary volume, which does not accurately reflect the total of space used on the two volumes.

The administrator can check the combined space available on the shadow volume pair and on each volume separately by using OES Remote Manager for Linux.

9.7.2 NSS Directory Quotas

An NSS directory quota limits the amount of space on a volume that can be consumed by all of the files and folders in that directory. Before you can set directory quotas, you must enable the volume's Directory Quotas attribute. As the administrator, you can view and configure directory quotas with the Files and Folders plug-in for iManager or Client for Open Enterprise Server. You can also use the OES File Access Rights Management (NFARM) utility to set Directory Quotas.

For information about setting a directory quota, see the following sections in the [OES 2018 SP3: NSS File System Administration Guide for Linux](#):

- ♦ **Storage plug-in:** “[Managing Directory Quotas](#)”
- ♦ **Files and Folders plug-in:** “[Viewing, Adding, Modifying, or Removing a Directory Quota](#)”

For a DST volume, you can set a directory quota only on the primary volume. When a secondary volume is in a shadow volume pair, the NSS volume is not bound to NCP and is not available for direct access by the NCP clients or the Files and Folders plug-in in iManager. The directory quota that you add, modify, or remove on the primary volume's instance of the directory is enforced only for the space consumed on the primary volume; it has no effect on the directory instance on the secondary volume.

The users can store up to the directory quota amount for the directory on the primary volume, and an unlimited amount up to the maximum volume size on the secondary volume.

Users see only the directory quota status for the primary volume. The directory quota information is not presented with a total for the directory across both volumes.

To manage directory quotas using NFARM, see [Information or Directory Quota](#) in the [OES 2018 SP3: NSS AD Administration Guide](#).

9.7.3 NSS User Quotas

User quotas are set on specific user names and specify how much data a user can store on a specific volume. Before you can set user quotas, you must enable the volume's User Quotas attribute. Configure the attribute on each volume in the shadow pair to enforce quotas on each volume. As the administrator, you can view and configure user quotas with the Storage plug-in for iManager. For information, see “[Managing User Space Quotas](#)” in the [OES 2018 SP3: NSS File System Administration Guide for Linux](#).

After the User Quotas attribute is enabled for a volume, NSS tracks file space usage for every user who owns one or more files on the volume. The usage is tracked even if the user has no quota assigned.

For a DST volume, both volumes are mounted in NSS. This allows you to set, modify, or remove a user quota for a user separately on the primary volume and the secondary volume. Quotas on each volume are enforced independently of the other. For example, if you remove a user quota on the primary volume, the user quota on the secondary volume is not automatically removed.

Users see only the user quota status for the primary volume. The user quota information is not presented with a total space usage across both volumes. Users can actually store up to the user quota amount set on each of the volumes, where each limit is enforced separately.

To manage user quotas using NFARM, see [User Quota](#) in the [OES 2018 SP3: NSS AD Administration Guide](#).

9.8 Using Opportunistic File Locking on DST Shadow Volume Pairs

Opportunistic locking is a feature of NCP Server that works to ensure that only one person at a time can open a file for write. In a DST shadow volume pair, NCP independently controls opportunistic locking for each volume. Users access files only via the merged view. There is no direct access by users to the secondary volume. When DST access a file on the secondary volume on behalf of multiple users, each request is related to the connection for a given user. NCP treats opportunistic locks for the request as if it came directly from the user.

For information about opportunistic locking and how it works, see [“Using Opportunistic Locking for NCP File Handling”](#) in the *OES 2018 SP3: NCP Server for Linux Administration Guide*.

9.9 Using OES Cluster Services with DST Shadow Volume Pairs

DST supports using DST shadow volumes with OES Cluster Services for Linux for clusters of up to 16 nodes. Clustering is supported for NSS volumes on shared Fibre Channel and iSCSI devices. Users can access files via NCP and via either OES CIFS or SMB/CIFS.

The following caveats apply:

- ◆ The primary cluster resource must be joined to the Active Directory domain in order for the Active Directory users to access the DST shadow volume pair.
- ◆ All nodes where you plan to fail over the shadow volume must be running OES 2015 or later and be configured for DST. The nodes must have the same configuration of file systems, access protocols, and so on.
- ◆ DST and the NCP Server services are not cluster aware. They must be installed and configured separately on each node in the cluster.
- ◆ Global policies for DST must have the same settings on each node in the server. To manage a global DST policy for a given node, open OES Remote Manager for Linux by using the IP address of the node, not the cluster resource. For information about configuring DST global policies, the [Chapter 4, “Installing Dynamic Storage Technology,” on page 35](#).
- ◆ To manage shadow volume policies in a cluster, open OES Remote Manager for Linux by using the IP address of the cluster resource. You can also open OES Remote Manager by using the IP address of the physical node where the cluster resource is currently mounted if you know which node it is on.
- ◆ The individual shadow volume’s policies fail over along with the shadow volume.
- ◆ The primary volume and the secondary volume are managed in the primary cluster resource load and unload scripts. This allows the configuration to be failed over or cluster migrated to a different node as a single resource.

For planning information about installing and configuring shadow volumes in a cluster, see [Chapter 17, “Configuring DST Shadow Volume Pairs with OES Cluster Services,” on page 179](#)

9.10 Using Novell Distributed File Services with DST Shadow Volume Pairs

Novell Distributed File Services (DFS) is installed automatically as part of the NSS file system. Dynamic Storage Technology supports using DFS junctions on the primary NSS volume in a shadow volume pair. The primary volume can also be the target of a junction. Primary NSS volumes that contain DFS junctions or are junction targets can reside in a OES Cluster Services cluster.

DST does not support using DFS junctions on the secondary volume. The secondary volume cannot be the target of a junction.

Junctions that are created on a primary NSS volume behave normally as they would for a single NSS volume. The policy enforcer is designed to prevent junctions from being moved between the primary and secondary volumes when policies are run.

[Table 9-3](#) summarizes the supported DST configurations for use with DFS:

Table 9-3 *DST Support for Novell DFS Features*

Novell DFS Features	Primary NSS Volume	Secondary NSS Volume	Cluster	File Access Protocol
Junctions	Yes	No	Yes	NCP OES CIFS. DFS support must be enabled in CIFS. See “DFS Junction Support in CIFS Linux” in the <i>OES 2018 SP3: OES CIFS for Linux Administration Guide</i> . Linux Samba does not support DFS junctions for NSS volumes.
Junction targets	Yes	No	Yes	Not applicable
Move/split volumes	No	No	No	Not applicable

When you use DST shadow volumes in combination with Novell DFS junctions, consider the following caveats:

- ♦ Junctions are broken when they reside on secondary NSS volumes. If you use an existing NSS volume as a secondary volume, delete junctions on it before you create the shadow volume pair. Make a note of the paths of the junctions and their targets. After the shadow volume pair is working, you can re-create the junctions in the same path on the primary volume.
- ♦ If you use an existing NSS volume as a secondary volume, any junctions pointing to it are broken when you create the shadow volume pair. You must create a new junction that points to the same location on the primary volume of that shadow relationship. After the new junction is working, delete the junction that points to the secondary volume.
- ♦ Do not create a shadow relationship for an NSS volume if a DFS move volume or split volume job is in progress.
- ♦ You must remove the shadow volume before you start a DFS move or split volume job.

9.11 Using OES Storage Services Auditing Client Logger (VLOG) Utility with DST Shadow Volume Pairs

The OES Storage Services Client Logger (VLOG) utility can be used to audit user access to files on the DST shadow volume pair. You should see the following events in VLOG:

- ♦ A directory has an instance on the primary volume and secondary volume. When a user performs an operation on a directory (such as open, close, create, delete, modify, or rename), you should see two entries in VLOG that record the separate action that occurs on each volume.
- ♦ A file resides either on the primary volume or the secondary volume. When a user performs an operation on a file (such as open, close, create, delete, modify, or rename), you should see one entry in VLOG for the volume where the file currently resides.
- ♦ DST managed movement of files between the primary volume and secondary volume is not captured in the audit log. For example, a DST policy moves a file from the primary volume to the secondary volume (or vice versa), or a file is shifted from the secondary volume to the primary volume based on the Last Accessed global policy.

To capture user access events in VLOG for the secondary volume, ensure that you do not enable the `-b (blockNssEventsOfVol)` option for the secondary volume. This option filters the output based on the volume name.

9.12 Using Virus Checking Utilities with DST Shadow Volume Pairs

You can continue use of existing virus checking utilities that currently execute successfully against the designated file systems on the primary volume. DST is transparent to this operation. Because the only access to the secondary volume is through the primary volume, there is no need for a virus checking operation directly on the secondary volume unless the shadow volume relationship is removed and the volume acts independently again.

9.13 Using Backup Utilities with DST Shadow Volume Pairs

Applications that directly access the local Linux file system see the primary file tree and the secondary file tree as independent directories. Thus, 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.

Using shadow volumes allows backups of important data to be made faster and more frequently, because you can apply different backup policies for the primary volume and secondary volume. For example, the server administrator can partition the volume's data into two categories:

- ♦ Important data that needs to be maintained on quality storage and backed up frequently
- ♦ Less important data that can be stored on less expensive storage and backed up less frequently

An analysis or inventory of a volume's data shows that a large portion of it is seldom used. Having a shadow volume allows the server administrator to spend more time on the most important data and spend less on the less important data. The frequently used data can be backed up nightly. The

seldom-used data can be backed up weekly or monthly. Getting the less important data out of the way enables the backups of your important data to run more quickly and efficiently. Partitioning your data in this way can significantly reduce the cost of hosting it.

Because the most important files are located in the primary storage area, disaster recovery can also be faster. The server administrator can restore the critical files by restoring the primary storage area first, and then restore the secondary storage area. This lets users quickly get the files they need most, and infrequently used files are restored as they work. In addition, more fault tolerant replication solutions can be deployed for the primary storage area where it matters most.

If you back up the NSS volume by using the NSS Extended Attributes (`XAttr`) settings to preserve the NetWare metadata (`netware.metadata`) for file system rights and attributes, this information can be restored only directly to an NSS volume. For information about using the NSS / `ListXAttrNWmetadata` option and the security considerations involved, see [“ListXAttrNWmetadata Option”](#) in the *OES 2018 SP3: NSS File System Administration Guide for Linux*.

Backup utilities might also work against the ShadowFS merged view. However, because the ShadowFS merged view is a Linux mount point and is not seen by the backup software as an NSS volume, the file system rights and attributes are not preserved. It is not recommended or supported to back up NSS files from this Linux mount point.

10 Creating and Managing DST Shadow Volumes for NSS Volumes

Dynamic Storage Technology (DST) supports shadow volume pairs with two OES Storage Services (NSS) volumes on Open Enterprise Server (OES). This section describes how to create and manage shadow volume pairs with NSS volumes.

- ◆ [Section 10.1, “Understanding DST Volume Pairs,” on page 85](#)
- ◆ [Section 10.2, “Creating a DST Shadow Volume with NSS Volumes,” on page 88](#)
- ◆ [Section 10.3, “Giving Users a Merged View of the Shadow Volume,” on page 96](#)
- ◆ [Section 10.4, “Replicating the Secondary File Tree Structure to the Primary Volume,” on page 97](#)
- ◆ [Section 10.5, “Configuring the NCP/NSS Bindings for an NSS Volume,” on page 97](#)
- ◆ [Section 10.6, “Viewing a List of NCP Shares,” on page 100](#)
- ◆ [Section 10.7, “Mounting and Dismounting DST Shadow Volumes,” on page 101](#)
- ◆ [Section 10.8, “Viewing the Name and Path Information for a Shadow Volume,” on page 101](#)
- ◆ [Section 10.9, “Viewing Information about a Shadow Volume,” on page 101](#)
- ◆ [Section 10.10, “Auditing File Move Events for the Shadow Volume,” on page 106](#)
- ◆ [Section 10.11, “Backing Up DST Shadow Volumes,” on page 106](#)
- ◆ [Section 10.12, “Removing the Shadow Relationship for a Non-Clustered DST Shadow Volume,” on page 110](#)

10.1 Understanding DST Volume Pairs

The DST shadow volume is a virtual NCP (NetWare Core Protocol) volume that consists of a primary storage area and a secondary storage area. The primary and secondary areas use NSS volumes.

- ◆ [Section 10.1.1, “Primary Volume,” on page 86](#)
- ◆ [Section 10.1.2, “Secondary Volume,” on page 86](#)
- ◆ [Section 10.1.3, “Merged View,” on page 86](#)
- ◆ [Section 10.1.4, “How Directories Are Created in the Shadow Volume,” on page 86](#)
- ◆ [Section 10.1.5, “Global DST Policies,” on page 86](#)
- ◆ [Section 10.1.6, “Shadow Volume Policies,” on page 87](#)
- ◆ [Section 10.1.7, “File Inventory for the Shadow Volume,” on page 87](#)
- ◆ [Section 10.1.8, “Moving Specified Files between Volumes,” on page 87](#)

10.1.1 Primary Volume

The primary volume is an NSS volume that is mounted on an OES server that is running DST. Shadow volumes are known by their primary volume name. Typically, the primary volume is on the higher-performance device.

When the primary volume has a state of **Shadowed**, the volume ID that is assigned as its NCP volume ID represents the DST shadow volume pair of volumes. The secondary volume does not have a separate volume ID while it is in the shadow relationship.

10.1.2 Secondary Volume

The secondary volume is an NSS volume that is mounted on the same OES server as the primary volume. This volume is also referred to as the *shadow path*. The secondary volume is also referred to as the *secondary file tree*.

The secondary volume is typically a new volume. It should have a similar setup as the primary volume for key attributes settings, such as Salvage, Encryption, and Lookup Namespace. For guidelines and caveats about using NSS volume attributes with Dynamic Storage Technology, see [Table 9-1, “DST Support for NSS Volume Attributes,” on page 74](#).

10.1.3 Merged View

The primary file tree and the secondary file tree have the same directory structure. A file can be located in either the primary file tree or the secondary file tree. The merged view presents these two file trees as a single file tree, as shown in [Figure 1-1, “User View of the File System Directory,” on page 15](#).

All trustee management should be performed from the merged view.

The NCP clients and management tools see a merged view of files on the DST shadow volume when they access the primary volume. OES CIFS also provides a merged view for CIFS users that access CIFS shares on the primary volume.

If Novell Samba (or a supported native Linux protocol) is used with DST shadow volumes, a SMB/CIFS user sees the merged view that is provided by the ShadowFS component of DST.

10.1.4 How Directories Are Created in the Shadow Volume

New directories are created in the primary file tree and secondary file tree to achieve the same directory structure. For more information see [Section 7.1, “Replicating Branches of the Primary File Tree in the Secondary File Tree,” on page 49](#).

10.1.5 Global DST Policies

Global DST policies are NCP Server parameters that govern the behavior of DST. They apply to all shadow volume pairs on a given server. Before you configure shadow volumes on the server, ensure that you configure the global policies listed in [Table 10-1](#):

Table 10-1 Global DST Policies

NCP Server Parameter for DST	For Information
REPLICATE_PRIMARY_TREE_TO_SHADOW	Section 7.1, “Replicating Branches of the Primary File Tree in the Secondary File Tree,” on page 49
SHIFT_MODIFIED_SHADOW_FILES	Section 7.2, “Shifting Files from the Secondary File Tree to the Primary File Tree,” on page 50
SHIFT_ACCESSED_SHADOW_FILES	
SHIFT_DAYS_SINCE_LAST_ACCESS	
DUPLICATE_SHADOW_FILE_ACTION	Section 7.3, “Resolving Instances of Duplicate Files,” on page 55
DUPLICATE_SHADOW_FILE_BROADCAST	

10.1.6 Shadow Volume Policies

Shadow volume policies manage how files are distributed across the shadow volume’s primary and shadow areas. A Shadow Volume policy allows you to specify when the policy is enforced (one time, hourly, daily, weekly, and so on), which volumes the policy applies to, which direction files are moved (primary to shadow or shadow to primary), and which files are moved (file type, modify date, access date, size, and so on). Multiple policies can be applied to the same volumes and multiple policies can be scheduled to run concurrently.

For information about creating or modifying Dynamic Storage Technology policies for shadow volumes, see [Chapter 11, “Creating and Managing Policies for Shadow Volumes,” on page 113](#).

10.1.7 File Inventory for the Shadow Volume

You can generate an inventory of the files located on the two volumes by selecting the **Inventory** link next to the primary volume on the Dynamic Storage Technology Options page. This provides statistics broken out for both volumes and for each volume separately. For information, see [Section 14.4, “Viewing Statistics for the Shadow Volume,” on page 166](#)

10.1.8 Moving Specified Files between Volumes

The Inventory page allows you to navigate through the statistics reports to determine a list of files to be moved between the two volumes (primary to secondary, or secondary to primary). For information, see [Section 14.5, “Using Inventory Detail Reports to Move, Copy, or Delete Files on the Shadow Volume,” on page 166](#).

10.2 Creating a DST Shadow Volume with NSS Volumes

A DST shadow volume links two NSS volumes. Typically, one of the volumes contains data and one is newly created. For information about how to create NSS volumes, see the *OES 2018 SP3: NSS File System Administration Guide for Linux*.

This section describes how to create unshared DST shadow volumes. For information about using shared NSS volumes to create a shared DST shadow volume in a cluster environment, see [Chapter 17, “Configuring DST Shadow Volume Pairs with OES Cluster Services,”](#) on page 179.

IMPORTANT: The following procedures use `VOL1` for the primary storage area, and `ARCVOL` as the secondary storage area. Ensure that you substitute the actual names of the NSS volumes you are using in each of the steps.

- ◆ [Section 10.2.1, “Checking the Availability of an NSS Volume for Primary,”](#) on page 88
- ◆ [Section 10.2.2, “Checking the Availability of an NSS Volume for Secondary,”](#) on page 89
- ◆ [Section 10.2.3, “Using an Existing Volume as Secondary,”](#) on page 90
- ◆ [Section 10.2.4, “Disabling the NCP/NSS Bindings for the Secondary Volume,”](#) on page 93
- ◆ [Section 10.2.5, “Adding a Shadow to the Primary NSS Volume \(Linking the NSS Volumes\),”](#) on page 94
- ◆ [Section 10.2.6, “Moving Data between the Two Volumes,”](#) on page 95

10.2.1 Checking the Availability of an NSS Volume for Primary

Ensure that the volume you want to use as the primary is available and mounted.

- 1 Open OES Remote Manager for Linux in a web browser, then log in to the DST server as the `root` user.
- 2 Select **View File System > Dynamic Storage Technology Options** to view a list of mounted volumes.
- 3 On the Dynamic Storage Technology page, ensure that the NSS volume that you want to use as the primary volume appears in the **Volume Information** list with a status of **Add Shadow**.




If it is listed and mounted, continue with [“Checking the Availability of an NSS Volume for Secondary”](#) on page 89.

Otherwise, the NSS volume might be unmounted, or its NCP/NSS bindings might be disabled.

- 4 Select **Manage NCP Services > Manage Shares** to view a list of active volumes, then do one of the following:

Volume	Status	Action
Listed	Not mounted	Click the Mount button next to the volume name, then continue with “Checking the Availability of an NSS Volume for Secondary” on page 89.



Volume	Status	Action
Not listed	Unknown	<p>Click NCP/NSS Bindings to view the Available NSS Volumes list, then check to see if the volume is listed.</p> <p>Listed: Continue with Step 5.</p> <p>Not Listed: Continue with Step 6.</p>

- 5 If the volume is in the **Available NSS Volumes** list and its **NCP Accessible** value is set to **No**, the volume's NCP/NSS binding is disabled.
 - ♦ **Volume is in use as a secondary in a DST volume pair:** The most likely reason for the setting to be disabled is that the volume is already being used as the secondary volume in another shadow volume. If the volume is used in a DST pool cluster resource, the volume's setting might appear in the list even when its resource is offline on the selected node. Because the volume is already used as a secondary volume, it cannot be used a primary volume. Return to [Step 1](#) and begin again.
 - ♦ **Volume was recently unlinked from another DST volume pair:** If you are certain that the volume is not being used as the secondary volume in another DST shadow volume pair, you can enable the NCP/NSS Bindings setting:
 1. Select **Yes** in the **NCP Accessible** column for the NSS volume, then click **Save Selection** to save and apply the change.
 2. If the volume is not automatically mounted, select **Manage NCP Services > Manage Shares** to view the **Volume Information** list, then click the **Mount** button next to the volume name to mount it.
 3. Continue with [“Checking the Availability of an NSS Volume for Secondary” on page 89](#).
- 6 If the volume does not appear in the **Volume Information** list, and it does not appear on the NCP/NSS Bindings page in the **Available NSS Volumes** list, the volume's pool is deactive or its pool cluster resource is offline.

Exit OES Remote Manager, then use NSS tools to activate the pool and mount the volume, or bring the pool cluster resource online. When the volume is mounted, return to [Step 1](#) and begin again.
- 7 Continue with [Section 10.2.2, “Checking the Availability of an NSS Volume for Secondary,” on page 89](#).

10.2.2 Checking the Availability of an NSS Volume for Secondary

Ensure that the secondary volume is available and mounted.

- 1 Open OES Remote Manager for Linux in a web browser, then log in to the DST server as the `root` user.
- 2 Select **View File System > Dynamic Storage Technology Options** to view a list of mounted volumes.
- 3 On the Dynamic Storage Technology page, ensure that the NSS volume that you want to use as the secondary volume appears in the **Volume Information** list with a status of **Add Shadow**.
The secondary volume must be mounted in NCP in order to perform the next step.

For example, the volume `SHARE` is the NSS volume that is planned to be used for the secondary volume. The volume is in the **Volume Information** list with a **Shadow Status** value of **Add Shadow**.

Volume Information		
Volume Name	Shadow Status	
① SHARE	Add Shadow	Inventory
① VOL2	Add Shadow	Inventory
① DATA	Add Shadow	Inventory
① VOL1	Add Shadow	Inventory
① SYS	Add Shadow	Inventory

- 4 Continue with one of the following, depending on whether the secondary volume contains data:

Primary NSS Volume	Secondary NSS Volume	Continue with
New, empty volume	New, empty volume	Section 10.2.4, “Disabling the NCP/NSS Bindings for the Secondary Volume,” on page 93
Existing volume	New, empty volume	Section 10.2.4, “Disabling the NCP/NSS Bindings for the Secondary Volume,” on page 93
New, empty volume	Existing volume	Section 10.2.3, “Using an Existing Volume as Secondary,” on page 90

10.2.3 Using an Existing Volume as Secondary

In a typical DST shadow volume, either both volumes are new, or the primary volume contains data and the secondary volume is empty. However, you can also use a new volume as primary and an existing volume as secondary.

When you create the DST volume, two default behaviors govern how the trustees and directory structure are replicated on the secondary volume.

- ◆ DST automatically synchronizes the trustee database from the primary volume to the secondary volume.
- ◆ The global policy for the `REPLICATE_PRIMARY_TREE_TO_SHADOW` parameter determines when the file tree structure is replicated from the primary volume to the secondary volume.

The default behavior is designed to handle the typical pair configurations. When the secondary volume contains the data, the existing trustees are overwritten when the trustees database is synchronized from the new primary volume. Instances of the directories are created on the primary volume as files are moved from the secondary volume to the primary volume. Empty directories are not replicated to the primary, so the file structure on the primary is not an exact match of the secondary.

This section describes some techniques you can use to get the existing volume's trustees settings and file tree structure to the primary volume:

- ♦ [“Manually Copying the Trustees Database to the New Volume” on page 91](#)
- ♦ [“Manually Copying the Trustees Database to the New Volume and Replicating the Secondary File Tree Structure to the Primary” on page 92](#)
- ♦ [“Getting DST to Copy the Trustees Database and Replicate the File Tree Structure to the New Volume” on page 92](#)

Manually Copying the Trustees Database to the New Volume

If the primary volume is a new empty volume, the trustee synchronization that occurs when you first link the volumes effectively removes the trustee settings on the existing volume. As a result, you must reconfigure trustees, trustee rights, and inherited rights from the merged view. If the secondary volume is the one with data, you should copy its existing trustee database to the primary volume before you create the shadow volume relationship. Then when the volumes are linked, the existing settings are used for the DST volume pair.

For all NCP volumes (NSS and NCP on Linux POSIX volumes), the trustee information is obtained at volume mount time from the `._NETWARE/.trustee_database.xml` file. For an NSS volume, the Linux path to the file is `/media/nss/volume_name/._NETWARE/.trustee_database.xml`.

Before you create a DST shadow volume pair with a new volume as primary and an existing volume as secondary:

- 1 In OES Remote Manager, log in as the `root` user.
- 2 Do not allow users to access the volumes while you perform the following steps.
- 3 Select **Manage NCP Services > Manage Shares** to view a list of active volumes.
- 4 Dismount the two NSS volumes that you will be using for the DST shadow volume pair from NCP Server by selecting the **Unmount** button next to each volume.

This dismounts the volumes from NCP, but they are still mounted by NSS.

- 5 In a file browser, rename or delete the `/media/nss/primary_volume_name/._NETWARE/.trustee_database.xml` file on the primary volume.
- 6 Open a terminal console as the `root` user, then copy the trustee database file from the secondary volume location to the primary volume location by entering the following at a terminal console prompt:

```
cp /media/nss/secondary_volume_name/._NETWARE/.trustee_database.xml /  
media/nss/primary_volumename/._NETWARE/.trustee_database.xml
```

For example:

```
cp /media/nss/ARCVOL/._NETWARE/.trustee_database.xml /media/nss/VOL1/  
._NETWARE/.trustee_database.xml
```

- 7 Select **Manage NCP Services > Manage Shares** to view a list of active volumes.
- 8 Mount the primary volume and secondary volume for NCP Server by selecting the **Mount** button next to each volume.
- 9 At the terminal console prompt, enter the following command to synchronize the NSS trustee information that is now on the primary volume with NCP Server:

```
ncpcon nss resync=primary_volumename
```

- 10 Continue with [Section 10.2.4, “Disabling the NCP/NSS Bindings for the Secondary Volume,”](#) on page 93.

Manually Copying the Trustees Database to the New Volume and Replicating the Secondary File Tree Structure to the Primary

You can use the following procedure to copy the trustees database to the new volume before you create a DST shadow volume pair with the existing volume as secondary, and then replicate the secondary file tree structure to the primary volume after you create the pair.

- 1 Before you create the DST shadow volume pair, copy the trustees database from the existing volume to the new volume as described in [“Manually Copying the Trustees Database to the New Volume”](#) on page 91.
- 2 After you create the DST shadow volume pair, replicate the secondary file tree structure to the primary volume as described in [Section 10.4, “Replicating the Secondary File Tree Structure to the Primary Volume,”](#) on page 97.

Getting DST to Copy the Trustees Database and Replicate the File Tree Structure to the New Volume

You can use the procedure in this section to get the both the trustees database and the file tree structure from an existing volume to a new volume before configuring the existing volume as secondary.

These instructions use `OLDVOL1` and `NEWVOL1` for the volume names. Modify the commands as needed to use the actual volume names on your system.

- 1 Log in to the server as the `root` user, then open a terminal console.
- 2 In OES Remote Manager, enable the `REPLICATE_PRIMARY_TREE_TO_SHADOW` global policy that allows the file tree structure to be replicated automatically from the primary to the secondary when the DST pair is created.

For information, see [Section 7.1, “Replicating Branches of the Primary File Tree in the Secondary File Tree,”](#) on page 49.

- 3 Temporarily set up a DST shadow volume pair with the old volume as primary and the new volume as secondary. Do not allow users to access the volumes at this time.

For information, see [Section 10.2.4, “Disabling the NCP/NSS Bindings for the Secondary Volume,”](#) on page 93 and [Section 10.2.5, “Adding a Shadow to the Primary NSS Volume \(Linking the NSS Volumes\),”](#) on page 94.

When you create the shadow volume pair, the `.trustee_database.xml` file is automatically copied from the primary volume (`/media/nss/OLDVOL1/._NETWARE/.trustee_database.xml`) to the secondary volume (`/media/nss/NEWVOL1/._NETWARE/.trustee_database.xml`), and the settings are synchronized with NCP Server.

The entire file tree structure of the old volume is replicated to the new volume. This can take a few minutes to an hour, depending on the number of directories and depth of the structure.

- 4 As the `root` user, open the `/media/nss/NEWVOL1/._NETWARE/.trustee_database.xml` file in a text editor, visually verify that it now contains trustees information from the old volume, then close the file.
- 5 Remove the DST shadow relationship between the two volumes.
For information, see [Section 10.12, “Removing the Shadow Relationship for a Non-Clustered DST Shadow Volume,”](#) on page 110.
- 6 Continue with [Section 10.2.4, “Disabling the NCP/NSS Bindings for the Secondary Volume,”](#) on page 93.

10.2.4 Disabling the NCP/NSS Bindings for the Secondary Volume

- 1 Open OES Remote Manager for Linux in a web browser, then log in to the DST server as the `root` user.
- 2 Select **View File System > Dynamic Storage Technology Options** to view a list of mounted volumes.
- 3 Select **Manage NCP Services > Manage Shares**, then click **NCP/NSS Bindings**.
- 4 In the **Available NSS Volumes** list, select **No** in the **NCP Accessible** column for the NSS volume that you want to use as the secondary volume.

Available NSS volumes		
NCP Accessible	Volume Name	Mount point
Yes: <input type="radio"/> No: <input checked="" type="radio"/> <input type="button" value="Save Selection"/>	ARCVOL	/media/nss/ARCVOL
Yes: <input checked="" type="radio"/> No: <input type="radio"/> <input type="button" value="Save Selection"/>	VOL1	/media/nss/VOL1

- 5 Click **Save Selection** to save and apply the change.
- 6 Go to the Dynamic Storage Technology page, and verify that the secondary volume (for example, ARCVOL) is no longer listed.
- 7 Continue with [Section 10.2.5, “Adding a Shadow to the Primary NSS Volume \(Linking the NSS Volumes\),”](#) on page 94.


10.2.5 Adding a Shadow to the Primary NSS Volume (Linking the NSS Volumes)

- 1 Open OES Remote Manager for Linux in a web browser, then log in to the DST server as the `root` user.
- 2 Use one of the following methods to go to the volume's Share Information page of the NSS volume that you want to use as the primary storage area.
 - ♦ Select **View File System > Dynamic Storage Technology Options** to go to the Dynamic Storage Options page, then click the **Add Shadow** link next to the volume name of the NSS volume.

For example, click the **Add Shadow** link for VOL1.

Volume Information		
Volume Name	Shadow Status	
ⓘ DATA	Add Shadow	Inventory
ⓘ VOL1	Add Shadow	Inventory
ⓘ SYS	Add Shadow	Inventory

- ♦ Select **Manage NCP Services > Manage Shares** to open the Manage Shares page, then click the **Information** (ⓘ) icon next to the volume name of the NSS volume.
- 3 On the volume's Share Information page, scroll down to the **Volume Tasks** area, then click **Add Shadow Volume**.



- 4 Specify the following information for the secondary storage area for the DST shadow volume, then click **Create** to define the shadow volume.

Create Shadow for Volume VOL1 ?

The dialog box has a title bar "Create Shadow for Volume VOL1" with a question mark icon on the right. It contains a "Shadow Path:" label followed by a text input field. Below the input field is a checkbox labeled "Create if not present". At the bottom of the dialog are two buttons: "Create" and "Cancel".

- ♦ **Shadow Path:** Type the Linux path for the NSS volume that you want to use as the secondary storage area. The default Linux path where NSS volumes are mounted is `/media/nss/volumename`.
For example, to specify the NSS volume named ARCVOL as the secondary storage area, type `/media/nss/ARCVOL` in the **Shadow Path** field.
- ♦ **Create if not present:** For NSS volumes, the volume must already exist. Ensure that this option is deselected (not checked) when shadowing NSS volumes.

IMPORTANT: This option is a placeholder for future capabilities to support shadow volumes for NCP volumes on Linux POSIX file systems (such as Ext3, Reiser, and XFS).

- 5 On the volume's Share Information page, ensure that the **File System Shadow Path** information shows the shadow path you specified in [Step 4](#).

VOL1 Share Information ?

Information	
Description	Value
File system path	/media/nss/VOL1
File system shadow path	/media/nss/ARCVOL
Loaded name spaces	DOS LONG
File system type	NSS
NCP volume ID	254
Status	mounted online cluster resource salvageable user quotas
Sector Size	512
Sectors per Cluster	8
Capacity	4.92 GB
Used space	632 KB
Advanced Information	<input type="button" value="View"/>

- 6 Select **View File System > Dynamic Storage Technology Options** to go to the Dynamic Storage Options page, then verify that the **Shadow Status** for the volume is set to **Shadowed** and the **View Log** link is available.
- 7 Continue with [Section 10.2.6, "Moving Data between the Two Volumes,"](#) on page 95.

10.2.6 Moving Data between the Two Volumes

- 1 In OES Remote Manager, select **View File System > Dynamic Storage Technology Options** to go to the Dynamic Storage Options page, then create one or multiple shadow volume policies for the shadow volume.

Shadow volume policies can be configured to move files according to the time since the file was last modified, accessed, or changed; by file names; by file types; or by file size. You can schedule policies to run automatically, or you can run them on demand.

For information about creating and scheduling shadow volume policies, see [Chapter 11, "Creating and Managing Policies for Shadow Volumes,"](#) on page 113.

- 2 (Optional) Move selected data on demand by running customized inventory reports, then using the inventory detail reports to move selected files to either volume according to the time since the file was last modified, accessed, or changed; by file names; by file types; or by file size.

For information, see [Section 14.5, “Using Inventory Detail Reports to Move, Copy, or Delete Files on the Shadow Volume,”](#) on page 166.

10.3 Giving Users a Merged View of the Shadow Volume

Users see a merged view of the data by accessing a share on the primary volume. The following user access is supported:

- ♦ [Section 10.3.1, “NCP,”](#) on page 96
- ♦ [Section 10.3.2, “OES CIFS,”](#) on page 96
- ♦ [Section 10.3.3, “Novell Samba with ShadowFS and FUSE,”](#) on page 96

10.3.1 NCP

Configure file access for the NCP users on the primary NSS volume, just as you would for a single NSS volume. NCP automatically provides a merged view of the data.

10.3.2 OES CIFS

Configure file access for the CIFS users on the primary NSS volume, just as you would for a single NSS volume. CIFS automatically provides a merged view of the data.

Beginning with OES 2015 or later, Active Directory users can also access files residing in shadow volume using the merged view provided by CIFS.

10.3.3 Novell Samba with ShadowFS and FUSE

Configure file access for the Novell Samba users by configuring ShadowFS and FUSE, then create a share on the primary NSS volume. Linux Samba and the SMB/CIFS users must be enabled for Linux User Management.

For information, see [Chapter 13, “Using ShadowFS to Provide a Merged View for Novell Samba Users,”](#) on page 145.

10.4 Replicating the Secondary File Tree Structure to the Primary Volume

To ensure that the entire existing file tree structure (including empty directories) on the secondary volume is replicated to the primary volume, an administrator can browse (enumerate) the directory from an NCP client that is mapped to the merged view. You might need to perform this task in the following situations:

- ♦ You used a new volume as primary and an existing volume as secondary, and did not replicate the directory structure before you created the DST volume pair.
- ♦ You restored data to the secondary volume, and you want to ensure that empty directories are replicate to the primary.

To replicate the secondary file tree to the primary:

- 1 On an NCP client, map a drive to the root of the DST shadow volume pair.
NCP provides a merged view access when you access the primary volume of the pair.
- 2 On the NCP client in the Windows Explorer browser, browse to the topmost subdirectory of the mapped drive.
Instances of the top level of subdirectories are created on the primary volume.
- 3 Launch a Windows terminal console, then change directory to the drive letter that you mapped to the DST volume pair.
- 4 At the command prompt, enter the following command to enumerate the directory:

```
dir *.* /s
```

Instances of all subdirectories, including empty directories, are created on the primary volume.

10.5 Configuring the NCP/NSS Bindings for an NSS Volume

The NCP/NSS Bindings parameter for an NSS volume governs whether the volume is automatically mounted on system restart in NCP Server. When the parameter is enabled, the NSS volume is automatically mounted for NCP Server. When it is disabled, the NSS volume is not mounted for NCP Server. The NCP/NSS Bindings parameter is enabled by default, making the volume NCP accessible.

NSS volumes are automatically mounted by default on system restart, first in NSS, then in NCP Server. This is the desired behavior for all independent NSS volumes that are not in shadow volumes, and for NSS volumes that you use as primary storage locations in a DST shadow volumes. When an NSS volume is used as the secondary storage area in a DST shadow volume, you want the NSS volume to be mounted in NSS, but not in NCP Server. This allows DST to control access to the secondary storage area via the primary storage area. Files in the secondary storage area cannot be directly accessed by users.


After you remove a shadow volume, the NCP/NSS Bindings parameter for the NSS volume that was used as the secondary storage area remains disabled until you enable it. You must enable the bindings and mount the volume in order to enable users to access the now independent volume.

- ♦ [Section 10.5.1, “Disabling the NCP/NSS Bindings for an NSS Volume,” on page 98](#)
- ♦ [Section 10.5.2, “Enabling the NCP/NSS Bindings for an NSS Volume,” on page 99](#)
- ♦ [Section 10.5.3, “Enabling or Disabling NCP/NSS Bindings by Editing the /etc/opt/novell/ncp2nss.conf File,” on page 100](#)

10.5.1 Disabling the NCP/NSS Bindings for an NSS Volume

The volume’s NCP/NSS Bindings parameter must be disabled for NSS volumes that you use as secondary storage locations in a DST shadow volumes.

- 1 In OES Remote Manager for Linux, select **Manage NCP Services > Manage Shares**.
- 2 In the **Configuration** area of the NCP Shares page, click **NCP/NSS Bindings**.



- 3 In the **Available NSS Volumes** list, locate the NSS volume that you want to disable.
- 4 In the **NCP Accessible** column, click **No** to make the NSS volume not accessible to NCP so that it is not mounted in NCP after it is mounted in NSS.

Available NSS volumes		
NCP Accessible	Volume Name	Mount point
Yes: <input type="radio"/> No: <input checked="" type="radio"/> <input type="button" value="Save Selection"/>	ARCVOL	/media/nss/ARCVOL
Yes: <input checked="" type="radio"/> No: <input type="radio"/> <input type="button" value="Save Selection"/>	VOL1	/media/nss/VOL1


- 5 Beneath the volume’s setting for **NCP Accessible**, click **Save Selection** to save and apply the new setting.
- 6 Verify that the NSS volume is not available for NCP by selecting **Manage NCP Services > Manage Shares** to view a list of active volumes.

If the NCP/NSS Bindings parameter is successfully disabled, the NSS volume should not appear in the **Volume Information** list.

10.5.2 Enabling the NCP/NSS Bindings for an NSS Volume

The volume's NCP/NSS Bindings parameter must be enabled for NSS volumes that you use as primary storage locations in a DST shadow volumes, and for all independent NSS volumes that are not in shadow volumes. This is the default.

- 1 In OES Remote Manager for Linux, select **Manage NCP Services > Manage Shares**.



- 2 In the **Configuration** area of the NCP Shares page, click **NCP/NSS Bindings** to open the NCP/NSS Bindings page.
- 3 In the **Available NSS Volumes** list, locate the NSS volume that you want to enable.

Available NSS volumes		
NCP Accessible	Volume Name	Mount point
Yes: <input type="radio"/> No: <input checked="" type="radio"/> Save Selection	ARCVOL	/media/nss/ARCVOL
Yes: <input checked="" type="radio"/> No: <input type="radio"/> Save Selection	VOL1	/media/nss/VOL1

- 4 If the volume's **NCP Accessible** setting is **No**, click **Yes** to make the NSS volume accessible to NCP so that the volume is automatically mounted in NCP after it is mounted in NSS.

Available NSS volumes		
NCP Accessible	Volume Name	Mount point
Yes: <input checked="" type="radio"/> No: <input type="radio"/> Save Selection	VOL1	/media/nss/VOL1
Yes: <input type="radio"/> No: <input checked="" type="radio"/> Save Selection	ARCVOL	/media/nss/ARCVOL

- 5 Beneath the volume's setting for **NCP Accessible**, click **Save Selection** to save and apply the new setting.

The volume appears in the **Active Shares** list on the NCP Shares page.

- 6 Verify that the NSS volume is available for NCP by selecting **Manage NCP Services > Manage Shares** to view a list of active volumes.

If the NSS/NCP Bindings parameter is enabled, the NSS volume appears in the **Volume Information** list, and a **Mount** button is displayed next to it.

- 7 Click **Mount**.

When the volume is successfully mounted, the volume's name is hyperlinked, and an **Unmount** button is displayed next to it.

10.5.3 Enabling or Disabling NCP/NSS Bindings by Editing the `/etc/opt/novell/ncp2nss.conf` File

When the NCP/NSS bindings parameter is disabled for a volume, NCP Server adds an `EXCLUDE_VOLUME` entry to the `/etc/opt/novell/ncp2nss.conf` file. You can manually disable or enable the NSS volume's NCP/NSS bindings parameter by adding or removing this entry from the file, then restarting the NCP2NSS daemon.

- 1 Open the `/etc/opt/novell/ncp2nss.conf` configuration file in a text editor.
- 2 Do one of the following:
 - ♦ **Disable the NCP/NSS Binding:** Add an `EXCLUDE_VOLUME` entry for the volume you plan to use as secondary NSS volume in order to exclude the volume from being automatically mounted for NCP Server.

```
EXCLUDE_VOLUME nss_volumename
```

Replace `nss_volumename` with the name of the NSS volume. For example, to disable the bindings for the NSS volume named `ARCVOL:`, add the following line. Note that you do not include the colon after the volume name.

```
EXCLUDE_VOLUME ARCVOL
```

- ♦ **Enable the NCP/NSS Binding:** Locate the `EXCLUDE_VOLUME` entry for the NSS volume, then remove that line from the file.
- 3 Save the file.
 - 4 Restart the eDirectory daemon by entering the following commands:

```
rcnssd stop
```

```
rcnssd start
```

- 5 Restart the NCP/NSS IPC daemon to synchronize the changes you made to the `/etc/opt/novell/ncp2nss.conf` file:
 - 5a At the terminal console prompt, enter

```
systemctl restart ncp2nss.service
```

10.6 Viewing a List of NCP Shares

In OES Remote Manager, the NCP Server plug-in appears as the **Manage NCP Services** role in the left panel. This allows you to mount or unmount NCP volumes, NSS volumes, and DST shadow volumes from the NCP Server. Unmounting an NSS volume from NCP does not dismount the volume from NSS.

- 1 Open OES Remote Manager, then log in to the DST server as the `root` user.
- 2 Use either of the following methods to view a list of NCP shares and their status:
 - ♦ Select **Manage NCP Services > Manage Shares**.
 - ♦ Select **View File System > Dynamic Storage Technology Options** to open the Volume Information report.

10.7 Mounting and Dismounting DST Shadow Volumes

To mount or dismount the DST shadow volume for NCP Server, you mount or dismount the primary storage area. Unmounting an NSS volume from NCP does not dismount the volume from NSS.

To mount a shadow volume:

- 1 In OES Remote Manager, click **Manage NCP Services > Manage Shares**, then click the **Mount** button next to the volume name of the primary storage area for the DST shadow volume you want to mount.

To dismount a shadow volume:

- 1 In OES Remote Manager, click **Manage NCP Services > Manage Shares**, then click the **Unmount** button next to the volume name of the primary storage area for the DST shadow volume you want to dismount.

10.8 Viewing the Name and Path Information for a Shadow Volume

You can quickly get name and path information for the member volumes in the DST shadow volume by using the `ncpcon volume` command.

- 1 Log in to the DST server as the `root` user, then open a terminal console.
- 2 At the terminal console prompt, enter

```
ncpcon volume <primary volume name>
```

This information is also available in OES Remote Manager as described in [Section 10.9.3, “Viewing the Share Information for a Shadow Volume,”](#) on page 104.

10.9 Viewing Information about a Shadow Volume

In OES Remote Manager on the Dynamic Storage Technology Options page, the **Volume Information** report (as shown in [Figure 10-1](#)) contains information about all NCP volumes on the server. This includes NSS volumes (which are by default NCP volumes) and NCP shares on Linux POSIX file systems such as Ext3, Reiser, and XFS.

IMPORTANT: DST supports NSS volumes to be used in shadow volumes at this time.

Figure 10-1 Volume Information Report

Dynamic Storage Technology Options



Dynamic Storage Technology allows you optimize the use of your storage by automatically moving data to storage best optimized for the data type or frequency of use. You can create one or more policies to manage, and optimize the use of your storage.

Volume Information			
Volume Name	Shadow Status		
① SHARE	Add Shadow	Inventory	
① VOL2	Add Shadow	Inventory	
① DATA	Add Shadow	Inventory	
① VOL1	Shadowed	Inventory	View Log
① SYS	Add Shadow	Inventory	

The report does not distinguish between the underlying file systems for the NCP volumes. Ensure that you create shadows only for NCP volumes based on the NSS file system. You can identify whether a volume is an NSS volume by clicking the **Information** icon next to the volume name, then viewing its underlying file system type.

To understand the information provided in the report, see the following sections:

- ◆ [Section 10.9.1, “Accessing the Volume Information Report,” on page 102](#)
- ◆ [Section 10.9.2, “Viewing the Shadow Status of a Volume,” on page 102](#)
- ◆ [Section 10.9.3, “Viewing the Share Information for a Shadow Volume,” on page 104](#)

10.9.1 Accessing the Volume Information Report

- 1 Open OES Remote Manager, then log in as the `root` user.
- 2 Select **View File System > Dynamic Storage Technology Options** to open the Volume Information report.

10.9.2 Viewing the Shadow Status of a Volume

In the **Volume Information** report, the **Shadow Status** column displays whether or not a volume has a shadow. There are three states:

Table 10-2 Shadow Status in the Volume Information Report

Shadow State	Description
No Shadow	The NSS <code>_ADMIN</code> volume cannot be shadowed, and displays a status of No Shadow.

Shadow State	Description
Add Shadow	<p>The volume is an NSS volume or an NCP volume that is eligible for shadowing. You must separately verify that the volume satisfies the guidelines and caveats that are specified in Chapter 9, “Planning for DST Shadow Volume Pairs and Policies,” on page 65.</p> <p>IMPORTANT: Select the Add Shadow link only for NCP volumes where the underlying file system is the NSS file system.</p>
Shadowed	<p>The volume is the primary volume in a DST shadow volume. To identify the secondary storage area for this volume, click the Information icon next to the volume name to go to the Share Information page, then view the File System Shadow Path.</p>

10.9.3 Viewing the Share Information for a Shadow Volume

The Share Information page displays details about the NCP volume, such as its Linux file system path, the file system path of its shadow area (if it is shadowed), the file system type, and capacity.

Figure 10-2 NCP Volume Share Information

VOL1 Share Information

Information	
Description	Value
File system path	/media/nss/VOL1
File system shadow path	/media/nss/ARCVOL
Loaded name spaces	DOS LONG
File system type	NSS
NCP volume ID	254
Status	mounted online cluster resource salvageable
Sector Size	512
Sectors per Cluster	8
Capacity	4.92 GB
Used space	604 KB
Advanced Information	<input type="button" value="View"/>

Salvageable files: None

-
-
- [Perform Inventory](#)

Table 10-3 describes each of the reported parameters on the Share Information page:

Table 10-3 NCP Volume Share Information

Parameter	Description
File System Path	The mount point of the selected volume.
File System Shadow Path	If the selected volume is shadowed, this is the mount point of its secondary storage area.
File System Type	The underlying file system type, such as NSS, Ext3, Reiser, or XFS.
NCP Volume ID	<p>The unique identifier given to the volume by the NCP engine. Values range between 0 and 254 (up to 255 volumes mounted concurrently).</p> <p>When the primary volume has a state of Shadowed, its NCP volume ID represents the DST shadow volume pair of volumes. There is not a separate NCP volume ID assigned to the secondary volume while it is in the shadow volume relationship.</p>
Status	The status of the selected volume, such as if it is mounted and online or offline for the NCP engine. For NSS volumes, it also shows which attributes are enabled, such as user quotas, directory quotas, and salvage.
Capacity	The total amount of space allocated to the volume.
Advanced Information	<p>Click View to reveal the following information:</p> <p>Local Cache: Shows the current status of cache parameters, such as trustee count, cached files, evicted files, cached folders, cache retrieved, and cache retrieved locked.</p> <p>Pool Name: For NSS, the name of the NSS pool where the volume resides.</p> <p>Pool Attributes: For NSS, the attribute identifier for the volume's pool.</p> <p>GUID: The eDirectory globally unique identifier for the selected volume.</p>
Open Files	<p>Reports the connection number (station) of the NCP client connection, the typeless fully distinguished eDirectory user name (such as <code>username.context</code>) who opened the connection, and the files that are currently open for that connection.</p> <p>You manage NCP connections to the primary storage area of the DST shadow volume. Users do not connect directly to the secondary storage area. To manage connections, go to the Manage NCP Services role, then click Manage Connections.</p>

10.10 Auditing File Move Events for the Shadow Volume

For volumes with a **Shadow Status** of **Shadowed**, all file moves between the primary volume and the secondary volume are logged to the shadow volume's audit file. An audit log for a DST shadow volume is located in the `._NETWARE` directory located at the root of the primary volume. For NSS volumes, the default file path for the log is `/media/nss/volumename/._NETWARE/volumename.audit.log`.

For example, if the primary area is named `VOL1`, the audit file is `/media/nss/VOL1/._NETWARE/VOL1.audit.log`.

- 1 In OES Remote Manager for Linux, log in as the `root` user.
- 2 Select **View File System > Dynamic Storage Technology Options**, locate the volume in the list, then click the **View Log** link next to it.
- 3 When you are prompted, select whether to view the file in a text editor, or to save a copy to your local computer.

The "local computer" is the computer where you are running the web browser for accessing the server via OES Remote Manager.

10.11 Backing Up DST Shadow Volumes

- ◆ [Section 10.11.1, "Planning Your Backup Solution," on page 106](#)
- ◆ [Section 10.11.2, "Planning Your Restore Solution," on page 107](#)
- ◆ [Section 10.11.3, "Using the /etc/NCPVolumes XML File for Backup," on page 109](#)
- ◆ [Section 10.11.4, "Configuring the Backup Attribute for NSS Volumes," on page 109](#)
- ◆ [Section 10.11.5, "Configuring a Backup for Trustee Information on NSS Volumes on Linux," on page 110](#)

10.11.1 Planning Your Backup Solution

Applications that directly access the local Linux file system see the primary file tree and the secondary file tree as independent directories. The backup utility does not see the merged view of the file tree that the end user sees. Thus, backup tools can apply one backup policy to the primary file tree and a different backup policy to the secondary file tree. In a DST volume, the only operations that take place directly on the secondary volume are backup (or remove and archive) and restore functions.

Using shadow volumes allows backups of important data to be made faster and more frequently because you can apply different backup policies for the primary volume and secondary volume.

For example, the server administrator can partition the volume's data into two categories:

- ◆ Important data that needs to be maintained on quality storage and backed up frequently.
- ◆ Less important data that can be stored on less expensive storage and backed up less frequently.

An analysis or inventory of a volume's data shows that a large portion of it is seldom used. Having a shadow volume allows the server administrator to spend more on the most important data and spend less on the less important data. The frequently used data can be backed up nightly. The seldom-used data can be backed up weekly or monthly.

Getting the less important data out of the way enables the backups of your important data to run more quickly and efficiently. Partitioning your data in this way can significantly reduce the cost of hosting it.

Because the most important files are located in the primary storage area, disaster recovery can also be faster. The server administrator can restore the critical files by restoring the primary storage area first, then restore the secondary storage area. This lets users quickly get the files they need most, and they do not need to wait while files they do not usually need are restored. In addition, more fault tolerant replication solutions can be deployed for the primary storage area where it matters most.

Ensure that policies are not being run during the backup window.

10.11.2 Planning Your Restore Solution

You can restore the data separately to each volume by using the backups you made of each area. If ShadowFS is running, you can also restore the data by using the ShadowFS local mount point in `/media/shadowfs/volumename` that presents a merged file tree that includes both volumes. The advantages and disadvantages of each restore option are described below.

- ♦ [“Restoring Data Separately to the NSS Volumes” on page 107](#)
- ♦ [“Restoring Data to the ShadowFS File Tree” on page 108](#)

Restoring Data Separately to the NSS Volumes

Consider the following advantages and disadvantages when restoring data separately to the NSS volumes. You restore data backed up from the primary path to the primary NSS volume. You restore data backed up from the secondary path to the secondary NSS volume.

- ♦ [“Advantages” on page 107](#)
- ♦ [“Disadvantages” on page 108](#)

Advantages

- ♦ Files are restored directly to the primary volume and secondary volume where they were when the files were backed up, so there is no need for the information to be transferred again through policies.
- ♦ There is no performance hit when you restore directly to each volume like there is when restoring to the ShadowFS file tree.

- ◆ The restoration size is not an issue because you are restoring to the proper volume rather than through the ShadowFS file tree view.
- ◆ You can back up the NSS volume by using the NSS Extended Attributes (`XAttr`) settings to preserve the NetWare metadata (`netware.metadata`) for file system rights and attributes, and restore that information to each volume as you restore data. For information about `XAttr`, see [“Extended Attributes \(XAttr\) Commands”](#) the *OES 2018 SP3: NSS File System Administration Guide for Linux*.

Disadvantages

- ◆ Ensure that no policy runs are in progress while data is being restored.
- ◆ Potential conflicts might occur if you restore duplicate versions of the file on each of the volumes. The duplicate files are resolved by DST global policies instead of being resolved by the backup software. By default, the duplicate files are allowed to coexist, and a conflict message is broadcast to users. For information about duplicate file resolution, see [Section 7.3, “Resolving Instances of Duplicate Files,”](#) on page 55.
- ◆ When you restore partial data, you need to know whether the most recent version of the data is located on the backup for the primary volume or the secondary volume.

Restoring Data to the ShadowFS File Tree

If ShadowFS is running, you can also restore the data by using the ShadowFS local mount point in `media/shadowfs/volumename` that presents a merged file tree that includes both volumes.

- ◆ [“Advantages”](#) on page 108
- ◆ [“Disadvantages”](#) on page 108

Advantages

- ◆ The backup software sees both volumes through the merged file tree view. You can restore the primary volume, secondary volume, or both volumes through this view, and let any duplicates be handled by your backup software.
- ◆ Whether the data is on the backup for the primary volume or the backup for the secondary volume, if both are restored, the users’ data is restored.

Disadvantages

- ◆ Ensure that no policy runs are in progress while data is being restored.
- ◆ The FUSE technology used by ShadowFS is slower than using the NCP view, but the backup software cannot see the NCP view.
- ◆ If you back up the NSS volume by using the NSS Extended Attributes (`XAttr`) settings to preserve the NetWare metadata (`netware.metadata`) for file system rights and attributes, this information cannot be restored through the shadowfs merged view of the data because `XAttr` requires that the destination location be an NSS volume. The shadowfs view is a mount point and is not seen by the backup software as an NSS volume. For information about `XAttr`, see [“Extended Attributes \(XAttr\) Commands”](#) the *OES 2018 SP3: NSS File System Administration Guide for Linux*.

- ◆ All files restored through the ShadowFS file tree view are copied to the primary volume. The data that you restore from the backup for the secondary volume is not returned to the secondary volume until you run policies or use inventory scans to move the data back to the secondary volume.
- ◆ Because all data is restored to the primary volume when you restore through the ShadowFS file tree view, it is possible to run out of space. The primary volume must be large enough to accommodate holding both volumes worth of data unless you restore in phases; that is, restore some directories, then shift data to the secondary, then restore more directories.

10.11.3 Using the /etc/NCPVolumes XML File for Backup

A backup utility can use the `/etc/NCPVolumes` XML file to easily locate each mounted NCP volume and find its primary and secondary file trees. The file contains an entry for each mounted volume. It lists the volume's name and the path for the volume's primary file tree (`PRIMARY_ROOT`). If the volume is a shadow volume, it also shows the path for the secondary file tree (`SHADOW_ROOT`).

For example, the following XML entry defines the DST shadow volume named `VOL1`. The volumes are NSS volumes, with `VOL1` as the primary storage location, and `ARCVOL` as the secondary storage location.

```
<VOLUME>
  <NAME>VOL1</NAME>
  <PRIMARY_ROOT>/media/nss/VOL1</PRIMARY_ROOT>
  <SHADOW_ROOT>/media/nss/ARCVOL</SHADOW_ROOT>
</VOLUME>
```

10.11.4 Configuring the Backup Attribute for NSS Volumes

You can use Novell Storage Management Services tools for backup and restore of NSS volumes. You can back up each NSS volume separately, and restore them separately. You need to be aware of the relationship on restore because you can get duplicate files. However, the mechanics of the backup and restore with SMS are the same as they are with any NSS volume. Refer to the SMS documentation for information about how to use SMS for NSS backup and restore.

The NSS Backup attribute must be enabled on the NSS volumes if you use SMS tools for backup of NSS volumes. The attribute is enabled by default when you create a new NSS volume.

To enable the Backup attribute for an existing NSS volume:

- 1 In iManager, click **Storage > Volumes**.
- 2 Select a server to manage to view a list of the NSS volumes on it.
- 3 In the **Volumes** list, select the volume that you want manage, then wait for the page to refresh to show the volume's details.
- 4 Click **Properties** to view the settings for the volume attributes.
- 5 On the **Attributes** tab, select the **Backup** attribute, then click **Apply**.

10.11.5 Configuring a Backup for Trustee Information on NSS Volumes on Linux

If you plan to use a backup utility with DST, you might need to add an NSS attribute that allows for backing up and restoring file system trustee assignments, trustee rights, and inherited rights filters. NSS provides the `nss /ListXattrNWMetadata` switch to enable this capability. For information, see “[ListXattrNWmetadata Option](#)” in the *OES 2018 SP3: NSS File System Administration Guide for Linux*.

10.12 Removing the Shadow Relationship for a Non-Clustered DST Shadow Volume

Removing a DST shadow volume simply removes the relationship between the primary and secondary storage area. It does not remove the underlying volumes themselves. The files remain on whichever storage area they are on at the time when you remove the shadow relationship.

- ◆ [Section 10.12.1, “Preparing to Remove a Shadow Volume,” on page 110](#)
- ◆ [Section 10.12.2, “Removing the Shadow Volume Relationship by Using OES Remote Manager for Linux,” on page 111](#)
- ◆ [Section 10.12.3, “Removing a Shadow Volume by Editing Configuration Files,” on page 112](#)

IMPORTANT: If you are using clustered shadow volumes, see [Section 17.10, “Removing the Shadow Relationship for a Clustered DST Volume Pair,” on page 215](#).

10.12.1 Preparing to Remove a Shadow Volume

Before you remove a shadow volume relationship, ensure that you shift data between the two volumes that make up the shadow volume, according to where you want the data to reside after the DST shadow volume relationship is removed.

- 1 In OES Remote Manager for Linux, log in as the `root` user.
- 2 Select **View File System > Dynamic Storage Technology Options**, locate the volume in the list, then click the **Inventory** link next to it.

View the volume inventory for the shadow volume to determine the space in use and the available space for both the primary and the secondary areas of the shadow volume. Ensure that there is sufficient free space available in either location for the data that you plan to move to that location.

- 3 Use any combination of the following techniques to shift data between the two areas:
 - ◆ **Shadow Volume Policies:** Run an existing shadow volume policy by using the **Execute Now** option in the **Frequency** area of the policy. You can also create a new shadow volume policy that moves specific data, and run the policy by using the **One Time** and **Execute Now** options in the **Frequency** area of the policy.

For information about configuring policies to move data between the primary and secondary areas, see [Chapter 11, “Creating and Managing Policies for Shadow Volumes,” on page 113](#).

- ♦ **Inventories:** Use the detailed inventory reports or customized inventories to move specific files to either area.

For information about using the volume customized inventory options to move data between the primary and secondary areas, see [Section 14.6, “Generating a Custom Inventory Report,”](#) on page 167.

10.12.2 Removing the Shadow Volume Relationship by Using OES Remote Manager for Linux

- 1 In OES Remote Manager for Linux, log in as the `root` user.
- 2 Select **Manage NCP Services > Manage Shares** to go to the NCP Shares page.
- 3 Ensure that you know which NSS volume is being used as the secondary volume so that you can manage it independently later.
 - 3a On the NCP Shares page, locate the primary NSS volume in the **Active Shares** list, then click the **Information** icon next to the share name.
 - 3b On the primary volume’s Share Information page, view the volume information in the **File System Shadow Path**.

In the following example, `ARCVOL` is an NSS volume that is the secondary storage area in the shadow volume.

File system path	<code>/media/nss/VOL1</code>
File system shadow path	<code>/media/nss/ARCVOL</code>

- 4 On the NCP Shares page, locate the primary NSS volume in the **Active Shares** list, then click the **Unmount** button next to the share name.
- 5 On the Manage Shares page, click the **Information** (i) icon next to the volume name of the NSS volume to access the **Remove Shadow Action Options**.
- 6 On the volume’s Share Information page under **Volume Tasks > Remove Shadow Action Options**, click **Remove Shadow**.

After the shadow volume is removed, the page refreshes to report a successful removal.
- 7 Select **Share Management** to go to the NCP Shares page, locate the volume that was the primary volume in the **Active Shares** list, then click the **Mount** button next to it.
- 8 Verify that the shadow volume was removed by using one of the following methods:
 - ♦ Select **View File System > Dynamic Storage Technology Options** to go to the Dynamic Storage Options page. The former primary volume now has an **Add Shadow** link next to it instead of a **Shadowed** link.
 - ♦ Select **Manage NCP Services > Manage Shares**, then click the **Information** icon next to the former primary volume’s name. The **File System Shadow Path** field displays `n/a` (not applicable).
- 9 Enable the NCP/NSS Bindings on the volume that was used as the secondary volume (for example, `ARCVOL`) in order to mount the volume for NCP.

For information, see [Section 10.5.2, “Enabling the NCP/NSS Bindings for an NSS Volume,”](#) on page 99.

10.12.3 Removing a Shadow Volume by Editing Configuration Files

- 1 Open a terminal console, then log in as the `root` user.
- 2 Edit the `/etc/opt/novell/ncpserv.conf` file to remove the following entry for your volume, then save your changes.

```
SHADOW_VOLUME primary_volumename secondary_volume_path
```

For example:

```
SHADOW_VOLUME VOL1 /media/nss/ARCVOL
```

- 3 Stop and restart the eDirectory `ndsd` daemon for the changes to take effect by entering

```
systemctl stop ndsd.service
```

```
systemctl start ndsd.service
```

- 4 Make the secondary NSS volume available for mounting in NCP by removing the `EXCLUDE_VOLUME` entry for the volume in the `/etc/opt/novell/ncp2nss.conf` file.

If necessary, edit the `/etc/opt/novell/ncp2nss.conf` file to remove the following entry for it:

```
EXCLUDE_VOLUME nss_volumename
```

An entry is automatically removed from the `/etc/opt/novell/ncp2nss.conf` file by using OES Remote Manager for Linux to set the **Manage NCP Services > Manage Shares > NCP/NSS Bindings > NCP Accessible** option to Yes for the NSS volume. For instructions, see [Section 10.5.2, “Enabling the NCP/NSS Bindings for an NSS Volume,”](#) on page 99.

- 5 Stop and restart the eDirectory `ndsd` daemon for the changes to take effect by entering

```
systemctl stop ndsd.service
```

```
systemctl start ndsd.service
```

- 6 Restart the NCP/NSS IPC daemon to synchronize the changes you made to the `/etc/opt/novell/ncp2nss.conf` file.

- 6a At the terminal console prompt, enter

```
systemctl restart ncp2nss.service
```


11 Creating and Managing Policies for Shadow Volumes

This section describes how to configure and manage Dynamic Storage Technology policies for shadow volumes on an Open Enterprise Server (OES) server.

- ♦ [Section 11.1, “Understanding Shadow Volume Policy Options,” on page 113](#)
- ♦ [Section 11.2, “Creating a Shadow Volume Policy,” on page 121](#)
- ♦ [Section 11.3, “Modifying a Shadow Volume Policy,” on page 123](#)
- ♦ [Section 11.4, “Running a Policy On Demand,” on page 124](#)
- ♦ [Section 11.5, “Viewing DST Policies and Policy Status,” on page 124](#)
- ♦ [Section 11.6, “Viewing Information about the Files Moved During a Policy Run,” on page 125](#)
- ♦ [Section 11.7, “Stopping a Running Policy,” on page 126](#)
- ♦ [Section 11.8, “Deleting a Shadow Volume Policy,” on page 127](#)

For information about setting global policies for DST on the server, see [Chapter 4, “Installing Dynamic Storage Technology,” on page 35](#).

11.1 Understanding Shadow Volume Policy Options

Shadow Volume policies manage how files are distributed across the shadow volume’s primary and secondary areas. A Shadow Volume policy allows you to specify when the policy is enforced (one time, hourly, daily, weekly, and so on), which volumes the policy applies to, which direction files are moved (primary area to its secondary area, or secondary area to its primary area), and which files are moved (by file name, file type, time stamps, or file size).

DST policies are configured in OES Remote Manager for Linux. DST provides the following policy options:

- ♦ [Section 11.1.1, “Last Executed,” on page 114](#)
- ♦ [Section 11.1.2, “Files moved,” on page 114](#)
- ♦ [Section 11.1.3, “Space moved,” on page 114](#)
- ♦ [Section 11.1.4, “Total files moved,” on page 114](#)
- ♦ [Section 11.1.5, “Total space moved,” on page 114](#)
- ♦ [Section 11.1.6, “Description,” on page 114](#)
- ♦ [Section 11.1.7, “Start Time,” on page 114](#)
- ♦ [Section 11.1.8, “End Time,” on page 115](#)
- ♦ [Section 11.1.9, “Start Day,” on page 115](#)
- ♦ [Section 11.1.10, “Frequency,” on page 115](#)
- ♦ [Section 11.1.11, “Command Status,” on page 116](#)

- ◆ [Section 11.1.12, “Volume Selection,” on page 116](#)
- ◆ [Section 11.1.13, “Volume Operations,” on page 117](#)
- ◆ [Section 11.1.14, “Subdirectory Restrictions,” on page 117](#)
- ◆ [Section 11.1.15, “Search Criteria,” on page 119](#)
- ◆ [Section 11.1.16, “Stop,” on page 121](#)

11.1.1 Last Executed

For an existing policy, the **Last Executed** parameter reports the last time the policy ran successfully. This parameter is not configurable.

11.1.2 Files moved

Files moved specifies the total number of files moved in the last execution of the DST policy. This parameter is not configurable.

11.1.3 Space moved

Space moved specifies the total amount of space moved in the last execution of the DST policy. This parameter is not configurable.

11.1.4 Total files moved

Total files moved specifies the total number of files moved in all the iterations of the DST policy. This parameter is not configurable.

11.1.5 Total space moved

Total space moved specifies the total amount of space moved in all the iterations of the DST policy. This parameter is not configurable.

11.1.6 Description

Description is the user-defined name for the policy. It should be descriptive of the policy it represents, and meaningful to the administrator. This name appears in the **Dynamic Storage Technology Policies** table on the main **Dynamic Storage Technology Options** page.

Description (required):

11.1.7 Start Time

Start Time specifies the time of day to begin a run to enforce the policy. For hourly policies, the policy enforcement begins at the selected minutes past each hour. Time is specified based on a 24-hour clock. For example, 18:00 (6:00 p.m.) is the default start time.

Start Time: :

11.1.8 End Time

End Time specifies the time of day to stop work on an enforcement run. Specifying an end time for a scheduled run allows you to prevent policy enforcement from happening during busy work hours. Time is specified based on a 24-hour clock. For example, 07:00 (7:00 a.m.) is the default end time.

End Time: :

If the policy enforcement process is still running when the end time is reached, the policy's queued work is paused until the next scheduled run. When the policy run begins at its next scheduled time, it continues with the queued work, and adds new work to the end of the queue.

11.1.9 Start Day

For policies that run weekly, **Start Day** specifies the day of the week to enforce the policy. You can specify only one day of the week for a given policy. Options are **Saturday** (default), **Sunday**, **Monday**, **Tuesday**, **Wednesday**, **Thursday**, or **Friday**.

Start Day: (for weekly commands)
 (for one time or monthly commands)

For policies that are run one time or monthly, **Start Day** specifies the month and day of the month when the policy is scheduled to be enforced.

11.1.10 Frequency

Frequency specifies how often the policy is enforced when the **Command Status** is set to **Active**. [Table 11-1](#) describes each frequency option. The **Execute Now** option can be selected or deselected in combination with any one of the scheduled frequency options.

Frequency:

- One Time
- Hourly
- Daily
- Weekly
- Monthly
- Execute now

Table 11-1 Frequency Options for DST Policies

Option	Description
One Time	Whenever the policy's Command Status is set to Active , the policy runs one time, then changes the Command Status to Inactive . You can activate the policy to run again by changing its status.

Option	Description
Hourly	The policy enforcement process runs once each hour. It begins at the number of minutes past the hour specified by the Start Time . The process continues until it is done, or until the number of minutes past the hour specified by the End Time . Unfinished work is queued until the next run.
Daily	The policy enforcement process runs once each day. It begins at the time specified by the Start Time . The process continues until it is done, or until the time specified by the End Time . Unfinished work is queued until the next run.
Weekly (default)	The policy enforcement process runs once each week. It begins on the day of the week specified by the Start Day . The process continues until it is done, or until the time specified by the End Time . Unfinished work is queued until the next run.
Monthly	The policy enforcement process runs once each month. It begins on the month and day specified by the Start Day , then it runs every month afterwards on that day of the month. The process continues until it is done, or until the time specified by the End Time . Unfinished work is queued until the next run.
Execute Now	Select this option to run the policy now, in addition to its regularly scheduled runs. The policy enforcement process is initiated within a few minutes after the policy's Command Status is set to Active and saved (submitted). The process continues until it is done, or until the time specified by the End Time . Unfinished work is queued until the next run.

11.1.11 Command Status

Command Status governs whether a policy is actively enforced or inactive. Inactive policies can be changed back to active. New policies can be created and set to inactive without running them. Options are **Active** (default) and **Inactive**.

Command Status: Active
 Inactive

11.1.12 Volume Selection

Volume Selection allows you to specify one or more shadow volume pairs where the policy will run. You can select one or multiple shadow volumes from a drop-down list of existing shadowed volumes, or select **All Shadowed Volumes**. You can have multiple policies associated with a given shadow volume. A given policy can apply to multiple shadow volumes. An individual volume policy is stored on the volume. An All-Shadowed-Volumes policy is stored on the server.

Volume Selection:

When you work with DST shadow volumes in a cluster, you should create separate policies for the shadow volume pair that exists in a given cluster resource. A given policy can apply to multiple shadow volumes in the cluster resource. You can have multiple policies associated with a given shadow volume in the cluster resource. An individual volume policy is stored on the volume, and

automatically fails over with the cluster resource. However, an All-Shadowed-Volumes policy is stored on the server. You must set up the same policy on all nodes, or copy the policy's file to all nodes.

11.1.13 Volume Operations

Volume Operations specifies the direction the files are moved between the primary storage location (primary area) and the secondary storage location (shadow area).

Volume Operations:



-  Move selected files from primary area to shadow area.
-  Move selected files from shadow area to primary area.

Table 11-2 Volume Operations for a Policy

Option	Description
Move selected files from primary area to shadow area (default)	When the policy is enforced, all files on the primary storage location that meet all of the search criteria are moved from the primary storage location to the secondary storage location.
Move selected files from shadow area to primary area	When the policy is enforced, all files on the secondary storage location that meet all of the search criteria are moved from the secondary storage location to the primary storage location.

11.1.14 Subdirectory Restrictions

In the **Subdirectory Restrictions** area, you specify the **Scope** and **Subdirectory List** information that determines whether the policy applies to everything in a volume, only to a specified directory (and its contents), or to all directories but the one specified.

Subdirectory Restriction:

Scope:

Subdirectory list: (Example: /subdir1/subdir2)

The **Scope** options allows you to specify included paths or excluded paths in a given policy, but not both. [Table 11-3](#) describes the options for **Scope**:

Table 11-3 *Subdirectory Restrictions for the Scope of a Policy*

Option	Description
None (default)	The policy is enforced for all subdirectories in the volume. Do not specify a path.
Apply only in subdirectory	The policy is enforced only for a specified subdirectory and its contents. You can specify multiple paths to be included when running the policy.
Exclude subdirectory	<p>The policy is enforced for all subdirectories in the volume, except for the specified subdirectory and its contents. You can specify multiple paths to be excluded when running the policy.</p> <p>Beginning in OES 11 SP1, the Exclude Subdirectory option also allows you to specify a directory name that might exist in multiple places on a volume. You indicate this intended action by specifying only a directory name with no forward slashes, and the directory name must contain at least one wildcard (such as <code>?</code> and <code>*</code>). All instances of directories that match the specified directory name are excluded from the policy run.</p>

The **Subdirectory List** allows up to 8 subdirectory paths on the primary volume to be specified for being included or excluded in a policy when it runs. Each additional path that you specify requires another pass through the data, so it increases the time needed to enforce the policy.

Specify the subdirectory paths relative to the root of the DST volume, not the full Linux path. Wildcards are not allowed in a subdirectory path. Each path must point to a valid subdirectory in the file system.

Precede subdirectory paths with a forward slash (`/`). For example:

```
/subdir1/subdir2
```

Directory names with spaces in them are supported in the subdirectory paths. For example:

```
/projects/project abc/dev
```

Beginning in OES 11 SP1, the Exclude Subdirectory option supports using wildcards to specify a directory that might exist in multiple places on a volume. For example, to exclude all GroupWise archive subdirectories, specify the following directory name with wildcards:

```
of???arc
```

[Table 11-4](#) demonstrates how to take advantage of the ability to specify subdirectory paths and directories that have multiple occurring instances in the DST volume.

Table 11-4 *Sample Extended Subdirectory Entries and Intended Actions*

Exclude Subdirectory Entry	Intended Action for the Policy
<code>/test</code>	The preceding forward slash indicates that this is a subdirectory path relative to the root of the volume. This entry excludes only the <code>/test</code> directory located at the root level of the DST volume.

Exclude Subdirectory Entry	Intended Action for the Policy
tes?	The absence of a forward slash and the presence of a question mark wildcard (?) in the directory name indicates that this is a directory that might have multiple instances as subdirectories in the volume. This entry excludes all instances of directories with 4 characters in the name that match the first 3 characters, and any character in the 4th position of the directory name.
tes*	The absence of a forward slash and the presence of an asterisk wildcard (*) in the directory name indicates that this is a directory that might have multiple instances as subdirectories in the volume. This entry excludes all instances of directories with names of any length that match the first 3 characters, and any characters to the end of the directory name.
test	No actions are taken for this entry. It is not preceded by a forward slash, so it does not qualify as a subdirectory path. It does not contain a wildcard, so it does not qualify as a directory entry.

11.1.15 Search Criteria

Files must match all of the specified criteria in order to be moved between the primary storage location and secondary storage location. Criteria options include file name or extension, time stamp, and file size. The conditions are combined (and-ed) together, which means that all conditions must be true for a file before it is queued for moving to the other location. Specify any of the following search criteria:

- ◆ [“Search Pattern” on page 119](#)
- ◆ [“Time Stamp Restrictions” on page 120](#)
- ◆ [“File Size Restriction” on page 121](#)

Search Pattern

Search Pattern allows you to set criteria based on the file name or extension. You can specify characters and wildcards to search by file name. You can specify files by types by specifying a wildcard and an extension, such as *.mp3. The default entry is *.* , which applies the policy to all file names and all file types.

Search Pattern (comma delimited extensions):

You can specify up to 50 extensions in a given policy. Separate the multiple entries with a comma and no spaces. For example, to specify multiple image file extensions, type the following:

*.bmp, *.jpg, *.png, *.tif

You can specify files with spaces in the name. Enter the file name without quotes.

filename with spaces.txt, another file with spaces.jpg, yet another file.doc

Time Stamp Restrictions

Time Stamp Restrictions identifies which time stamps to use when applying the policy.

Time Stamp Restrictions:

Time Stamp:

Last Modified Time

Last Accessed Time

Last Changed Time

Time from now:

Direction:

Days:

Weeks:

Months:

Years:

The time stamp types are:

- ◆ **Last Time Modified:** Time of last content modification for the selected file.
- ◆ **Last Time Accessed:** Time of last access.
- ◆ **Last Time Changed:** Time of last file status change.

The default is no time restriction (all Time Stamp options are deselected), so the default policy applies the policy for all existing files.

These time stamps are defined by POSIX and supported by Linux. Many operations change more than one time stamp. NCP can modify the access time and the modify time, but cannot control whether the change time is reset. The Last Time Changed value is controlled automatically. For example, if you copy a file from one location to another, NCP preserves the access and modify times, but the change time is reset because the file's path changed. That is, it had a status change but the file was not opened for access and its data was not modified.

You must also specify the specific time period to use in **Time from Now**. Direction options are **Greater than** and **Less than**. Specify a direction, then select one of the time periods described in [Table 11-5](#).

Table 11-5 Time Periods for the Time Stamp Restrictions in a Policy

Option	Description
Days	Specify 0 to 14 days. 0 days (the default) disables the option.
Weeks	Specify 0 to 10 weeks. 0 weeks (the default) disables the option.
Months	Specify 0 to 24 months. 0 months (the default) disables the option.
Years	Specify 0 to 24 years. 0 years (the default) disables the option.

For example, you can select all files that have a modified time greater than 6 months by selecting **Last Time Modified** in the **Time Stamp** field, **Greater than** for the **Direction** field, and 6 in the **Months** field.

File Size Restriction

Specifies the range of file sizes to search. **Direction** specifies to look for files that are greater than or less than the specified size in KB. Specify a value of 0 KB to disable the file size restriction. The default is no size restriction, which applies the policy for files of all sizes.

File Size Restriction:

Direction:

Size (KB):

11.1.16 Stop

A **Stop** button is available on the policy's View/Edit Shadow Volume policy page when the policy is running. You can use this option to stop an individual currently running policy. For information, see [Section 11.7, "Stopping a Running Policy," on page 126](#).

You can stop all currently running policies by using the **Stop all currently running policies** option on the Dynamic Storage Technology Options page.

11.2 Creating a Shadow Volume Policy

- ◆ [Section 11.2.1, "Prerequisite," on page 121](#)
- ◆ [Section 11.2.2, "Guidelines for Shadow Volume Policies," on page 121](#)
- ◆ [Section 11.2.3, "Creating a Shadow Volume Policy," on page 121](#)

11.2.1 Prerequisite

In order to configure policies that apply only to a specific shadow volume, the shadow volume must already be defined.

11.2.2 Guidelines for Shadow Volume Policies

- ◆ For each Dynamic Storage Technology shadow volume, you must establish at least one policy that controls how files are migrated from the primary storage area to the secondary storage area of the shadow volume, or vice versa.
- ◆ Any given shadow volume policy is best kept to a simple goal. Complex combinations of rules in a single policy can lead to confusion on how they are executed.
- ◆ You can have multiple policies associated with a given shadow volume.
- ◆ A given policy can apply to multiple shadow volumes.
- ◆ Multiple policies can be scheduled to be run concurrently.

11.2.3 Creating a Shadow Volume Policy

- 1 In OES Remote Manager for Linux, select **View File System**, then select **Dynamic Storage Technology Options** to open the Dynamic Storage Technology Options page.

Initially, no policies are defined, so you do not see a policy report.

No Dynamic Storage Technology policies defined.

[Create a new policy](#)

After one or more policies are defined, the policies are reported in a table.

Dynamic Storage Technology Policies			
Name	Volume	Last Executed	Total Files Moved
ProjectABC Exclude contracts	All Shadowed Volumes	Not executed	0

[Create a new policy](#)

- 2 Beneath the list of **Dynamic Storage Technology Policies**, click **Create a New Policy** to open a page where you can configure a new storage policy.

The screenshot shows the 'Create New Shadow Volume Policy' page in the OES Remote Manager. The page is divided into a left sidebar with navigation links and a main content area with configuration fields. The configuration fields include:

- Description (required):** A text input field.
- Start Time:** A time selection dropdown (18:00).
- End Time:** A time selection dropdown (07:00).
- Start Day:** A day selection dropdown (Saturday).
- Frequency:** Radio buttons for One Time, Hourly, Daily, Weekly, and Monthly.
- Command Status:** Radio buttons for Execute now, Active, and Inactive.
- Volume Selection:** A dropdown menu (All Shadowed Volumes).
- Volume Operations:** Radio buttons for Move selected files from primary area to shadow area and Move selected files from shadow area to primary area.
- Subdirectory Restriction:** A dropdown for Scope (None) and a text input for Subdirectory list.
- Search Pattern (comma delimited extensions):** A text input field.
- Time Stamp Restrictions:** Checkboxes for Last Modified Time, Last Accessed Time, and Last Changed Time.
- Time from now:** A dropdown for Direction (Greater than) and input fields for Days, Weeks, Months, and Years.
- File Size Restriction:** A dropdown for Direction (Greater than).

- 3 On the **Create New Shadow Volume Policy** page, specify a name for the policy in the **Description** field.

The name should be descriptive of the policy it represents, and meaningful to the administrator. For example, suppose you plan to create a policy for a shadow volume used by Project ABC, and exclude the path to the `contracts` directory. You might name the policy *Project ABC Exclude contracts*.

- 4 On the **Create New Shadow Volume Policy** page, configure policy settings.

For information about policy options, see [Section 11.1, "Understanding Shadow Volume Policy Options,"](#) on page 113.

- 5 At the top of the page, specify the **Command Status** as **Active** or **Inactive**.

A policy's state must be active in order for it to run.

- If you want the policy changes to be enforced sooner than the next scheduled run, ensure that you select the **Execute Now** check box in the **Frequency** area.

The policy run is triggered to begin within a few minutes after you save (submit) the policy.

- Click **Submit** (at the bottom of the page) in order to save the policy, and to schedule it if it is active.

The new policy is listed in the **Dynamic Storage Technology Policies** report on the Dynamic Storage Technology Options page.

Dynamic Storage Technology Policies			
Name	Volume	Last Executed	Total Files Moved
ProjectABC Exclude contracts	All Shadowed Volumes	Not executed	0

[Create a new policy](#)

11.3 Modifying a Shadow Volume Policy

You can modify a shadow volume policy at any time. For example, if the planned migration activity for a policy is not completed in the allowed time, you can adjust the policy run times and frequency until it meets your workload needs. Modified policies take effect the next time the policy runs, and do not affect currently running processes.

- In OES Remote Manager for Linux, select **View File System**, then select **Dynamic Storage Technology Options** to open the Dynamic Storage Technology Options page.

Dynamic Storage Technology Policies			
Name	Volume	Last Executed	Total Files Moved
ProjectABC Exclude contracts	All Shadowed Volumes	Not executed	0

[Create a new policy](#)

- In the list of **Dynamic Storage Technology Policies**, click the **Name** link for the policy in order to view and modify the individual settings for the policy.
- On the View/Edit Shadow Volume Policy page, view and modify the policy settings.

Last Executed: Reports the last time this policy was run successfully.

Files moved: Files moved from primary area to shadow area or from shadow area to primary area the last time this policy was executed.

Space moved: Size of space moved from primary area to shadow area or from shadow area to primary area the last time this policy was executed.

Total files moved: Cumulative number of files moved from primary area to shadow area or from shadow area to primary area by all instances of this policy execution.

Total space moved: moved: Cumulative size of space moved from primary area to shadow area or from shadow area to primary area by all instances of this policy execution.

For information about policy settings, see [Section 11.1, “Understanding Shadow Volume Policy Options,”](#) on page 113.

- Specify the **Command Status** as **Active** or **Inactive**.

A policy’s state must be active in order for it to run.

- 5 If you want the policy changes to be enforced sooner than the next scheduled run, select **Execute Now** in the **Frequency** area.

If the policy is not currently running, the policy runs within a few minutes after you click **Submit** in [Step 6](#).

If the policy is currently running, the updated policy does not run until the current run stops. That means the updated policy process is triggered within a few minutes after the currently running process completes or reaches the scheduled **End Time**.

- 6 If you make any changes, you must click **Submit** (at the bottom of the page) in order for the changes to take effect at the next scheduled run.

11.4 Running a Policy On Demand

You can run a policy on demand by enabling the **Execute Now** option in the policy's **Frequency** settings. If the policy is not currently running, the policy run is triggered within a few minutes after you save (submit) the policy change. Otherwise, it begins a few minutes after the currently policy run ends.

- 1 In OES Remote Manager for Linux, select **View File System**, then select **Dynamic Storage Technology Options** to open the Dynamic Storage Technology Options page.

Dynamic Storage Technology Policies			
Name	Volume	Last Executed	Total Files Moved
ProjectABC Exclude contracts	All Shadowed Volumes	Not executed	0

[Create a new policy](#)

- 2 In the list of **Dynamic Storage Technology Policies**, click the **Name** link for the policy to open the View/Edit Shadow Volume Policy page.
- 3 In the **Frequency** area, select **Execute Now**.
- 4 Scroll to the bottom of the page, then click **Submit**.

If the policy is not currently running, the policy runs within a few minutes after you click **Submit**.

If the policy is currently running, the on-demand run begins a few minutes after the currently running process completes or reaches the policy's scheduled **End Time**.

11.5 Viewing DST Policies and Policy Status

After you create DST policies, the Dynamic Storage Technology Policies table reports a list of policies, and information such as the shadow volumes to which the policy applies, when the policy was last executed, and the total number of files moved in the last run for that policy.

- 1 In OES Remote Manager for Linux, select **View File System**, then select **Dynamic Storage Technology Options** to open the Dynamic Storage Technology Options page.

Initially, no policies are defined.

No Dynamic Storage Technology policies defined.

[Create a new policy](#)

After one or more policies are defined, the policies are reported in a table.

Dynamic Storage Technology Policies			
Name	Volume	Last Executed	Total Files Moved
ProjectABC Exclude contracts	All Shadowed Volumes	Not executed	0

[Create a new policy](#)

- 2 View the following summary of information about all current policies on the server:

Parameter	Description
Name	The administrator-defined description of the policy. You specify the name in the Description field in the policy form.
Volume	A list of shadow volumes to which the policy applies. These are specified in Volume Selection field of the policy form.
Last Executed	The time the policy was last enforced.
Total Files Moved	The number of files moved between the primary storage location and the secondary storage location the last time the policy ran.

- 3 Click the **Name** link for the policy to view or modify the individual settings for the policy.
- 4 On the View/Edit Shadow Volume Policy page, view or modify the policy settings.
For information about policy settings, see [Section 11.1, “Understanding Shadow Volume Policy Options,”](#) on page 113.
- 5 If you make any changes, you must click **Submit** (at the bottom of the page) in order for the changes to take effect.

11.6 Viewing Information about the Files Moved During a Policy Run

If a file is moved during a policy run, the event is logged in the primary volume’s log file. It is also tracked in the volume’s audit log file (`/media/nss/<primary_volume_name>/ .NETWARE/ <primary_volume_name name>.audit.log`).

To view the primary volume’s log file by using OES Remote Manager:

- 1 In OES Remote Manager, select **View File System**, then select **Dynamic Storage Technology Options** to open the Dynamic Storage Technology Options page.
- 2 Under **Volume Information**, locate the shadow volume, then click the link to its log file.
- 3 In the log file, look for entries for the file moves.

For example, the following entry shows that the `/finance/rosebud_annual_report.pdf` file was successfully moved from the primary volume to the secondary volume:

```
<libncpengine name="volumeAuditOperations" timestamp="Mon Jun 22
09:22:22 2017 AM IST" errno="0">
<Move_status type="string">Successfully moved file</Move_status>
<Direction type="string">primary to shadow</Direction>
<File_path type="string">/finance/rosebud_annual_report.pdf</File_path>
</libncpengine>
```

11.7 Stopping a Running Policy

You can stop all currently running policies, or to stop an individual running policy. When a policy run begins, the policy is enforced on each of the shadow volumes that are specified in the policy's **Volume Selection** parameter. Therefore, the stop command applies to all shadow volumes that are associated with the policy. It is not possible to stop the policy for only one of multiple associated shadow volumes.

It takes some time (several seconds to a few minutes) for a policy to stop gracefully. For each of its associated shadow volumes, the policy run stops after it completes the move for the file that is currently being moved. The list of files to be moved for each associated shadow volume is discarded.

You cannot restart the policy from the point where you stopped a policy run. The next time that the policy is started, it scans its associated shadow volumes to create new lists of files to be moved.

Use the procedures in this section to stop one or all of the currently running shadow volume policies.

- ♦ [Section 11.7.1, “Stopping All Running Shadow Volume Policies,” on page 126](#)
- ♦ [Section 11.7.2, “Stopping a Running Individual Shadow Volume Policy,” on page 127](#)

11.7.1 Stopping All Running Shadow Volume Policies

The **Stop all running policies** option on the Dynamic Storage Technology Options page can be used to stop all currently running Shadow Volume storage policies. This option is available whether or not there are any currently running policies.

- 1 In OES Remote Manager for Linux, select **View File System**, then select **Dynamic Storage Technology Options** to open the Dynamic Storage Technology Options page.
- 2 Click **Stop all running policies**.
- 3 Click **Yes** to confirm that you want to stop all currently running Shadow Volume storage policies.

The status for each policy changes after its run is stopped gracefully for each of its associated shadow volumes.

11.7.2 Stopping a Running Individual Shadow Volume Policy

The **Stop** button on the View/Edit Shadow Volume Policy page can be used to stop a currently running individual Shadow Volume Policy rather than stopping all running policies. The **Stop** button is visible only while the policy is running.

- 1 In OES Remote Manager for Linux, select **View File System**, then select **Dynamic Storage Technology Options** to open the Dynamic Storage Technology Options page.
- 2 In the list of **Dynamic Storage Technology Policies**, click the **Name** link for the policy in order to view the policy.
- 3 On the View/Edit Shadow Volume Policy page, scroll to the bottom of the page.
If the policy is idle, the **Stop** button is not shown.
If the policy is currently running, the **Stop** button is available.
- 4 Click **Stop**.
- 5 Click **Yes** to confirm that you want to stop the selected currently running Shadow Volume policy.
The policy's status changes after the run is stopped gracefully for each of its associated shadow volumes.

11.8 Deleting a Shadow Volume Policy

You can delete a shadow volume policy at any time. If a policy is currently running, the policy is deleted after the process completes its run or reaches the previously set End Time.

- 1 In OES Remote Manager for Linux, select **View File System**, then select **Dynamic Storage Technology Options** to open the Dynamic Storage Technology Options page.
- 2 In the list of **Dynamic Storage Technology Policies**, click the **Name** link for the policy in order to view the policy.

Dynamic Storage Technology Policies			
Name	Volume	Last Executed	Total Files Moved
ProjectABC Exclude contracts	All Shadowed Volumes	Not executed	0

[Create a new policy](#)

- 3 On the View/Edit Shadow Volume Policy page, scroll to the bottom of the page, click **Delete**, then click **Yes** to confirm the deletion.
If the policy is not currently running, it is deleted immediately.
If the policy is currently running, it is deleted after the process stops.

12 Managing Directory and User Space Quotas on DST Volumes

The Dynamic Storage Technology (DST) quotas management system is a unified quotas management system that allows you to manage your disk space usage policy for DST shadow volume pairs. The `ncpcon quotas` command allows you to view and synchronize directory quotas settings and user quotas settings from the primary volume to the secondary volume. The NSS file system independently enforces the quotas set on each volume.

- ◆ [Section 12.1, “Understanding DST Volume Quotas Management,” on page 129](#)
- ◆ [Section 12.2, “NCPCON Quotas Command,” on page 132](#)
- ◆ [Section 12.3, “Setting User Space Quotas on DST Volumes,” on page 134](#)
- ◆ [Section 12.4, “Setting Directory Quotas on DST Volumes,” on page 137](#)
- ◆ [Section 12.5, “Getting Quotas from an Old Secondary Volume to a New Primary Volume,” on page 140](#)
- ◆ [Section 12.6, “Viewing User and Directory Quotas for a DST Volume,” on page 143](#)

12.1 Understanding DST Volume Quotas Management

You can leverage the OES Storage Services (NSS) directory and user space quotas management to control how space is allocated in a Dynamic Storage Technology shadow volume pair. These space restriction systems work independently, with the lower value being the most restrictive if multiple types of space restrictions apply. NSS allocates the space as it is needed; the quota does not reserve the space. The NSS file system enforces the quotas set on a volume independently of other volumes on the system. For information about NSS quotas, see [“Managing Space Quotas for Volumes, Directories, and Users”](#) in the *OES 2018 SP3: NSS File System Administration Guide for Linux*.

Management of quotas for volumes in a DST pair can be a challenge because NSS and NCP management tools are designed to work on independent volumes. DST provides a unified quota management system that allows you to manage your disk space usage policy for DST shadow volume pairs. With the `ncpcon quotas` command, you can view the existing directory quotas and user quotas on the primary volume, and then easily set related quotas on the secondary volume based on those settings. The NSS file system independently enforces the quotas set on each volume.

The `ncpcon quotas` command provides the following capabilities:

- ◆ You can view the amount of space that is allocated for an NSS directory or user on the primary volume, the secondary volume, and the combined space restriction.
- ◆ You can synchronize the NSS directory or user quotas from the primary volume to the secondary volume. You can use the same settings, or you can specify a percentage to set smaller or larger quotas on the secondary volume.

- ♦ To remove space restrictions for a user on the secondary volume, you can remove a user quota on the primary volume (which sets the user quota as unlimited), and then synchronize the unlimited quota value to the secondary volume.
- ♦ To remove space restrictions on a directory on the secondary volume, you can set an unlimited directory quota on the primary volume, and then synchronize the unlimited quota value to the secondary volume.

In OES 11 SP1 and earlier, DST does not offer a unified quota system for the DST shadow volume pair. [Table 12-1](#) and [Table 12-2](#) summarize the impact of the `ncpcon quotas` command on quotas management for DST shadow volume pairs.

Table 12-1 Directory Quotas Management and Functionality for DST Shadow Volume Pairs

Directory Quotas Management and Functionality for DST Volume Pairs in OES 11 SP1 and Earlier	Impact of the <code>ncpcon quotas</code> Command in OES 2015 and Later
You must enable the Directory Quotas attribute for each NSS volume in the DST shadow volume pair.	No change.
An administrator can use NCP tools to add, modify, or delete directory quotas on the primary volume. In NSS, each directory's quota information is stored as metadata with the directory in the file system. If there is no directory quota, no information is stored. Only directories with active quota settings have quota information. If you disable space restrictions for a directory, NSS automatically removes the quota setting from the directory. If you set a directory quota to a very large value (or unlimited), you effectively remove space restrictions from the directory, and the directory still has an active quota setting.	No change. The <code>ncpcon quotas sync</code> command does not synchronize deletions of directory quotas from the primary volume to the secondary volume. When you delete a directory quota, NSS removes the quota setting from the directory in the file system. In order to remove space restrictions on the directory instance on the secondary volume, you can set an unlimited quota for the directory instance on the primary volume, and synchronize the unlimited setting to the secondary volume.
NCP tools cannot be used to manage directory quotas on the secondary volume, because the secondary volume's NCP/NSS binding is disabled while it is in the DST relationship. The secondary volume is mounted in NSS but not in NCP.	As the <code>root</code> user, an administrator can synchronize the directory quotas settings to the secondary volume with the same or proportionally larger or smaller space allocation. Quotas synchronization can also be configured to run in a <code>cron</code> job.
An administrator can view directory quotas set on the primary volume. Secondary quotas settings are not available.	As the <code>root</code> user, an administrator can view directory quotas set on the primary volume, the secondary volume, and the combined space restriction for each directory.
A user who is a directory trustee can view the directory quota set on the primary volume by using an NCP client. Directory quotas settings for the secondary volume are not available.	No change.
The NSS file system independently enforces the directory quotas set on each volume.	No change.

Table 12-2 User Quotas Management and Functionality for DST Shadow Volume Pairs

User Quotas Management and Functionality for DST Volume Pairs in OES 11 SP1 and Earlier	Impact of the <code>ncpcon quotas</code> Command in OES 2015 and Later
You must enable the User Space Restrictions attribute for each NSS volume in the DST shadow volume pair.	No change.
An administrator can use NSS tools to add, modify, or delete the user quotas on the primary volume.	No change.
In NSS, user quota information is stored in the volume's <code>/_admin/Manage_NSS/Volume/<volname>UserInfo.xml</code> file.	<p>The <code>ncpcon quotas sync</code> command can synchronize deletions of user quotas from the primary volume to the secondary volume. When a user quota is deleted, NSS does not delete the quota entry for the user. Instead, NSS changes the user quota setting in the volume's <code>/_admin/Manage_NSS/Volume/<volname>UserInfo.xml</code> file from some restriction to unlimited. The command synchronizes the unlimited setting to the secondary volume, which effectively removes the space restrictions for the user.</p> <p>The <code>novcifs --user-quota-sync</code> command duplicates all of the user quotas that are set on the primary volume to the secondary volume.</p>
An administrator can use NSS tools to add, modify, or delete user quotas on the secondary volume.	As the <code>root</code> user, an administrator can synchronize the user quotas settings from the primary volume to the secondary volume with the same or proportionally larger or smaller space allocation.
This has no relationship to the settings made on the primary volume.	<p>Quotas synchronization can also be configured to run in a <code>cron</code> job. The interval to run the <code>cron</code> job depends on the nature of a company's setup and how often you add or change quotas in your environment. If quotas are regularly added throughout the day and require a quick synchronization, you might run the synchronization every hour. Otherwise, you might run it only once a day at a time when users are typically not active. If the quotas are managed only on certain days of the week or month, you can accordingly adjust how often and when the <code>cron</code> synchronization occurs.</p> <p>The load and time needed for quota synchronization depends on the number of directory quotas and user quotas to be synchronized. The more quotas there are to synchronize, the greater are the load and time required for the process. The load can be fairly intensive, but it is fairly brief. You can use a <code>cron</code> job to synchronize quotas at a time when users are typically not active.</p>

User Quotas Management and Functionality for DST Volume Pairs in OES 11 SP1 and Earlier	Impact of the <code>ncpcon quotas</code> Command in OES 2015 and Later
<p>Administrators can view the user quotas separately on the primary volume and the secondary volume, but not both in the same browser window.</p> <p>There is no report of the combined space restriction for each user.</p>	<p>As the <code>root</code> user, an administrator can view user quotas set on the primary volume, the secondary volume, and the combined space restriction for each user.</p>
<p>The NSS file system independently enforces the user quotas set on each volume.</p>	<p>No change.</p>

12.2 NCPCON Quotas Command

NCPCON supports the `ncpcon quotas` command for Novell Dynamic Storage Technology shadow volume pairs on OES 2015 and later.

Syntax

Log in to the server as the `root` user and launch a terminal console. You can run NCPCON console commands without entering the console by prefacing the command with `ncpcon`. You can also use the `ncpcon quotas` command from within `cron` jobs to automate the synchronization of quotas on the primary volume to the secondary volume.

```
ncpcon quotas help
ncpcon quotas view <nss_volume_name> <d|u> [c]
ncpcon quotas sync <ALL|MISSING|PERCENT> <nss_volume_name> [percent_value] <d|u> [q]
```

Command Options

quotas help

Display help for the quotas command at the command prompt.

quotas view

Show the assigned NSS directory quotas or user quotas for a specified NSS volume that is used in a DST volume pair.

```
ncpcon quotas view <nss_volume_name> <d|u> [c]
```

Replace *nss_volume_name* with the name (such as VOL1) of the primary volume or the secondary volume.

Specify either directory (`d`) or user (`u`) after the volume name to indicate the type of quotas to display. You can use the combined (`c`) option to display the specified type of quotas for both volumes.

OPTIONS

c

(Optional) Shows a combined view for the specified type of quotas on the primary volume and secondary volume of a DST shadow volume pair.

d

Applies the operation to NSS directory quotas.

u

Applies the operation to NSS user quotas.

EXAMPLES

```
ncpcon quotas view VOL_D u
```

Show the user quotas for NSS volume `VOL_D`. `VOL_D` can be the primary volume or secondary volume in a DST shadow volume pair.

```
ncpcon quotas view VOL1 d c
```

Show a combined view of the directory quotas on the primary volume and secondary volume of a specified DST shadow volume.

quotas sync

Synchronize the NSS directory quotas or user quotas from the primary volume to the secondary volume of a DST shadow volume pair.

```
ncpcon quotas sync <ALL|MISSING|PERCENT> <nss_volume_name> [percent_value] <d|u> [q]
```

Replace *nss_volume_name* with the name (such as `VOL1`) of the primary volume.

You can specify to use the same settings, or specify a percentage to set smaller or larger quotas on the secondary volume. You can duplicate all settings, or duplicate settings only where they do not exist.

Specify either directory (`d`) or user (`u`) after the volume name to indicate the type of quotas to synchronize.

OPERATIONS

ALL

For all of the directory quotas or user quotas (whichever type is specified) that are currently set on the primary volume, duplicate the quotas settings on the secondary volume.

MISSING

For each of the directory quotas or user quotas (whichever type is specified) that are currently set on the primary volume, if a quota is not set on the secondary volume, duplicate the quota setting on the secondary volume. This option does not overwrite existing quotas on the secondary volume.

PERCENT

For each of the directory quotas or user quotas (whichever type is specified) that are currently set on the primary volume, set the quotas settings on the secondary volume as a specified percentage of the quota that exists on the primary volume. The percentage value must also be specified after the volume name.

A percent value of 100 is a one-to-one quota assignment. A percent value of 50 assigns a quota that is one-half the size of the quota set on the primary volume. A percent value of 200 assigns a quota that is twice the size of the quota set on the primary volume.

OPTIONS

d

Applies the operation to all active NSS directory quotas. No information is synchronized for directories where no quota is set.

u

Applies the operation to NSS user quotas.

percent_value

Required if the PERCENT operation is used. Specifies the value to use when calculating the quota for the secondary volume based on a percentage of the primary volume's quota.

q

(Optional) Indicates quiet mode. No output appears in the execution window.

EXAMPLES

```
ncpcon quotas sync ALL VOL_D u
```

For all of the NSS user quotas that are currently set on the primary volume `VOL_D`, duplicate the quotas setting on the secondary volume of a DST shadow volume pair.

```
ncpcon quotas sync PERCENT VOL1 50 d
```

For each of the NSS directory quotas that are currently set on the primary volume `VOL1`, set a quota that is one-half that size on the secondary volume of a DST shadow volume pair.

```
ncpcon quotas sync MISSING VOL1 u
```

For each of the NSS user quotas that are currently set on the primary volume `VOL1`, if a quota does not exist on the secondary volume, duplicate the quota setting on the secondary volume of a DST shadow volume pair.

NOTE: For AD users, use `novcifs` command to synchronize the user quotas from the primary volume to the secondary volume of a DST shadow volume pair.

```
--user-quota-sync <primary_volume>
```

For more information on command operations and examples, see [Synchronizing Users Quotas](#) in the [OES 2018 SP3: OES CIFS for Linux Administration Guide](#).

12.3 Setting User Space Quotas on DST Volumes

User space restrictions limit the space available to a user of the NSS volumes across all directories and files owned by the user. In a DST volume pair, the user quota set on the primary volume is enforced separately from the user quota set on the secondary volume. Before you can add quotas using the `ncpcon quotas` command, you must enable the User Space Quotas attribute for the primary volume and the secondary volume.

In NSS, user quota information is stored in the volume's `/_admin/Manage_NSS/Volume/<vol name>UserInfo.xml` file. You can add, modify, or delete user space quotas on the primary volume, and then use the `ncpcon quotas sync` command to synchronize the settings from the primary volume to the secondary volume. When a user space quota is deleted, NSS does not delete the quota entry for the user. Instead, NSS changes the user quota setting in the volume's `/_admin/`

Manage_NSS/Volume/<vol name>UserInfo.xml file from some restriction to unlimited. The command synchronizes the unlimited setting to the secondary volume, which effectively removes the space restrictions for the user.

- 1 Verify that the User Space Quotas attribute is enabled on each volume.
 - 1a In iManager, select **Storage > Volumes**.
 - 1b Click the **Search** icon, then browse to select a server to manage.
 - 1c In the **Volumes** list, select the volume, then wait for the volume’s details to be displayed.
 - 1d Click **Properties**.
 - 1e On the **Attributes** tab, verify that the **User Space Quotas** check box is selected. If it is deselected, select it, then click **Apply**.
 - 1f Repeat **Step 1c** to **Step 1e** to verify that the User Space Quotas attribute is enabled on the secondary volume.

- 2 Add, modify, or delete user quotas on the primary NSS volume.

For more information, see “[Managing User Space Quotas](#)” in the *OES 2018 SP3: NSS File System Administration Guide for Linux*.

- 2a In iManager, select **Storage > User Quotas**, click the **Search** icon, then browse to select the primary volume.

You can also click **Storage > Volumes**, select the server, select the primary volume, then click **User Quotas**.

- 2b View the **Users with Quotas** list.

Storage

User Quotas ?

Volume:

Users with Quotas **All Users**

View only the users with assigned quotas for the selected volume. Select New to specify space quotas for users with access rights for this volume. Select one or more users from the list to modify or delete their existing space quotas.

User Quotas			
New Edit Delete	2 Item(s)		
<input type="checkbox"/> Name	Quota	Used	Available
<input type="checkbox"/> bob.novell	500.00 MB	0.00 Bytes	500.00 MB
<input type="checkbox"/> john.novell	200.00 MB	0.00 Bytes	200.00 MB

2c Do one of the following:

- ◆ **Add user quota:** Click **New**, browse to select the user, specify a quota, then click **Finish**. The user name appears in the list with the user quota information for the volume.

Storage > User Quotas

Specify the distinguished user name (userid.context), or browse to select one or more users. Specify the upper limit of storage space to allow each of the selected users. Space is allocated to each user as it is needed.

Name:

Quota: GB

(ex. 200 MB)

<< Back Finish Cancel

- ◆ **Edit user quota:** Select the check box next to the user name, click **Edit**, specify a quota, then click **Finish**. The user name appears in the list with the modified user quota information for the volume.

Storage > User Quotas

Specify the upper limit of storage space to allow each of the users you selected on the Users with Quotas page. Space is allocated to each user as it is needed. If a user's data exceeds the new quota, the user can access files on the volume, but cannot write to it until the user's data no longer exceeds the new quota.

Name	Quota
bob.novell	500.00 MB

Quota: GB

(ex. 200 MB (set quotas), +50 MB (increase quotas), -20 MB (decrease quotas))

<< Back Finish Cancel

- ◆ **Delete user quota:** Select the check box next to the user name, click **Delete**, then click **OK** to confirm the quota deletion. The users quota is set to unlimited in the volume's /_admin/Manage_NSS/Volume/<vol name>UserInfo.xml file.

2d Repeat [Step 2c](#) for each user quota you want to manage on the primary volume.

3 Synchronize the user space quotas settings from the primary volume to the secondary volume.

3a Log in as the `root` user to the server, then open a terminal console.

3b Use the `ncpcon quotas sync` command to synchronize the user quotas to the secondary volume.

```
ncpcon quotas sync <ALL|MISSING|PERCENT> <primary_volume_name> [percent_value] u [q]
```

For example, using the ALL option synchronizes all user quotas settings from the primary volume to the secondary volume:

```
ncpcon quotas sync ALL VOLD u
```


Using the PERCENT option with a percent value of 200 assigns a quota for the user on the secondary volume that is twice the size of the quota set on the primary volume:

```
ncpcon quotas sync PERCENT VOLD 200 u
```

For command information and other examples, see [Section 12.2, “NCPCON Quotas Command,” on page 132](#).

- 3c** For AD users, use the `novcifs --user-quota-sync` command to synchronize the user quotas to the secondary volume.

```
--user-quota-sync <primary_volume>
```

For more information on examples, see [Synchronizing Users Quotas](#) in the [OES 2018 SP3: OES CIFS for Linux Administration Guide](#).

- 4** (Optional) In iManager, modify the space quota for the user on the primary volume.

It is not required that the user be assigned the same space quota on the primary volume as on the secondary volume. However, if you change it, the next time you synchronize the user space restrictions, the modified user space quota is applied to the secondary volume if you use the ALL or PERCENT options. In order to synchronize user quotas only for users on the secondary volume that do not have a quota set, use the MISSING option.

NOTE: To manage user quotas for AD users using NFARM, see [User Quota](#) in the [OES 2018 SP3: NSS AD Administration Guide](#).

12.4 Setting Directory Quotas on DST Volumes

Directory quotas limit the space available in an individual NSS directory. In a DST volume pair, the directory quota set on the directory instance on the primary volume functions separately from the directory instance on the secondary volume. Before you can add quotas using the `ncpcon quotas` command, you must enable the Directory Quotas attribute for the primary volume and the secondary volume.

In NSS, each directory's quota information is stored as metadata with the directory in the file system. If there is no directory quota, no information is stored. Only directories with active quota settings have quota information. If you disable space restrictions for a directory, NSS automatically removes the quota setting from the directory. If you set a directory quota to a very large value (or unlimited), you effectively remove space restrictions from the directory, and the directory still has an active quota setting.

In order to remove space restrictions on the directory instance on the secondary volume, you can set an unlimited quota for the directory instance on the primary volume, and synchronize the unlimited setting to the secondary volume. The setting can be larger than the secondary volume's maximum size, but the directory's maximum size will be limited to the secondary volume's actual maximum size and space availability when the quota is enforced.

- 1** Verify that the Directory Quotas attribute is enabled on each volume.
 - 1a** In iManager, select **Storage > Volumes**.
 - 1b** Click the **Search** icon, then browse to select a server to manage.
 - 1c** In the **Volumes** list, select the volume, then wait for the volume's details to be displayed.
 - 1d** Click **Properties**.

- 1e On the **Attributes** tab, verify that the **Directory Quotas** check box is selected. If it is deselected, select it, then click **Apply**.
 - 1f Repeat **Step 1c** to **Step 1e** to verify that the Directory Quotas attribute is selected on the secondary volume.
- 2 Add, modify, or remove directory quotas for the instance of the directory on the primary volume.

For more information, see “[Managing Directory Quotas](#)” in the *OES 2018 SP3: NSS File System Administration Guide for Linux*.

- 2a In iManager, open the Properties page for the folder instance on the primary volume. You can use either of the following methods:

- ◆ **Files and Folders:**

1. In **Roles and Tasks**, select **Files and Folders > Properties**.
2. On the Properties page, click the **Search** icon, then browse to select the folder you want to manage.
3. Click **OK** to open the Properties page for the folder.

- ◆ **Tree View:**

1. In the iManager toolbar, click the **View Objects** icon.
2. In the Tree View, click the **Browse** tab in the left panel.
3. Browse the tree to locate and select the folder you want to manage.
4. In the right panel, select the check box next to the folder, then select **Actions > Properties**.

- 2b On the **Information** tab of the folder’s Properties page, view the current status of the directory quota.

If a Directory Quota is set, the **Restrict Size** field is selected and the **Limit** field shows the quota size (in KB, MB, GB, TB, or PB).

If the Directory Quota is not set, the **Restrict Size** field is deselected and the **Limit** field is dimmed (grayed out).

- 2c Do one of the following:

- ◆ **Add directory quota:** On the **Information** tab, select **Restrict Size** to enable space restrictions for the selected directory. In the **Limit** field, type the directory quota in KB, MB, GB, TB, or PB.
- ◆ **Modify a directory quota:** On the **Information** tab, the **Restrict Size** check box should already be selected. In the **Limit** field, type the new directory quota in KB, MB, GB, TB, or PB.
- ◆ **Set an unlimited directory quota:** On the **Information** tab, select **Restrict Size** to enable space restrictions for the selected directory. In the **Limit** field, type the new directory quota in KB, MB, GB, TB, or PB. For example, you can set a value that exceeds the volume size.

You might set the quota as unlimited instead of disabling the quota if you need to remove the directory quota on the secondary volume. The setting can be larger than the secondary volume’s maximum size, but the directory’s maximum size will be limited to the secondary volume’s actual maximum size and space availability when the quota is enforced.

- ♦ **Remove a directory quota:** On the **Information** tab, deselect **Restrict Size** to disable space restrictions for the selected directory. The **Limit** field is automatically dimmed (grayed out). The quota information is removed from the directory metadata.

A deleted quota for a directory cannot be synchronized to the secondary volume. In order to remove space restrictions for a directory on the secondary volume, set an unlimited quota for the directory instance on the primary volume, and then synchronize the unlimited quota setting to the secondary volume.

- 2d** On the **Information** page, click **Apply** or **OK** to apply the changes.
- 2e** Repeat **Step 2a** to **Step 2d** for each directory where you want to set a quota on the primary volume.
- 3** Synchronize the directory space quotas settings from the primary volume to the secondary volume.
 - 3a** Log in as the `root` user to the server, then open a terminal console.
 - 3b** Use the `ncpcon quotas sync` command to synchronize the directory quotas from the primary volume to the secondary volume.

```
ncpcon quotas sync <ALL|MISSING|PERCENT> <primary_volume_name> [percent_value] d [q]
```

For example, using the `ALL` option synchronizes all directory quotas settings from the primary volume to the secondary volume:

```
ncpcon quotas sync ALL VOLD d
```

Using the `PERCENT` option with a percent value of 200 assigns a quota for the instance of the directory on the secondary volume that is twice the size of the quota set on the primary volume.

```
ncpcon quotas sync PERCENT VOLD 200 d
```

For command information and other examples, see [Section 12.2, “NCPCON Quotas Command,”](#) on page 132.

- 4** (Optional) In iManager, modify the space quota for the directory instance on the primary volume.

It is not required that the directory be assigned the same space quota on the primary volume as on the secondary volume. However, if you change it, the next time you synchronize the directory quotas, the modified directory quota is applied to the secondary volume if you use the `ALL` or `PERCENT` options. In order to synchronize directory quotas only for secondary directories that do not have a quota set, use the `MISSING` option.

NOTE: To manage directory quotas for AD users using NFARM, see [Information or Directory Quota](#) in the [OES 2018 SP3: NSS AD Administration Guide](#).

12.5 Getting Quotas from an Old Secondary Volume to a New Primary Volume

One scenario for setting up a DST volume is to use an empty volume as the primary and an existing volume that contains data as the secondary, as is described in [Section 1.3.2, “Empty Volume as Primary with an Existing Volume as Secondary,”](#) on page 21. As users open and save files from the merged view of the DST volume pair, active data is pulled forward to the primary volume.

[Table 12-3](#) shows an example of a new/old volume scenario where trustees, directory quotas, and user quotas are set on the secondary NSS volume. There is no automatic way to synchronize the setting from the secondary volume to the primary volume. The `ncpcon quotas` command does not synchronize quotas settings from secondary to primary. You can copy the trustees database manually from the old volume to the new volume before you set up the shadow relationship. You can manually configure directory quotas and user quotas on the primary volume.

Table 12-3 Initial Settings on the Old and New NSS Volumes for the DST Volume Pair

Volume Description	New NSS Volume	Old NSS Volume
Volume name	NEWVOL1	OLDVOL1
Contains data	No	Yes
Trustees defined	No	Yes
User Space Restrictions attribute	Enabled	Enabled
User quotas	None	Yes, for multiple users
Directory Quotas attribute	Enabled	Enabled
Directory quotas	None	Yes, for multiple directories

If you have many directory quotas or user quotas set on the old volume, the manual process to set them up on the new volume can be daunting. You can use a two-phase setup of the DST volume pair to get quotas from the old volume (such as `OLDVOL1`) to the new volume (such as `NEWVOL1`):

- ◆ **Phase 1:** Temporarily set up a DST shadow volume pair with the old volume as primary and the new volume as secondary. Do not allow users to access the merged view of the volume. The `._NETWARE/.trustee_database.xml` file is automatically copied from the primary volume to the secondary volume when you create the DST volume pair relationship. Use the `ncpcon quotas` command to synchronize the directory and user quotas to the new volume. The command creates the directory file tree structure as needed to set the quotas on the new secondary volume. Remove the DST shadow relationship between the two volumes.
- ◆ **Phase 2:** Create the DST volume pair with the new volume as primary and the old volume as secondary. Allow users to access the merged view of the volume.

Use the following procedure to get trustees and quotas settings from an old volume to a new volume that you want to use as the primary volume of a DST shadow volume pair. These instructions use `OLDVOL1` and `NEWVOL1` for the volume names. Modify the commands as needed to use the actual volume names on your system.

- 1 Log in to the server as the `root` user, then open a terminal console.
- 2 Verify that quotas attributes are enabled for the new volume `NEWVOL1`.


The following procedure uses NSSMU. You can alternatively use the Storage plug-in in Novell iManager.

- 2a At the command prompt, launch NSSMU:

```
nssmu
```


- 2b In the NSSMU Main menu, select **Volumes**, then press Enter.
- 2c In the **Volumes** list, select the new NSS volume `NEWVOL1`, then press Enter to open its Properties page.
- 2d Ensure that the **Directory quotas** attribute is enabled on the new NSS volume `NEWVOL1`.

If **Directory quotas** is set to **No**, press the Down-arrow to select its value, then press `y` to change the setting to **Yes**.



- 2e Ensure that the **User Space Restrictions** attribute is enabled on the new NSS volume `NEWVOL1`.

If **User Space Restrictions** is set to **No**, press the Down-arrow to select the value, then press `y` to change the setting to **Yes**.



2f Press the Down-arrow to select **Apply**, then press Enter to save your changes.

2g Press Esc twice to exit NSSMU.

- 3** (Optional) Launch OES Remote Manager and log in as the `root` user, make sure that the `REPLICATE_PRIMARY_TREE_TO_SHADOW` global policy is enabled which allows the file tree structure to be replicated automatically from the primary to the secondary when the DST pair is created. This setting affects all DST volume pairs on the server. If you want to fully replicate the file tree structure only for this scenario, you can replicate the structure later in [Step 13 on page 143](#).

For information about setting the global policy, see [Section 7.1, “Replicating Branches of the Primary File Tree in the Secondary File Tree,” on page 49](#).

- 4** Create a DST volume with the old volume `OLDVOL1` as the primary volume and the new volume `NEWVOL1` as the secondary volume.

For information, see [Section 10.2, “Creating a DST Shadow Volume with NSS Volumes,” on page 88](#).

When you create the shadow volume pair, the `.trustee_database.xml` file should automatically be copied from the primary volume (`/media/nss/OLDVOL1/._NETWARE/.trustee_database.xml`) to the secondary volume (`/media/nss/NEWVOL1/._NETWARE/.trustee_database.xml`).

- 5** As the `root` user, open the `/media/nss/NEWVOL1/._NETWARE/.trustee_database.xml` file in a text editor, visually verify that it now contains trustees information from the old volume, then close the file.
- 6** Synchronize the directory quotas from the primary `OLDVOL1` to the secondary `NEWVOL1`. At the command prompt, enter

```
ncpcon quotas sync all OLDVOL1 d
```

If the `REPLICATE_PRIMARY_TREE_TO_SHADOW` global policy is disabled (0), the directory structure on the primary `OLDVOL1` was not automatically replicated to the secondary `NEWVOL1` when you defined the DST volume pair. In this case, the `ncpcon quotas` command creates the directory structure as needed on the secondary `NEWVOL1` to match the directories that have quotas already set on the primary `OLDVOL1`, and synchronizes the directory quotas for them.

If the `REPLICATE_PRIMARY_TREE_TO_SHADOW` global policy is enabled (1), the directory structure on the primary `OLDVOL1` was automatically replicated to the secondary `NEWVOL1` when you defined the DST volume pair. In this case, the `ncpcon quotas` command synchronizes the directory quotas to the matching directories on the new volume.

- 7** View the directory quotas to verify that quotas were written to the secondary `NEWVOL1`. At the command prompt, enter

```
ncpcon quotas view OLDVOL1 d c
```

- 8** Synchronize the user quotas from the primary `OLDVOL1` to the secondary `NEWVOL1`. At the command prompt, enter

```
ncpcon quotas sync all OLDVOL1 u
```

- 9** View the user quotas to verify that quotas were written to the secondary `NEWVOL1`. At the command prompt, enter

```
ncpcon quotas view OLDVOL1 u c
```

- 10 Remove the DST shadow relationship between the two volumes.

For information, see [Section 10.12, “Removing the Shadow Relationship for a Non-Clustered DST Shadow Volume,”](#) on page 110.

- 11 Verify that the trustees database file on the new volume contains trustee information.

In the unlikely event that the file is empty, you must manually copy the trustees database file from the old volume to the new volume as described in [“Manually Copying the Trustees Database to the New Volume”](#) on page 91.

- 12 Create a DST volume with the new volume as the primary and the old volume as the secondary.

From the previous configuration, the new volume has the following settings from the old volume:

- ◆ Trustees settings
- ◆ The replicated file tree structure (according to the conditions discussed in [Step 6](#))
- ◆ The directory quotas settings
- ◆ The user quotas settings

As users open and save files from the merged view, the active files are relocated to the new primary volume, which is the desired behavior of this scenario. All of the trustees settings are enforced based on the settings on the primary. All of the directory quotas are enforced independently by the NSS file system on the primary volume and the secondary volume. All of the user quotas are enforced independently by the NSS file system on the primary volume and the secondary volume.

From this point forward, you can follow the normal approved methods of setting trustees and quotas from the merged view. You can use the `ncpcon quotas` command to synchronize directory quotas settings and user quotas settings from the primary volume to the secondary volume.

- 13 (Optional) If the file tree structure was not fully replicated (that is, the `REPLICATE_PRIMARY_TREE_TO_SHADOW` global policy was disabled during the temporary configuration), you can fully replicate the old file tree structure to the new volume by enumerating the directory as described in [Section 10.4, “Replicating the Secondary File Tree Structure to the Primary Volume,”](#) on page 97.

12.6 Viewing User and Directory Quotas for a DST Volume

When you use iManager to view user space quotas or directory quotas for a DST volume, the tools report only the quotas set on the primary volume.

- ◆ [Section 12.6.1, “Viewing User Quotas from the Storage Plug-In in iManager,”](#) on page 144
- ◆ [Section 12.6.2, “Viewing Directory Quotas from the Files and Folders Plug-In in iManager,”](#) on page 144
- ◆ [Section 12.6.3, “Viewing User and Directory Quotas in NFARM,”](#) on page 144

12.6.1 Viewing User Quotas from the Storage Plug-In in iManager

The Storage plug-in in iManager is not aware of the DST relationship between the NSS volumes in a DST volume pair. You can use the Storage plug-in to view user space quotas set on NSS volumes. When you view user space quotas for the primary volume of a DST volume pair, the tool reports the usage for the primary volume only. It does not report user quotas set on the secondary volume, and it does not consider the user's data on the secondary volume when it calculates quota usage.

12.6.2 Viewing Directory Quotas from the Files and Folders Plug-In in iManager

The Files and Folders plug-in in iManager is not aware of the DST relationship between the NSS volumes in a DST volume pair. You can use the Files and Folders plug-in in iManager to view directory quotas set for directories on NSS volumes that are mounted in both NSS and NCP. Because the secondary volume is not mounted in NCP while it is in a DST volume pair, you cannot manage directory quotas on the secondary volume with this tool. When you view directory quotas for a directory on the primary volume of a DST volume pair, the Files and Folders **Directory Quota** option reports the usage for the primary volume only. It does not report the directory quotas set on the secondary volume, and it does not consider the directory's data on the secondary volume when it calculates quota usage.

12.6.3 Viewing User and Directory Quotas in NFARM

For information on viewing User and Directory quotas using NFARM, see [User Quota and Information or Directory Quota](#) in the [OES 2018 SP3: NSS AD Administration Guide](#).

13 Using ShadowFS to Provide a Merged View for Novell Samba Users

The Shadow File System (ShadowFS) provides a merged file tree view of the data on a Dynamic Storage Technology (DST) shadow volume for Novell Samba users. This section describes how ShadowFS works and how to configure it with Novell Samba.

IMPORTANT: OES CIFS works directly with NCP Server to provide a merged view for CIFS users. It is not necessary to install or set up ShadowFS. OES CIFS and Novell Samba cannot be installed on the same server.

- ◆ [Section 13.1, “Understanding ShadowFS,” on page 145](#)
- ◆ [Section 13.2, “Prerequisites for Using ShadowFS,” on page 146](#)
- ◆ [Section 13.3, “Preparing Your System for Using ShadowFS,” on page 146](#)
- ◆ [Section 13.4, “Installing ShadowFS and FUSE,” on page 147](#)
- ◆ [Section 13.5, “Setting Rights to ShadowFS Shares,” on page 149](#)
- ◆ [Section 13.6, “Creating a Samba Share,” on page 150](#)
- ◆ [Section 13.7, “Adding a User to Samba,” on page 150](#)
- ◆ [Section 13.8, “Connecting Users to the Share,” on page 151](#)
- ◆ [Section 13.9, “Testing Shadow Volume Policies,” on page 151](#)
- ◆ [Section 13.10, “Enabling or Disabling ShadowFS,” on page 152](#)
- ◆ [Section 13.11, “Starting and Stopping ShadowFS Manually,” on page 152](#)
- ◆ [Section 13.12, “Configuring Trustee Rights for Novell Samba Users,” on page 154](#)

13.1 Understanding ShadowFS

Shadow File System (ShadowFS) provides a merged file tree view of the DST volume for Novell Samba users. It allows users to access data on both locations via a share on the primary storage area by using the SMB/CIFS protocol instead of the NetWare Core Protocol (NCP). It is necessary to load ShadowFS only if Novell Samba is implemented on the server and SMB/CIFS users are given access to shadow volumes.

IMPORTANT: Performance for the Novell Samba clients to access the data via ShadowFS is slower than for NCP clients and OES CIFS clients.

The ShadowFS technology is implemented on the FUSE (File System in Userspace) virtual file system. FUSE is an open source software package that is installed automatically when you install Dynamic Storage Technology.

When ShadowFS loads, it checks the `/etc/opt/novell/ncpserv.conf` file to see what NCP shadow volumes exist, then it automatically creates a local mount point in `/media/shadowfs/volumename` that presents a merged file tree that includes both volumes. This local mount point allows Novell Samba and other local applications (including backup utilities) to see the same combined view that NCP clients see when they access a shadow volume. Each instance of ShadowFS runs as a separate process.

The ShadowFS configuration file is `/etc/opt/novell/shadowfs.conf`.

The ShadowFS log file is `/var/opt/novell/log/shadowfs.log`.

13.2 Prerequisites for Using ShadowFS

- ❑ Before using ShadowFS, ensure that the following services have been installed and configured:
 - ◆ NCP Server and Dynamic Storage Technology
 - ◆ NetIQ eDirectory
 - ◆ Novell Samba
 - ◆ Linux User Management
 - ◆ FUSE
 - ◆ Novell Remote Manager
 - ◆ Novell iManager for Linux

For information about these services, see [Chapter 3, “Planning Your Dynamic Storage Technology Server Environment,”](#) on page 29.

- ❑ There must be at least one functional shadow volume on the server that is mounted in NCP. For information, see [Section 10.2, “Creating a DST Shadow Volume with NSS Volumes,”](#) on page 88.

13.3 Preparing Your System for Using ShadowFS

Configure Novell Samba to prepare your system for using ShadowFS. For detailed instructions for installing, configuring, or setting up Novell Samba, see the [OES 2018: Novell Samba Administration Guide](#).

- 1 Verify that Novell Samba services are installed and functioning properly:
 - ◆ Samba server is running.
 - ◆ Shares can be created.
 - ◆ Users can access Samba shares.

Use the Novell Samba plug-in for iManager to configure and verify Samba services. In iManager, go to the **File Protocols > Samba > General** page with the server selected.

- 2 Novell Samba users must be Linux-enabled through Linux User Management in order to access data.

IMPORTANT: You must Linux-enable users before adding a Samba Password policy assignment for the Samba server. If you attempt to add a user to a group, and the user is not already Linux-enabled, you cannot continue.

The users must be members of a primary group that is Linux-enabled on the target server or workstation object where both the Primary Group ID and Primary Group Name are assigned to the user. This is the primary group that is later assigned rights to the Samba share. Only primary groups can be assigned as the Directory group for the Samba share.

Adding users to Samba automatically Linux-enables them with Linux User Management (LUM), and it also Samba enables them. You can also Linux-enable users by using Linux User Management.

To verify Linux-enabled users, go to the **Modify User > Linux Profile > General** page with the server selected. Ensure that the values match the users' Group Assignment.

- 3 Ensure that users have a Samba Password policy assignment at the eDirectory user, group or container level.
- 4 Ensure that users have a Universal Password.

Users must have a Universal Password set in order for Samba to work properly.

- 5 Linux-enable the group with Linux User Management.

You must assign a Unix Workstation object for the group. To verify, use iManager to go to the **Modify Group > Linux Profile > General** page, confirm that the **Enable Linux Profile** option is enabled, and confirm that a **Unix Workstation** object is assigned and has a Group ID.

NOTE: For the purposes of testing, you can PAM-enable services on the server, so that test users can SSH into the server and validate access to directory paths to shares. For information about configuring SSH for a user, see *SSH Services on OES* in the *OES 2018 SP3: Planning and Implementation Guide*.

13.4 Installing ShadowFS and FUSE

ShadowFS and FUSE are installed automatically when you install Dynamic Storage Technology. The following instructions are provided if you need to install it manually.

- 1 Open YaST as the `root` user.
- 2 In YaST, select **Software Management**.
- 3 In **Software Management**, search for *shadow* to find the `novell-shadowFS` package.
- 4 Select **novell-shadowFS**, click **Install**, click **Accept** to install it, then when prompted, accept its dependencies (such as FUSE).
- 5 Start ShadowFS by entering the following at a terminal console as the `root` user:

```
systemctl start novell-shadowfs.service
```

IMPORTANT: Ensure that you run only a single instance of `shadowfs` at a time. Do not enter the command multiple times.

For example, if the primary storage location is an NSS volume named `VOL1` and the secondary storage location is an NSS volume named `ARCVOL`, the output would look similar to this:

```
# systemctl status novell-shadowfs.service
● novell-shadowfs.service - Novell shadowfs
   Loaded: loaded (/usr/lib/systemd/system/novell-shadowfs.service;
disabled; vendor preset: disabled)
   Active: active (running) since Tue 2017-11-28 12:46:29 IST; 1s ago
     Process: 31403 ExecStopPost=/usr/bin/rm -f /var/lock/subsys/shadowfs
(code=exited, status=0/SUCCESS)
     Process: 31401 ExecStopPost=/usr/sbin/rmmod fuse (code=exited,
status=1/FAILURE)
     Process: 31391 ExecStop=/opt/novell/ncpserv/sbin/umountshadow.sh
(code=exited, status=32)
     Process: 2567 ExecStartPost=/usr/bin/touch /var/lock/subsys/shadowfs
(code=exited, status=0/SUCCESS)
     Process: 2553 ExecStart=/opt/novell/ncpserv/sbin/shadowfs
(code=exited, status=0/SUCCESS)
     Process: 2551 ExecStartPre=/sbin/modprobe fuse (code=exited,
status=0/SUCCESS)
    Tasks: 4 (limit: 512)
   CGroup: /system.slice/novell-shadowfs.service
           └─2563 /opt/novell/ncpserv/sbin/shadowfs
             └─2566 /opt/novell/ncpserv/sbin/shadowfs
```

```
Nov 28 12:46:28 blr8-117-174 systemd[1]: Starting Novell shadowfs...
Nov 28 12:46:29 blr8-117-174 shadowfs[2553]: SHIFT_ON_MODIFY: 1
Nov 28 12:46:29 blr8-117-174 shadowfs[2553]: SHIFT_ON_ACCESS: 0
Nov 28 12:46:29 blr8-117-174 shadowfs[2553]:
SHIFT_DAYS_SINCE_LAST_ACCESS: 1
Nov 28 12:46:29 blr8-117-174 shadowfs[2553]: Primary Tree 0: /media/
nss/VOL1
Nov 28 12:46:29 blr8-117-174 shadowfs[2553]: Shadow Tree 0: /media/nss/
ARCVOL
Nov 28 12:46:29 blr8-117-174 shadowfs[2553]: shadowfs root 0: /media/
shadowfs/VOL1
Nov 28 12:46:29 blr8-117-174 systemd[1]: Started Novell shadowfs.
```

Loading ShadowFS creates the ShadowFS root volume `/media/shadowfs/VOL1` where it creates the ShadowFS volumes. If multiple NCP volumes have shadow volumes, each of them is shadowed with ShadowFS and is reported. You cannot control whether to shadow only one or some of them.

13.5 Setting Rights to ShadowFS Shares

Grant POSIX rights for users so they can access files on the ShadowFS volume via the SMB/CIFS protocol. Rights are granted based on need. You set rights so that users can read, write, and execute in the ShadowFS volume's root location in the `/media/shadowfs` directory. Do not set POSIX rights to the actual NCP shares for the primary and secondary volumes.

- 1 Open a terminal console, then log in as the `root` user.
- 2 Go to the ShadowFS volume root location of `/media/shadowfs` by entering the following at the terminal prompt:

```
cd /media/shadowfs
```

- 3 Set directory ownership for the group-level access to the ShadowFS volume root by entering the following:

```
chown :groupname shadowfs_volumename
```

For example, if the *groupname* is `marketing` and the *shadowfs_volumename* is `USERS`, enter

```
chown :marketing USERS
```

- 4 Set POSIX rights for the directory group by entering the following:

```
chmod mode shadowfs_volumename
```

For example, to grant POSIX read, write, and execute permissions for the user and group levels, and to set read and execute only for the others (world) level, set the *mode* to `775` by entering:

```
chmod 775 USERS
```

You are setting directory rights for `/media/shadowfs/USERS` as `drwxrwxr-x`.

- 5 Visually verify POSIX rights by entering

```
ll
```

Continuing the example, the results should look like this:

```
drwxrwxr-x  3 root marketing  80 May 16 15:48 USERS
```

- 6 Verify that the SMB/CIFS user can access the ShadowFS volume and can create directories.
 - 6a Decide which user identity you want to use to test the setup. For example, you could assign the admin user as a user of the SMB/CIFS group, or use iManager to create a temporary user identity for a test user in the group.
 - 6b Use iManager to ensure that the test user is Linux-enabled with Linux User Management, and grant the user SSH rights for accessing the server.

For information about configuring SSH for a user, see [“SSH Services on OES”](#) in the *OES 2018 SP3: Planning and Implementation Guide*.
 - 6c Use iManager to set eDirectory permissions on the volume or path for the test user.
 - 6d Use Secure Shell (SSH) to log in to the volume as a user in the group.

For example, use `ssh` to connect to the server and log in:

```
ssh username@server.context.com
```

```
password:*****
```

- 6e** Go to the ShadowFS volume location by entering

```
cd /media/shadowfs/USERS
```

The user should be able to `cd` to and see the directory. If not, recheck the preceding steps to ensure that you followed the steps correctly.

- 6f** As the user, create a directory. For example, enter

```
mkdir username
```

If the directory `/media/shadowfs/USERS/username` is created, the rights are working as expected.

13.6 Creating a Samba Share

Create a Samba share that points to the newly created ShadowFS root, so that users can access it. Rights do not need to be set at the Primary and Shadow volumes themselves, unless they are not visible or accessible to the user or group assignment.

- 1 Log in to iManager as the administrator user.
- 2 In iManager, click **File Protocols > Samba > Shares**.
- 3 Select a server to manage.
- 4 On the **Shares** page, click **New**.
- 5 Specify the following information:
 - ♦ **Share name:** Specify a share name that does not conflict with existing shares that are defined in the `smb.conf` file. To continue earlier examples in this section, `USERS` has been used, so the Samba share name must differ. For example, `usertest`.
 - ♦ **Path:** Specify the context-sensitive path of the ShadowFS root location for the `USERS` volume, such as `/media/shadowfs/USERS`.
 - ♦ **Comment:** Specify a description of the share, such as “User file storage for Windows users.”
 - ♦ **Inherit ACLs:** Enable this option to allow POSIX inheritance of access control lists and rights.
- 6 Click **Finish** to create the Samba share.

If the share is created successfully, it is listed on the **Shares** page.

13.7 Adding a User to Samba

If Linux-enabled users who need access are not already added to Samba, add them to the Samba server.

- 1 Log in to iManager as the administrator user.
- 2 In iManager, click **File Protocols > Samba > Shares**.
- 3 Select a server to manage, then click the **Users** tab.

- 4 On the **Shares > Users** page, click **Add**, then locate and select the users you want to add to Samba.

If a user is added successfully, the user name is listed on the **Users** page. The user should be listed with the default Samba user group *hostname-W-SambaUserGroup* and with the primary Linux-enabled user group to which the user was added earlier.

Users are automatically added to *hostname-W-SambaUserGroup* when they are added as Samba users via the Samba Management plug-in for iManager. If a user is already a member of another Linux-enabled group, adding the user to Samba adds the Samba group as the user's primary group.

If the user's previous primary group gave the user specific access to PAM-enabled services, the user likely loses those access rights, because the default Samba group gives users no rights to any PAM-enabled services. If this occurs, you can remove the user from the default Samba user group and reassign the user back to his or her previous primary group. This is done by modifying the user's properties.

- 5 If you need to modify a user's properties, go to **User > Modify > Linux Profile**, and change the **Primary Group Name** back to the previous group name. This also changes the **Primary Group ID**.
- 6 If you encounter problems with Samba, you can start, stop, or restart the Samba server from the **File Protocols > Samba > Shares > General** page.

13.8 Connecting Users to the Share

At this point, the Samba share users should be able to attach to server from a Windows client or other CIFS/SMB client. The procedure in this section explains the steps for a Windows XP client. Use a similar method on other Windows operating systems.

- 1 On a Windows 7 computer, open Windows Explorer, then right-click **Computer** or **Network**.
- 2 Select **Map network drive** from the shortcut menu.
- 3 Specify the **Drive** letter for the connection.
- 4 In the **Folder** field, specify the location as `\\servername\Samba_sharename`, for example, `\\svr1\usertest`.

Connecting to the server can take a few seconds to minutes, depending on network speed, discovery of server and share, and so on.

- 5 Select the **Reconnect at logon** and **Connect using different credentials** appropriately, then click **Finish**.
- 6 When prompted, enter your user name (DN only, not FDN) and password.
If the connection is good, an Explorer window opens for the mapped location.
- 7 Ensure that the rights are working by creating a new folder.

13.9 Testing Shadow Volume Policies

If you are not familiar with policies on shadow volumes, you should test them against a test data set to understand how to use them to your advantage.

Add files of several different types to the new share, then either create a DST policy to move the files, or do an inventory to search for specific file types, then move them to the shadow.

SSH in as the user, or root, and look at the primary, shadow, and Shadowfs root paths to see if things are where you expect them to be.

13.10 Enabling or Disabling ShadowFS


By default, ShadowFS and FUSE are not started unless you start them manually. You can set a global policy for **ShadowFS Configuration** that starts them automatically.

IMPORTANT: If you use shadow volumes in a cluster, ensure that you set the same global policies on each OES node in the cluster.

- ♦ [Section 13.10.1, “Loading ShadowFS and FUSE,” on page 152](#)
- ♦ [Section 13.10.2, “Verifying ShadowFS Commands,” on page 152](#)

13.10.1 Loading ShadowFS and FUSE

- 1 In OES Remote Manager for Linux, select **View File System**, then select **Dynamic Storage Technology Options**.
- 2 In the **ShadowFS Configuration** area, view the current setting for **Load ShadowFS**.



This option executes the `systemctl start novell-shadowfs.service` command, and puts the necessary commands in the boot sequence to start ShadowFS on system restart.

- 3 Enable **Load ShadowFS** by selecting the check box.
- 4 In the **ShadowFS Configuration** area, click **Submit** to save and apply the change.

13.10.2 Verifying ShadowFS Commands

In **ShadowFS Configuration** on the Dynamic Storage Technology page, you can enable the **Load ShadowFS** check box to execute the `systemctl start novell-shadowfs.service` command. Enabling this option also puts the commands `shadowfs` and `fuse` startup commands in the boot sequence.

You can verify that the commands are available by viewing the script in a text editor. The following lines should be in the `novell-shadowfs` script

```
ExecStartPre=/sbin/modprobe fuse
ExecStart=/opt/novell/ncpserv/sbin/shadowfs
```

13.11 Starting and Stopping ShadowFS Manually

FUSE (`fuse`) and ShadowFS (`shadowfs`) are required when Novell Samba users are accessing NSS volumes via SMB/CIFS. If FUSE and ShadowFS stop running, you must start them manually. Only one instance of `shadowfs` should be running at a time.

13.11.1 Starting FUSE and ShadowFS

Loading ShadowFS creates a ShadowFS root `/media/shadowfs/<volumename>` directory for each of the mounted DST shadow volumes. The `volumename` is the same as volume name of the primary volume. The ShadowFS root directory contains the merged file tree view of the primary and secondary locations in the DST volume. A root is created for all of the mounted DST volumes; you cannot control whether to shadow only one or some of them.

- 1 On the server, open a terminal console, then log in as the `root` user.
- 2 At the terminal console prompt, start ShadowFS by entering

```
systemctl start novell-shadowfs.service
```

The output identifies the primary volume, secondary volume, and the `shadowfs` volume.

For example, if the primary storage location is an NSS volume named `VOL1` and the secondary storage location is an NSS volume named `ARCVOL`, the output would look similar to this:

```
# systemctl status novell-shadowfs.service
● novell-shadowfs.service - Novell shadowfs
   Loaded: loaded (/usr/lib/systemd/system/novell-shadowfs.service;
 disabled; vendor preset: disabled)
   Active: active (running) since Tue 2017-11-28 12:46:29 IST; 1s ago
     Process: 31403 ExecStopPost=/usr/bin/rm -f /var/lock/subsys/shadowfs
 (code=exited, status=0/SUCCESS)
     Process: 31401 ExecStopPost=/usr/sbin/rmmod fuse (code=exited,
 status=1/FAILURE)
     Process: 31391 ExecStop=/opt/novell/ncpserv/sbin/umountshadow.sh
 (code=exited, status=32)
     Process: 2567 ExecStartPost=/usr/bin/touch /var/lock/subsys/shadowfs
 (code=exited, status=0/SUCCESS)
     Process: 2553 ExecStart=/opt/novell/ncpserv/sbin/shadowfs
 (code=exited, status=0/SUCCESS)
     Process: 2551 ExecStartPre=/sbin/modprobe fuse (code=exited,
 status=0/SUCCESS)
    Tasks: 4 (limit: 512)
   CGroup: /system.slice/novell-shadowfs.service
           └─2563 /opt/novell/ncpserv/sbin/shadowfs
             └─2566 /opt/novell/ncpserv/sbin/shadowfs
```

```
Nov 28 12:46:28 blr8-117-174 systemd[1]: Starting Novell shadowfs...
Nov 28 12:46:29 blr8-117-174 shadowfs[2553]: SHIFT_ON_MODIFY: 1
Nov 28 12:46:29 blr8-117-174 shadowfs[2553]: SHIFT_ON_ACCESS: 0
Nov 28 12:46:29 blr8-117-174 shadowfs[2553]:
SHIFT_DAYS_SINCE_LAST_ACCESS: 1
Nov 28 12:46:29 blr8-117-174 shadowfs[2553]: Primary Tree 0: /media/
nss/VOL1
Nov 28 12:46:29 blr8-117-174 shadowfs[2553]: Shadow Tree 0: /media/nss/
ARCVOL
Nov 28 12:46:29 blr8-117-174 shadowfs[2553]: shadowfs root 0: /media/
shadowfs/VOL1
Nov 28 12:46:29 blr8-117-174 systemd[1]: Started Novell shadowfs.
```

13.11.2 Starting FUSE and ShadowFS with novell-shadowfs

- 1 On the server, open a terminal console, then log in as the `root` user.
- 2 At the terminal console prompt, start FUSE and ShadowFS by entering

```
systemctl start novell-shadowfs.service
```

13.11.3 Stopping Shadowfs

- 1 On the server, open a terminal console, then log in as the `root` user.
- 2 At the terminal console prompt, stop the shadowfs process by entering

```
systemctl stop novell-shadowfs.service
```

If the process does not stop, you need to kill the process. Enter

```
killall shadowfs
```

13.12 Configuring Trustee Rights for Novell Samba Users

When you use ShadowFS to provide a merged view to the Novell Samba users, file access is controlled by the OES trustee model for user access. You must use NCP rights management tools to set trustees, just as you do for NCP clients. For example, you can use the Files and Folders plug-in to iManager, the Client for Open Enterprise Server, or the `ncpcon rights` command to set trustees, trustee rights, and inherited rights filters for files and folders.

14 Generating a File Inventory for DST Shadow Volumes

In OES Remote Manager for Linux, you can view reports for the DST shadow volume, with statistics and information about files in the primary file tree and secondary file tree.

- ♦ [Section 14.1, “Understanding the File Inventory for a Shadow Volume,” on page 155](#)
- ♦ [Section 14.2, “Creating the Shadow Volume Inventory,” on page 166](#)
- ♦ [Section 14.3, “Viewing a Saved NCP Volume Report,” on page 166](#)
- ♦ [Section 14.4, “Viewing Statistics for the Shadow Volume,” on page 166](#)
- ♦ [Section 14.5, “Using Inventory Detail Reports to Move, Copy, or Delete Files on the Shadow Volume,” on page 166](#)
- ♦ [Section 14.6, “Generating a Custom Inventory Report,” on page 167](#)

14.1 Understanding the File Inventory for a Shadow Volume

The inventory provides key statistics about the files in the selected volume, such as files scanned and the available space trends. The inventory includes the following information:

- ♦ [Section 14.1.1, “Inventory Summary,” on page 155](#)
- ♦ [Section 14.1.2, “Available Space Trends,” on page 156](#)
- ♦ [Section 14.1.3, “Graphical Profiles,” on page 156](#)
- ♦ [Section 14.1.4, “Tabular Profiles,” on page 163](#)
- ♦ [Section 14.1.5, “Inventory Detail Reports,” on page 163](#)
- ♦ [Section 14.1.6, “Custom Shadow Volume Options,” on page 164](#)

14.1.1 Inventory Summary

The inventory summary lists the number of files scanned on the primary storage area and the secondary storage area. It also lists key statistics for the primary storage area, the secondary storage area, and both areas combined as the shadow volume.


Key Statistics	Description
Total Subdirectories	The total number of subdirectories in the volume.
Total Files	The total number of files in the volume.
Space in Use	The amount of space currently in use in the volume for data and metadata. On NSS volumes where salvage is enabled, the space in use includes space used by deleted files and directories.

Key Statistics	Description
Space Available	The amount of free space in the volume.
File Types	The number of different file types in use throughout the entire volume.
Soft Link Files	The NSS file system and NCP Server do not support soft links to files. This is a placeholder for future non-NCP support.
Soft Link Subdirectories	The NSS file system and NCP Server do not support soft links to subdirectories. This is a placeholder for future non-NCP support.


The following figure is an example of the summary:

Shadow Volume Inventory

Primary



Shadow



Primary Volume Tree: /media/nss/VOL1
Shadow Volume Tree: /media/nss/ARCVOL1
Report generated on Fri Nov 17 01:19:28 2017
Elapsed Time(seconds): 1

- [Available space trend graph](#)
- [File type profiles](#)
- [File owner profiles](#)
- [Last modified profiles](#)
- [Last accessed profiles](#)
- [Change time profiles](#)
- [File size profiles](#)
- [Links to specific reports](#)
- [Custom Shadow Volume Options](#)

Key Statistics	Totals	Primary Area	Shadow Area
Total Subdirectories:	13	7	6
Total Files:	185	153	32
Space In Use:	629 MB	616 MB	13 MB
Space Available:	-	4,501 MB	5,104 MB
File Types:	24	24	19
Soft Link Files:	0	0	0
Soft Link Subdirectories:	0	0	0

14.1.2 Available Space Trends

The **Available Space Trends** report shows the trends for space usage on the primary storage area and the secondary storage area.

14.1.3 Graphical Profiles

The **Profiles** portion of the inventory report graphically displays information about the shadow volume. Graphical profiles are displayed by size in bytes and file count for the following categories:



- ◆ [“File Type Profiles” on page 157](#)
- ◆ [“File Owner Profiles” on page 158](#)
- ◆ [“Time Stamp Profiles” on page 159](#)
- ◆ [“File Size Profiles” on page 160](#)

File Type Profiles

File Type Profiles indicates storage space usage by file types that are actually in use on your system, such as LOG, TDF, DAT, XML, EXE, and so on.

The following figure is an example of the File Type Profiles graphs:

[File type profiles:](#)
[Data Tables:](#)




File Owner Profiles

File Owner Profiles indicates storage space usage by the designated owner of the file. It is not unusual in NCP to see the `root` user as the owner of files. For NCP volumes and NSS, file access is governed by the file system trustees assigned to the file, not the file owner. Trustees are users who have User objects defined in eDirectory, and who have been granted file system rights for the file. NCP tracks ownership via the user's eDirectory GUID.


The following figure is an example of the **File Owner Profiles** graphs.

[File owner profiles:](#)
[Data Tables:](#)

File Owners (By Bytes In Use)



File Owners (By File Count)



Time Stamp Profiles

Three time stamp profiles are generated:

- ♦ **Files Modified Profiles:** Modified dates indicate the last time someone changed the contents of a file.
- ♦ **Files Accessed Profiles:** Access dates indicate the last time someone accessed a file, but did not change the contents if this differs from the modified date.
- ♦ **Files Changed Profiles:** Change dates indicate the last time someone changed the metadata of a file, but did not change the contents if this differs from the modified date.

Time stamps are grouped by the following time periods:

Within Last Day

1 day to 1 week

1 week to 2 weeks

2 weeks to 1 month

1 month to 2 months

2 months to 4 months

4 months to 6 months

6 months to 1 year


1 year to 2 years

More than 2 years


The following figure is an example of the **File Modified Profiles** graphs. Similar graphs are created for **File Accessed Profiles** and **File Changed Profiles**.

[Last modified profiles:](#)
[Data Tables:](#)

Last Modified Times (By Bytes In Use)



Last Modified Times (By File Count)



File Size Profiles

File Size Profiles reports the size of files, grouped by the following size ranges:


- Less than 1 KB
- 1 KB to 4 KB
- 4 KB to 16 KB
- 16 KB to 64 KB

64 KB to 256 KB
256 KB to 1 MB
1 MB to 4 MB
4 MB to 16 MB
16 MB to 64 MB
64 MB to 256 MB
More than 256 MB


The following figure is an example of the **File Size Profiles** graphs:

[File size profiles:](#)
[Data Tables:](#)

File Size Chart (By Bytes In Use)



File Size Chart (By File Count)



14.1.4 Tabular Profiles

Statistical data used to create the graphs is also available in tables that report statistics for the primary area, the secondary area, and both areas combined as the shadow volume. The count for file entries for the primary area and shadow (secondary) area are linked to detail reports that list the files matching that particular category and group. From the file lists, you have the option to copy, move, or delete one or multiple files.

For example, the following figure shows a few lines of a file-type information table:

File Extension	Total Space In Use	Total File Count	Primary Space	Primary Files	Shadow Space	Shadow Files
bmp	0	6	0	4	0	2
html	316.6 MB	6	309.0 MB	5	7.6 MB	1
jar	123.7 MB	4	123.7 MB	4	0	0
jpg	3.5 MB	12	3.5 MB	12	0	0
pcap	1.6 MB	4	1.6 MB	4	0	0
pptx	264.4 KB	6	237.9 KB	5	26.5 KB	1
xlsx	55.9 KB	6	47.4 KB	5	8.5 KB	1
docx	372.0 KB	34	316.3 KB	30	55.7 KB	4
exe	178.8 MB	13	173.4 MB	11	5.4 MB	2

14.1.5 Inventory Detail Reports

An **Inventory Detail Report** lists all of the files that match a particular category and group for a file count entry in the tabular reports in the shadow volume inventory. You can select one or multiple files in the list, then select one of the following operations to be performed:

- ♦ Move the selected volumes to the other file tree.
- ♦ Move the selected files to a specified path on the server.
- ♦ Copy the selected files to a specified path on the server.
- ♦ Delete the selected files.

The action is performed on the selected files, and a confirmation list is displayed.


Volume Inventory ?

Deleted: /media/nss/VOL1/._NETWARE/trustee_work.dat
Deleted: /media/nss/VOL1/._NETWARE/shadow_policy.bak
Deleted: /media/nss/VOL1/DESKTOP.AFP/ICON/736F646D414E494

Total files deleted: 3

The following figure is an example of a detail report for file types that reside on the secondary volume:

Primary Directories



Inventory Detail Report for: /media/nss/ARCVOL1
 - Shadow Volume Tree
 All files matching selected filter:

Move Checked Files To:

Copy Checked Files To:

 [/media/nss/ARCVOL1/secondary \(1\).html](/media/nss/ARCVOL1/secondary (1).html)
 OWNER: root, Size: 7,995,392 (7.6 MB), Modified: Thu 16 Nov 2017 03:58:13 PM IST, Accessed: Thu 16 Nov 2017 03:58:13 PM IST,
 Changed: Thu 16 Nov 2017 03:58:13 PM IST,

Entries Found: 1

14.1.6 Custom Shadow Volume Options

The **Custom Shadow Volume Options** section of the volume inventory allows you to generate reports based on key statistics of interest, and perform actions on them.

- ◆ [“Volume Operations” on page 164](#)
- ◆ [“Search Patterns” on page 164](#)
- ◆ [“File Owner Restrictions” on page 165](#)
- ◆ [“Time Stamp Restrictions” on page 165](#)
- ◆ [“File Size Restrictions” on page 165](#)

Volume Operations

You can perform one of the following volume operations on the files that match the search criteria you specify:

- ◆ List primary area selected files
- ◆ Move selected files from primary area to shadow area.
- ◆ List shadow area selected files.
- ◆ Move selected files from shadow area to primary area.

Search Patterns

In **Search Patterns**, you can specify wildcards and characters to select files by file names or extensions.

File Owner Restrictions

In **File Owner Restrictions**, select **None** or a user name. The search applies only to files where the file owner matches the specified owner.

Time Stamp Restrictions

You can specify one or multiple time stamps to consider for the search:

- ◆ Last Modified Time
- ◆ Last Accessed Time
- ◆ Last Changed Time

If no time stamp is selected, time stamps are not considered in the search criteria.

If a time stamp is selected, you can specify one or multiple time ranges to consider for the search:

Within last day

1 day to 1 week

1 week to 2 weeks

2 weeks to 1 month

1 month to 2 months

2 months to 4 months

4 months to 6 months

6 months to 1 year

1 year to 2 years

More than 2 years

File Size Restrictions

You can specify one or multiple ranges of file sizes to consider for the search:

Less than 1 KB

1 KB to 4 KB

4 KB to 16 KB

16 KB to 64 KB

64 KB to 256 KB

256 KB to 1 MB

1 MB to 4 MB

4 MB to 16 MB

16 MB to 64 MB

64 MB to 256 MB

More than 256 MB

14.2 Creating the Shadow Volume Inventory

- 1 Open OES Remote Manager for Linux in a web browser, then log in as the `root` user.
- 2 Use one of the following methods to view the volume inventory:
 - ♦ Select **View File System > Dynamic Storage Technology Options**, locate the volume in the list, then click the **Inventory** link next to it.
 - ♦ Select **View File System > NCP Volume Inventory**, locate the volume in the **NCP Volumes Available for Inventory** list, then click the **Volume** link for the volume.

14.3 Viewing a Saved NCP Volume Report

An inventory report is saved when you run an inventory on an NCP volume. You can view the last saved report by going to the **Manage NCP Services > View Inventory Reports** page and clicking the **View Last Report > Display** option for the volume. Graphics are not available in a saved report. The saved report provides the same statistics as running **View File System > NCP Volume Inventory**.

14.4 Viewing Statistics for the Shadow Volume

- 1 In OES Remote Manager, access the volume inventory for the shadow volume.
For information, see [Section 14.2, “Creating the Shadow Volume Inventory,” on page 166](#).
- 2 In the inventory summary area, click a link to go directly to one of the following reports, or scroll to view the reports. For information about each statistical report, see [Section 14.1, “Understanding the File Inventory for a Shadow Volume,” on page 155](#).
 - ♦ Available space trend graph
 - ♦ File type profiles
 - ♦ File owner profiles
 - ♦ Last modified profiles
 - ♦ Last accessed profiles
 - ♦ Change time profiles
 - ♦ File size profiles
 - ♦ Links to specific reports
 - ♦ Custom shadow volume options
- 3 Click the **Data Tables** link for a profile to jump directly to the tabular display of the information that was used to generate the graph.

14.5 Using Inventory Detail Reports to Move, Copy, or Delete Files on the Shadow Volume

- 1 In OES Remote Manager, access the volume inventory for the shadow volume.
For information, see [Section 14.2, “Creating the Shadow Volume Inventory,” on page 166](#).

- 2 In the summary area, click **Links to Specific Reports**, or scroll down to the **Links to Specific Reports** section to view the tabular reports of information used to generate the profiles.
- 3 Review the following categories to locate the files of interest:
 - ◆ Last modified range
 - ◆ Last accessed range
 - ◆ Change time range
 - ◆ File size range
 - ◆ File owner
 - ◆ File extension
- 4 Click the link of the data entry for the files that you want to manage. Files are grouped by Primary area and by shadow (secondary) area.
- 5 In the **Inventory Detail Report**, select one or multiple files in the list, then do one of the following:
 - ◆ Move the selected volumes to the other file tree (primary or shadow (secondary) file tree).
 - ◆ Move the selected files to a specified path on the server.
 - ◆ Copy the selected files to a specified path on the server.
 - ◆ Delete the selected files.

14.6 Generating a Custom Inventory Report

You can customize the inventory report to limit the search sizes and times reported. The reporting criteria can be combinations of the specific categories described in [Section 14.1.6, “Custom Shadow Volume Options,” on page 164](#).

- 1 In OES Remote Manager, access the volume inventory for the shadow volume.
For information, see [Section 14.2, “Creating the Shadow Volume Inventory,” on page 166](#).
- 2 Scroll down to the **Custom Shadow Volume Options** area at the end of the shadow volume inventory.

Custom Shadow Volume Options Volume Operations:

- List primary area selected files.
- Move selected files from primary area to shadow area.
- List shadow area selected files.
- Move selected files from shadow area to primary area.

Search Pattern:

.

File Owner Restriction: None ▾

Time Stamp Restrictions:

Time Stamp:

- Last Modified Time
- Last Accessed Time
- Last Changed Time

Range:

- Within Last Day
- 1 Day - 1 Week
- 1 Week - 2 Weeks
- 2 Weeks - 1 Month
- 1 Month - 2 Months
- 2 Months - 4 Months
- 4 Months - 6 Months
- 6 Months - 1 Year
- 1 Year - 2 Years
- More than 2 Years

File Size Restriction:

- Less than 1KB
- 1 KB - 4 KB
- 4 KB - 16 KB
- 16 KB - 64 KB
- 64 KB - 256 KB
- 256 KB - 1 MB
- 1 MB - 4 MB
- 4 MB - 16 MB
- 16 MB - 64 MB
- 64 MB - 256 MB
- More than 256 MB

Start Scan

- 3 In **Volume Operations**, select one of the following actions to perform on the files that meet the search criteria you specify for the scan in later steps.
 - ◆ List primary area selected files
 - ◆ Move selected files from primary area to shadow area.
 - ◆ List shadow area selected files.
 - ◆ Move selected files from shadow area to primary area.
- 4 In **Search Patterns**, specify wildcards and characters to select files by file name or extension. The default is *.* , which does not restrict the search to specific file names or extensions; all files are considered.
- 5 (Optional) In **File Owner Restrictions**, select **None**, or select a user name from the drop-down list.

If **None** is selected, file ownership is not considered for the search. If a user name is specified, the search applies only to files where the file owner matches the specified owner.

6 (Optional) In **Time Stamp**, specify one or multiple time stamps to be searched. If none are selected, the time stamps are not considered when searching.

- ◆ Last Modified Time
- ◆ Last Accessed Time
- ◆ Last Changed Time

7 In **Range**, if you specified a time stamp restriction, specify one or multiple ranges to be searched.

Within last day
1 day to 1 week
1 week to 2 weeks
2 weeks to 1 month
1 month to 2 months
2 months to 4 months
4 months to 6 months
6 months to 1 year
1 year to 2 years
More than 2 years

8 (Optional) In **File Size Restrictions**, specify one or multiple file sizes to be searched.

Less than 1 KB
1 KB to 4 KB
4 KB to 16 KB
16 KB to 64 KB
64 KB to 256 KB
256 KB to 1 MB
1 MB to 4 MB
4 MB to 16 MB
16 MB to 64 MB
64 MB to 256 MB
More than 256 MB

9 After you specify the volume operation and search criteria, click **Start Scan**.

10 If you chose to list the files, an Inventory Detail Report is generated where you can move, copy, or delete files.

10a Select one or multiple files in the list, then select one of the following actions:

- ◆ **Move the selected volumes to the other file tree.**
- ◆ **Move the selected files to a specified path on the server.**
- ◆ **Copy the selected files to a specified path on the server.**
- ◆ Delete the selected files.

10b Click **OK** to confirm the action.

The action is performed on the selected files, then a confirmation list of the files and the number of files deleted is displayed.

Volume Inventory



Deleted: /media/nss/VOL1/_NETWARE/trustee_work.dat
Deleted: /media/nss/VOL1/_NETWARE/shadow_policy.bak
Deleted: /media/nss/VOL1/DESKTOP/AFP/ICON/736F646D414E494

Total files deleted: 3

If you chose to move selected files from one volume to another, the files that meet the search criteria are automatically moved, then a confirmation list of the files and the number of entries moved is displayed.

Volume Inventory



Custom file move from Primary tree to Shadow tree
All files matching selected filter:

Moved: /media/nss/VOL1/dstuser1111/primary (1).lnk
Moved: /media/nss/VOL1/dstuser1111/secondary (1).lnk
Moved: /media/nss/VOL1/primary (1).lnk
Moved: /media/nss/VOL1/secondary (1).lnk

Entries Moved: 4

- 11 If you view the inventory chart again after the move, you can see that the files that matched the specified criteria before the move are now reported on the other volume.

15 Troubleshooting for DST

This section describes issues and possible workarounds for Dynamic Storage Technology (DST) for Open Enterprise Server (OES).

- ♦ [Section 15.1, “My NCP server information is set to: LOCAL_CODE_PAGE CP437. Why is it not using UTF-8?” on page 171](#)
- ♦ [Section 15.2, “A File is listed twice in a directory,” on page 171](#)
- ♦ [Section 15.3, “Users cannot see some files and directories,” on page 172](#)
- ♦ [Section 15.4, “Cross-protocol locking stops working,” on page 172](#)
- ♦ [Section 15.5, “OES Remote Manager connection error when you are working on the DST Options page,” on page 172](#)
- ♦ [Section 15.6, “Unable to access files on the shadow volume in a DST volume pair,” on page 173](#)
- ♦ [Section 15.7, “An error message “Shadow Volume Path is a non-NSS Volume.” is displayed.,” on page 173](#)

15.1 My NCP server information is set to: LOCAL_CODE_PAGE CP437. Why is it not using UTF-8?

All interaction with the Linux file system uses UTF-8. However, for backward compatibility with older Novell Clients, most of the NCP clients use a server-defined local code page setting. The more recently defined Case 89 NCP clients use UTF-8. We recommend that you configure your client to use them. If all of your clients are using the newer UTF-8 Case 89 NCP clients, then there is no need to set the server's local code page.

15.2 A File is listed twice in a directory

If a file happens to be located in the same directory on both the primary and secondary storage, the file name is listed twice in the directory listing. However, all file operations are directed to the file on the primary system.

To resolve this problem, you can rename one instance of the file to make both versions of the file available under different names. Then open the files to determine which version to keep.

You can control how DST handles duplicate files by configuring global policies. For information, see [Section 7.3, “Resolving Instances of Duplicate Files,” on page 55.](#)

15.3 Users cannot see some files and directories

If the secondary storage location becomes unavailable, it appears to users that some of their files and directories are suddenly missing. When the secondary storage location is back online, the files and directories are visible again.

Users might also observe that some files appear to be missing if NCP Server is having performance issues. Some tuning of NCP Server caching is recommended depending on the server RAM, volume size, number of files, and number of trustees accessing the volumes. For information about tuning issues for NCP Server, see *TID 7004888: NCP Performance Tuning on Open Enterprise Server 2 Linux* (<http://www.novell.com/support/>) in the Micro Focus Knowledgebase.

15.4 Cross-protocol locking stops working

Cross-protocol locking allows Novell Samba users or CIFS users to concurrently access files along with NCP users by allowing only one user at any time to open the file for write. Multiple users who are accessing via these protocols can open a file for read only.

WARNING: Allowing users who access files via different protocols to concurrently open a file for write can lead to data corruption.

NCP Server for Linux provides cross-protocol locking for NCP users, CIFS users, and Linux SMB/CIFS users. OES CIFS supports cross-protocol locking to coordinate with NCP.

If cross-protocol locking is enabled for NCP Server for Linux but stops working for DST shadow volume pairs—that is, multiple users can open a file for read and write—it is probably because ShadowFS needs to be restarted. To resolve this problem, stop the shadowfs process, then start shadowfs. For information, see [Section 13.11, “Starting and Stopping ShadowFS Manually,” on page 152](#).

15.5 OES Remote Manager connection error when you are working on the DST Options page

When you are working on the Dynamic Storage Technology Options page, OES Remote Manager returns the following connection error: The connection to the server was reset while the page was loading.

To resolve this issue, you must restart the OES Remote Manager (`httpstkd`) and Apache (`rcapache2`) daemons.

15.6 Unable to access files on the shadow volume in a DST volume pair

On accessing files on the shadow volume, the "EACCESS (Permission denied)" error is displayed. This is because the ACLs (Inherited Rights Filter (IRF) and Trustees) on the DST shadow volume is not in sync with the primary volume.

To resolve this issue, use the `dst_verify_resync_trustee.py` utility to verify and sync the ACLs from the DST primary volume to secondary volume. For more information on this utility, see [Section A.10, "Verifying and Syncing ACLs \(Inherited Rights Filter \(IRF\) and Trustees\) from DST Primary Volume to Shadow Volume,"](#) on page 239.

15.7 An error message "Shadow Volume Path is a non-NSS Volume." is displayed.

On mounting the primary volume after pool deactivation, directory operations performed on a primary volume fail on the shadow volume. An error message stating shadow volume is a non-NSS volume is displayed in the `ncpserv.log`.

To resolve this issue, mount the shadow volume first and then the primary volume.

16 Security Considerations

This section describes security issues and recommendations for Dynamic Storage Technology (DST) for Open Enterprise Server (OES). It is intended for security administrators or anyone who is using DST and is responsible for the security of the system. It requires a basic understanding of DST. It also requires the organizational authorization and the administrative rights to carry out the configuration recommendations.

- ♦ [Section 16.1, “Client Access,” on page 175](#)
- ♦ [Section 16.2, “Linux-Enabled eDirectory Users,” on page 176](#)
- ♦ [Section 16.3, “Using File System Trustees and Rights,” on page 176](#)
- ♦ [Section 16.4, “Server-to-Server Access,” on page 176](#)
- ♦ [Section 16.5, “Hidden Directories and Files,” on page 177](#)
- ♦ [Section 16.6, “Shadow Volumes Audit Logs,” on page 177](#)
- ♦ [Section 16.7, “Shadow File System Audit Logs,” on page 177](#)
- ♦ [Section 16.8, “NCP Server Auditing and Log Files,” on page 177](#)
- ♦ [Section 16.9, “Using Secure Remote Connections,” on page 178](#)

16.1 Client Access

NCP clients can access a merged view of data on the shadow volume through the normal NCP Server.

OES CIFS clients can access a merged view of data on shadow volumes built with NSS volumes. OES CIFS leverages the NCP Server cache.

All the Active Directory users can now access the data on DST volumes via CIFS.

OES AFP does not support DST.

Linux SMB/CIFS clients can access a merged view of data on a shadow volume through Novell Samba used with ShadowFS and FUSE. These users must be Linux-enabled through Linux User Management.

OES FTP clients can access a merged view of data on a shadow volume through OES FTP used with ShadowFS and FUSE. OES FTP is Pure-FTPd modified to work with eDirectory. These users must be Linux-enabled through Linux User Management.

Other native Linux client protocols such as FTP, HTTP, and NFS are not supported.

16.2 Linux-Enabled eDirectory Users

Dynamic Storage Technology requires that all users of the shadow volume be users that are defined in eDirectory. For information, see the [About this Book and the Library](#).

SMB/CIFS users must be enabled for Linux with Linux User Management. This is true for NCP volumes on Linux POSIX file systems (Ext3 and Reiser) and for NSS volumes on Linux and NetWare. The [OES 2018 SP3: Linux User Management Administration Guide](#) describes how to Linux-enable users for an OES server.

16.3 Using File System Trustees and Rights

Dynamic Storage Technology requires that file system access control for data be managed by using the OES Trustee Model for file system trustees and trustee rights.

For all NCP volumes (NSS and NCP on Linux POSIX volumes), the trustee information is obtained at volume mount time from the `._NETWARE/.trustee_database.xml` file. When trustee changes are made, this trustee database file is updated. Because this file is located on the volume, it follows the volume from node to node as it moves around the cluster.

NCP trustee information is synchronized with the NSS file system. When an NCP user makes a trustee change, the NCP Server informs NSS of the change. When NSS changes a trustee assignment, it generates an event that the NCP Server listens for so NCP can keep up to date on NSS changes. When DST is involved, events from the secondary NSS volume are also noted, and trustee changes are also synchronized with it.

IMPORTANT: For NCP volumes, ensure that the **Inherit POSIX Permissions** option is disabled (the default setting). When this setting is disabled, the local Linux environment access is restricted to the `root` user and the file owner or creator, which is the most secure configuration. For information, see “[Configuring Inherit POSIX Permissions for an NCP Volume](#)” in the [OES 2018 SP3: NCP Server for Linux Administration Guide](#).

Rights and trustee management across multiple file systems should all be managed with the NCP tools. There are rights model mapping problems with using a POSIX rights model on NCP volumes, and vice versa.

All the Active Directory users can now access the data on DST volumes via CIFS. To manage the rights of the Active Directory trustees on the DST volumes, you can use OES File Access Rights Management (NFARM) utility or `rights` utility. For more information, see [Managing the Trustee Rights in the NSS File System](#) and `rights` in the [OES 2018 SP3: NSS File System Administration Guide for Linux](#).

16.4 Server-to-Server Access

iSCSI is the only protocol supported for server-to-server access that allows a remote volume to be used as a primary or secondary storage area for a shadow volume.

16.5 Hidden Directories and Files

- ♦ [Section 16.5.1, “Trustee Database,” on page 177](#)
- ♦ [Section 16.5.2, “Available Space Trends,” on page 177](#)

16.5.1 Trustee Database

A copy of the trustee database is placed in the `._NETWARE` subdirectory in both the primary tree and the shadow tree.

16.5.2 Available Space Trends

An available space trend data file is placed in the `._NETWARE` directory in both the primary tree and the shadow tree. It is used by the volume inventory option in OES Remote Manager for Linux.

16.6 Shadow Volumes Audit Logs

An audit log for a DST shadow volume is located in the `._NETWARE` directory at the root of the primary volume. For NSS volumes, the default file path for the log is `/media/nss/volumename/._NETWARE/volumename.audit.log`. All moves between the primary storage area and the secondary storage area are logged as events to the shadow volume’s audit log.

For example, if the primary area is named `VOL1`, the audit file is `/media/nss/VOL1/._NETWARE/VOL1.audit.log`.

16.7 Shadow File System Audit Logs

Audit logs for the Shadow File System are located in the `/var/opt/novell/log/shadowfs.log` file.

16.8 NCP Server Auditing and Log Files

The following log files are located in the `/var/opt/novell/log` directory:

- ♦ `ncpserv.log`
- ♦ `ncp2nss.log`
- ♦ `ncptop.log`
- ♦ `ncpcon.log`

Log files are managed by `logrotate`. For information on usage, see its man page (`man logrotate`).

The control files for `logrotate` are:

- ♦ `/etc/logrotate.d/novell-ncpserv-log`
- ♦ `/etc/logrotate.d/novell-ncpserv-audit`

- ♦ `/etc/logrotate.d/novell-ncp2nss-log`
- ♦ `/etc/logrotate.d/novell-ncp2nss-audit`

By default, the rollover size is 16 MB and 5 compressed copies are kept.

16.9 Using Secure Remote Connections

If the primary storage area or secondary storage area is connected across remote connections, the connection must be secure. For example, use a virtual private network (VPN) or a private WAN connection.

IMPORTANT: iSCSI is the only protocol supported for remote server-to-server connections.

Ensure that authentication, encryption, and data integrity are secure when accessing and transferring data across the network. For example, if sensitive data is written to the primary volume, that data might be written to the secondary volume, depending on shadow policies in place. If there is an anonymous NFS mount for the shadow volume, the data is transferred in the clear over the network, where it might be prone to attacks or capture. In this case, you want to ensure that only authenticated users are able to access the NFS mount and that the connection between the servers is secure.

17 Configuring DST Shadow Volume Pairs with OES Cluster Services

Dynamic Storage Technology shadow volume pairs can be configured as cluster resources with OES Cluster Services. This section describes two methods for configuring the cluster resource, how to manage the shadow volume in a cluster, and how to remove a shadow relationship from a cluster resource.

- ◆ [Section 17.1, “Planning for DST in a Cluster,” on page 179](#)
- ◆ [Section 17.2, “Planning for DST Shadow Volume Pairs and Policies in a Cluster,” on page 182](#)
- ◆ [Section 17.3, “Preparing the Nodes to Support DST in a Cluster Environment,” on page 187](#)
- ◆ [Section 17.4, “Configuring the DST Pool Cluster Resource with Two Cluster-Enabled Pools,” on page 187](#)
- ◆ [Section 17.5, “Configuring the DST Pool Cluster Resource with a Cluster-Enabled Pool and a Shared Pool,” on page 197](#)
- ◆ [Section 17.6, “Sample Scripts for a DST Pool Cluster Resource,” on page 207](#)
- ◆ [Section 17.7, “Configuring Shadow Volume Policies for the Clustered DST Volume Pair,” on page 209](#)
- ◆ [Section 17.8, “Renaming a Shared Pool in a DST Cluster Resource,” on page 209](#)
- ◆ [Section 17.9, “Renaming a Shared Volume in a DST Cluster Resource,” on page 211](#)
- ◆ [Section 17.10, “Removing the Shadow Relationship for a Clustered DST Volume Pair,” on page 215](#)
- ◆ [Section 17.11, “Upgrading a Cluster with DST Resources from OES 2 SP3 to OES 2015 or Later,” on page 223](#)

17.1 Planning for DST in a Cluster

In addition to the requirements described [Chapter 3, “Planning Your Dynamic Storage Technology Server Environment,” on page 29](#), use the requirements in this section to configure DST on nodes in a OES Cluster Services cluster.

- ◆ [Section 17.1.1, “Open Enterprise Server SP3,” on page 180](#)
- ◆ [Section 17.1.2, “OES Cluster Services,” on page 180](#)
- ◆ [Section 17.1.3, “NCP Server and Dynamic Storage Technology,” on page 180](#)
- ◆ [Section 17.1.4, “OES Storage Services File System,” on page 180](#)
- ◆ [Section 17.1.5, “OES Remote Manager for Linux,” on page 180](#)
- ◆ [Section 17.1.6, “Merged View Access with NCP,” on page 180](#)
- ◆ [Section 17.1.7, “Merged View Access with OES CIFS,” on page 181](#)

- ♦ [Section 17.1.8, “Merged View Access with Novell Samba and ShadowFS,”](#) on page 181
- ♦ [Section 17.1.9, “No Merged View Access for AFP,”](#) on page 182

17.1.1 Open Enterprise Server SP3

Ensure that each node in the cluster is running the same release version of OES.

17.1.2 OES Cluster Services

When you use clustered DST volumes, special steps are needed when you upgrade the cluster from OES 2 SP3 to OES 2015 or later. For information, see [Section 17.11, “Upgrading a Cluster with DST Resources from OES 2 SP3 to OES 2015 or Later,”](#) on page 223.

17.1.3 NCP Server and Dynamic Storage Technology

The NCP (NetWare Core Protocol) Server and the Dynamic Storage Technology software are not cluster aware. They must be installed on every node in the cluster where you plan to migrate or fail over the cluster resource that contains shadow volumes. You do not cluster NCP Server or DST services. You can cluster the DST shadow volume pair by creating a DST pool cluster resource that manages the primary and secondary disks, pools, and volumes.

17.1.4 OES Storage Services File System

Dynamic Storage Technology supports shadow volumes created with pairs of shared OES Storage Services (NSS) volumes. Install NSS on each node in the cluster. For information, see the [OES 2018 SP3: NSS File System Administration Guide for Linux](#).

You must create the two NSS pools and volumes on separate shared disks before you create the shadow volume relationship for the two volumes. The primary pool must be cluster-enabled. The secondary pool must be shared. You can alternatively cluster-enable the secondary pool, but its Cluster objects and IP address are not used by the secondary cluster resource while the two NSS volumes are in the shadow relationship.

17.1.5 OES Remote Manager for Linux

You do not use the Dynamic Storage Technology Options page in OES Remote Manager to create a clustered DST shadow volume pair. The `ncpcon mount` command in the load script creates the DST shadow volume pair on the node where the resource is brought online.

When you use OES Remote Manager for Linux to manage policies for the shadow volume, you typically connect to the IP address of the DST pair cluster resource. You can also connect to the IP address of the server node where the cluster resource is currently mounted.

17.1.6 Merged View Access with NCP

NCP Server allows NCP users to access a merged view of the clustered DST volume pair when the cluster resource is online. As with any clustered volume, the files are not available when the cluster resource is offline.

17.1.7 Merged View Access with OES CIFS

OES CIFS supports the Dynamic Storage Technology merged view of DST volume pairs composed of two NSS volumes. eDirectory and Active Directory users can access a merged view of the clustered DST volume pair when the cluster resource is online. As with any clustered volume, the files are not available when the cluster resource is offline. OES CIFS does not require users to be Linux enabled with Linux User Management (LUM).

You must install and configure OES CIFS on every node in the cluster where you plan to give users CIFS access to the shared cluster resource. For information, see [OES 2018 SP3: OES CIFS for Linux Administration Guide](#).

You want OES CIFS to be available on the node where the DST shadow volumes is active. To do this, you add OES CIFS as an advertising protocol for the primary NSS pool resource as you cluster-enable it. You do not enable OES CIFS shares for the secondary NSS volume, because users access data via the merged view and do not directly access data on the secondary volume.

In the primary NSS pool cluster resource load script, the following command binds OES CIFS to provide access to the shared resource through the virtual server IP address when the resource is mounted on a node.

```
exit_on_error novcifs --add --vserver=virtualserverFDN --ip-addr=virtualserverip
```

In the primary NSS pool cluster resource unload script, the following command unbinds OES CIFS from the DST pool cluster resource when the resource is failed over or cluster migrated to another node in the cluster.

```
ignore_error novcifs --remove --vserver=virtualserverFDN --ip-addr=virtualserverip
```

In the primary NSS pool cluster resource monitor script, the `CIFS_monitor` command helps to keep CIFS up and running.

```
exit_on_error rcnovell-cifs monitor
```

In addition, the following CIFS attributes are automatically added to the NCS:NCP Server object for the virtual server on the primary pool cluster resource:

- ◆ `nfapCIFSServerName` (read access)
- ◆ `nfapCIFSAttach` (read access)
- ◆ `nfapCIFSComment` (read access)
- ◆ `nfapCIFSShares` (write access)

For information, see “[Configuring CIFS with Cluster Services for an NSS File System](#)” in the [OES 2018 SP3: OES CIFS for Linux Administration Guide](#).

17.1.8 Merged View Access with Novell Samba and ShadowFS

ShadowFS and FUSE (File System in Userspace) can be used with Novell Samba to allow SMB/CIFS users to access a merged view of the clustered DST volume pair. Novell Samba is an alternative to OES CIFS; they cannot be used together on the same server. Novell Samba requires users to be Linux enabled with LUM.

You do not enable Novell Samba shares for the secondary NSS volume, because users access data via the merged view and do not directly access data on the secondary volume.

You must install and configure Novell Samba and ShadowFS for each node in the cluster. For information about setting up SMB/CIFS access on each node, see [Chapter 13, “Using ShadowFS to Provide a Merged View for Novell Samba Users,”](#) on page 145.

Additional commands for managing FUSE for the resource must be added manually in the cluster load/unload scripts of the primary pool cluster resource. You must also add the following lines in the load script of the primary NSS pool cluster resource to allow time for ShadowFS to start:

```
# If shadowfs is used, wait for shadowfs to start
for (( c=1; c<=10; c++ )) do
if [ ! -d /media/shadowfs/VOLUME/._NETWARE ]; then sleep 5; fi
done
```

You must add the following line to the unload script of the primary NSS pool cluster resource to unload the volume in FUSE:

```
#unload the volume in FUSE
ignore_error fusermount -u /media/shadowfs/VOLUME
```

17.1.9 No Merged View Access for AFP

OES AFP does not support a merged view of files on the DST volume pair. AFP users see only the files on the primary volume. Do not give AFP users direct access to the secondary volume.

17.2 Planning for DST Shadow Volume Pairs and Policies in a Cluster

In addition to the requirements in [Chapter 9, “Planning for DST Shadow Volume Pairs and Policies,”](#) on page 65, your setup must meet the requirements in this section when you use DST in a OES Cluster Services cluster.

- ◆ [Section 17.2.1, “DST Pool Cluster Resource,”](#) on page 183
- ◆ [Section 17.2.2, “Shadow Volume Definition in the /etc/NCPVolumes File,”](#) on page 183
- ◆ [Section 17.2.3, “Shadow Volume Definition in the ncpserv.conf File,”](#) on page 183
- ◆ [Section 17.2.4, “NCP2NSS Bindings for the Secondary Volume,”](#) on page 184
- ◆ [Section 17.2.5, “NCPCON Mount Command for the Load Script,”](#) on page 184
- ◆ [Section 17.2.6, “Load Order in the Load Script,”](#) on page 184
- ◆ [Section 17.2.7, “Unload Order in the Unload Script,”](#) on page 185
- ◆ [Section 17.2.8, “Monitoring Storage in the Monitor Script,”](#) on page 185
- ◆ [Section 17.2.9, “Additional Volumes in the Primary Pool,”](#) on page 185
- ◆ [Section 17.2.10, “Policies for DST Nodes and Volumes in a Cluster,”](#) on page 186

17.2.1 DST Pool Cluster Resource

The primary and secondary volumes must be able to fail over or cluster migrate together to other nodes in the cluster. Thus, a single DST pool cluster resource is used to manage the pair. Its resource scripts include commands that manage the two devices, pools, volumes.

The devices and pools that contain the primary volume and secondary volume in a clustered DST volume pair must be marked as shareable for clustering. The primary pool must be cluster-enabled for OES Cluster Services. The secondary pool must be shared. You can cluster-enable the pool that contains the secondary volume, but its individual pool resource IP address and Cluster objects are not used in the load and unload scripts for the DST pool cluster resource.

17.2.2 Shadow Volume Definition in the `/etc/NCPVolumes` File

In a cluster, the DST volume pair is defined in the `ncpcon mount` command of the load script for the DST pool cluster resource. When the resource is brought online, the volume is mounted, and an entry is added to the `/etc/NCPVolumes` file. When the resource is taken offline, the volumes are dismounted when the pools are deactivated, and the entry is removed.

```
<VOLUME>
  <NAME>VOL1</NAME>
  <PRIMARY_ROOT>/media/nss/VOL1</PRIMARY_ROOT>
  <SHADOW_ROOT>/media/nss/ARCVOL1</SHADOW_ROOT>
</VOLUME>
```

17.2.3 Shadow Volume Definition in the `ncpserv.conf` File

In a cluster, the DST volume pair is defined with the `ncpcon mount` command in the load script for the DST pool cluster resource. You do not create a clustered DST volume by using the Dynamic Storage Technology Options page in OES Remote Manager. When you bring the resource online on a node for the first time, a `SHADOW_VOLUME` line is automatically added to the `/etc/opt/novell/ncpserv.conf` file:

```
SHADOW_VOLUME primary_volumename secondary_volume_path
```

For example:

```
SHADOW_VOLUME VOL1 /media/nss/ARCVOL1
```

When the resource fails over or is cluster migrated to another node, the shadow volume definition remains defined on that server.

If you remove the shadow relationship from the cluster load script, the `SHADOW_VOLUME` entry is usually not needed in the `/etc/opt/novell/ncpserv.conf` file. To permanently unlink the two volumes, you must manually remove the line from the `/etc/opt/novell/ncpserv.conf` file and restart `ndsd` on each node. To disable clustering but keep the DST shadow volume pair on a specified node, you must manually remove the line from the configuration file and restart `ndsd` on all nodes except that one.

17.2.4 NCP2NSS Bindings for the Secondary Volume

The `EXCLUDE_VOLUME` line in the `/etc/opt/novell/ncp2nss.conf` file prevents the secondary NSS volume from being mounted in NCP. This allows the secondary volume to be mounted for NSS and Linux, but not in NCP. The users access the files on the secondary volume via the merged view of the DST volume pair, not directly.

In a cluster, the DST volume pair is defined with the `ncpcon mount` command in the load script for the DST pool cluster resource. When you bring the resource online on a node for the first time, an `EXCLUDE_VOLUME` line is automatically added to the `/etc/opt/novell/ncp2nss.conf` file, as well as the temporary exclusion table in cache on that node.

```
EXCLUDE_VOLUME secondary_volumename
```

For example:

```
EXCLUDE_VOLUME ARCVOL1
```

If you remove the shadow relationship from the cluster load script, the `EXCLUDE_VOLUME` entry is usually not needed in the `/etc/opt/novell/ncp2nss.conf` file. To permanently unlink the two volumes, you must manually remove the line from the `/etc/opt/novell/ncp2nss.conf` file and restart `ncp2nss` on each node. To disable clustering but keep the DST shadow volume pair on a specified node, you must manually remove the line from the configuration file and restart `ncp2nss` on all nodes except that one.

17.2.5 NCPCON Mount Command for the Load Script

The DST shadow volume pair is defined by the following `ncpcon mount` command in the DST pool cluster resource's load script. The volume pair is available only when the resource is online.

```
exit_on_error ncpcon mount primary_volumename=volID,SHADOWVOLUME=secondary_volumename
```

Both NSS volumes must already exist. The mount location is `/media/nss/<primary_volume_name>`.

Replace `volID` with a volume ID that is unique across all servers in the cluster. Valid values are 0 to 254. By convention, the IDs are assigned from 254 and downward for clustered volumes.

When the primary volume has a state of **Shadowed**, the volume ID that you assign as its NCP volume ID represents the DST shadow volume pair of volumes. The secondary volume does not have a separate volume ID while it is in the shadow relationship.

For example, the following command mounts the primary NSS volume named `VOL1` with a volume ID of 254. The primary volume is mounted for NSS and NCP at `/media/nss/VOL1`. The secondary volume is an existing NSS volume named `ARCVOL1`. It is mounted for NSS at `/media/nss/ARCVOL1`.

```
exit_on_error ncpcon mount VOL1=254,SHADOWVOLUME=ARCVOL1
```

17.2.6 Load Order in the Load Script

The secondary pool must be mounted before the primary pool. This helps to ensure that the pool is activated and available when the DST volume pair is mounted.

IMPORTANT: If the secondary volume is not available when the shadow volume pair is mounted, the cluster load script does not fail and does not provide a warning. The DST shadow volume is created and appears to be working when viewed from OES Remote Manager. However, until the DST shadow volume is mounted, the files on the secondary volume are not available to users and appear to be missing in the merged file tree view. After the secondary volume has successfully mounted, the files automatically appear in the merged file tree view.

If you observe that the pools are slow to mount, you can add a wait time to the load script before the mount command for the shadow volume pair.

For example, you add a `sleep` command with a delay of a few seconds, such as:

```
sleep 10
```

You can increase the sleep time value until it allows sufficient time for the pools to be activated and the volumes to be mounted in NSS before continuing.

IMPORTANT: If wait times are added to the load script or unload script, ensure that you increase the script timeout settings accordingly. Otherwise, the script might time out while you are waiting for the action.

17.2.7 Unload Order in the Unload Script

The primary pool must be deactivated before the secondary pool. This allows the DST volume pair to be dismounted before the secondary pool is deactivated.

17.2.8 Monitoring Storage in the Monitor Script

The monitor script for the DST pool cluster resource has monitoring commands for the primary pool, secondary pool, the primary volume, and advertising protocols for the primary volume.

You should not monitor the secondary volume in the monitor script. The `EXCLUDE_VOLUME` line in the `/etc/opt/novell/ncp2nss.conf` file makes it unavailable to NCP. Thus, the `ncpcon volume` command that is used to check its status is not able to see the secondary volume.

Ensure that you remove or comment out the following line from the resource monitoring script:

```
exit_on_error ncpcon volume primary_volume_name  
#exit_on_error ncpcon volume secondary_volume_name
```

17.2.9 Additional Volumes in the Primary Pool

If you add a volume to the primary pool for a clustered DST volume pair, the `mount` command is added twice in the primary pool's cluster load script, once after the primary pool's activation command and once after the secondary pool's activation command. You must manually delete the instance that occurs after the secondary pool's activation, and then offline and online the primary pool cluster resource to apply the modified load script.

For information, see “[Adding a Volume to a Clustered Pool](#)” in the *OES 2018 SP3: OES Cluster Services for Linux Administration Guide*.

17.2.10 Policies for DST Nodes and Volumes in a Cluster

In a cluster, the DST policies must be available on every node where a clustered DST pool cluster resource is brought online. As a best practice, you should create policies at the volume level for each clustered DST volume pair so that the volume's policies fail over with it when its DST pool cluster resource fails over, or when it is cluster migrated to a different node.

- ♦ [“Global Policies” on page 186](#)
- ♦ [“All-Shadow-Volumes Policies” on page 186](#)
- ♦ [“Volume Policies” on page 187](#)

Global Policies

Global policies are NCP Settings for DST that you set at the server level. They govern how DST behaves for all DST volume pairs mounted on the server. Global policies are not cluster aware.

Ensure that the same global DST policies are configured on each node where you want to fail over the DST pool cluster resources. To manage a global DST policy, open OES Remote Manager for Linux by using the IP address of the node. For information, see [Chapter 7, “Configuring DST Global Policies,” on page 49](#).

IMPORTANT: Whenever you modify global policies on a given node in the cluster, you must make those same changes on the other nodes.

All-Shadow-Volumes Policies

An all-shadow-volumes policy applies to any DST volume that is mounted on that server when the policy runs. The all-shadow-volumes policies are not cluster aware.

If you select **All Shadowed Volumes** when you create a policy, the policy information is stored in the `/usr/novell/sys/._NETWARE/shadow_policy.xml` file. Ensure that the same all-shadow-volumes policies are configured on each node where you want to fail over the DST pool cluster resources. You can create the same all-shadow-volumes policies on each node in the cluster, or you can create it on one node and copy the `shadow_policy.xml` file to all nodes where you plan to bring the DST pool cluster resource online. To manage an all-shadow-volumes policy, open OES Remote Manager for Linux by using the IP address of the node.

IMPORTANT: Whenever you modify all-shadow-volumes policies on a given node in the cluster, you must make those same changes on the other nodes.

Volume Policies

Volume policies apply only to specified DST volume pair. Volume policies are not cluster aware. They are stored with the volume and are available automatically on any node where the DST pool cluster resource is brought online. When you set up volume policies, the DST pool cluster resource must be online and the DST volume pair must be mounted.

If a policy applies to a specific volume, the policy information is stored in the `/media/nss/<primary_volumename>/._NETWARE/shadow_policy.xml` file. This file is stored on the volume itself and thereby automatically follows the volume as its DST pool cluster resource is failed over or cluster migrated to a different node. To manage a volume policy, open OES Remote Manager for Linux by using the IP address of the resource, or by using the IP address of the node where the resource is currently active.

17.3 Preparing the Nodes to Support DST in a Cluster Environment

For each OES 2015 or later server, perform the following tasks to prepare them for hosting DST pool cluster resources in a cluster:

- 1 Install NCP Server and Dynamic Storage Technology. For information, see [Chapter 4, “Installing Dynamic Storage Technology,” on page 35](#).
- 2 Install and configure OES Cluster Services for Linux. For information, see [“Installing, Configuring, and Repairing OES Cluster Services” in the OES 2018 SP3: OES Cluster Services for Linux Administration Guide](#).
- 3 For each node, configure the same DST global policies by using OES Remote Manager. For information, see [Chapter 7, “Configuring DST Global Policies,” on page 49](#).

17.4 Configuring the DST Pool Cluster Resource with Two Cluster-Enabled Pools

One way to set up the DST pool cluster resource is to cluster-enable both pools to create separate cluster resources. You copy the commands from the secondary pool resource scripts to the primary pool resource scripts in the proper load and unload order. The primary pool cluster resource manages the two pools and volumes.

The advantages of creating two cluster pool resources are:

- ♦ You can copy and paste the lines of code you need from one script to the other.
- ♦ The NCP Server object’s name for the secondary volume is automatically renamed to use the cluster name instead of the node hostname. In a migration scenario, you can later remove the shadow relationship and start using the secondary pool immediately as an independent pool cluster resource. Ensure that the volume ID on the volume is unique across all nodes.

The disadvantages of this approach are:

- ♦ The static IP address that is assigned to the secondary cluster pool resource is consumed but not used while the pool is in the shadow relationship.
- ♦ The secondary cluster pool resource appears with a status of **Offline** and is not used.

IMPORTANT: After you modify the primary pool cluster resource, you use this resource to manage the secondary pool and volume. Do not bring the secondary resource online.

Use the information in the following sections to set up the DST pool cluster resource.

- ♦ [Section 17.4.1, “Overview of the Two Pool Cluster Resources,” on page 188](#)
- ♦ [Section 17.4.2, “Viewing the Scripts for the Two Pool Cluster Resources,” on page 189](#)
- ♦ [Section 17.4.3, “Adding Commands for the Secondary Clustered Pool and Volume to the Primary Pool Cluster Resource,” on page 191](#)

17.4.1 Overview of the Two Pool Cluster Resources

For this method, you need two NSS volumes, each in its own clustered-enabled pool. For instructions for creating the clustered pools and the NSS volumes, see [Configuring and Managing Cluster Resources for Shared NSS Pools and Volumes](#) in the *OES 2018 SP3: OES Cluster Services for Linux Administration Guide*.

The cluster load scripts elsewhere in this section assume the following setup for NSS volumes in the clustered DST volume pair. Ensure that you use the actual information from your setup.

Parameter	Primary Cluster Resource	Secondary Cluster Resource
Server hostname for node 1	server38	server38
Cluster server name for node 1	NCS1	NCS1
Pool name	POOL1	ARCPPOOL1
NSS volume name	VOL1	ARCVOL1
Cluster resource virtual server name	NCS1-POOL1-SERVER	NCS1-ARCPPOOL1-SERVER (not used after you set up the primary resource for DST)
Cluster resource IP address	10.10.10.38 You use the IP address for the primary pool’s cluster resource for the shadow volume.	10.10.10.48 (not used after you set up the primary resource for DST)
Volume ID	254	253 (not used after you set up the primary resource for DST)

When the primary volume has a state of **Shadowed**, its NCP volume ID represents the DST shadow volume pair of volumes. A second NCP volume ID is not assigned to the secondary volume while it is in the shadow volume relationship. You use only the ID on the primary volume in the `nccpcon mount` command in the cluster resource load script.

IMPORTANT: In the cluster load and unload scripts, the `add_secondary_ipaddress` and `del_secondary_ipaddress` commands refer to the cluster resource's IP address that is "secondary" to the node's IP address. It is not related to the DST volume's terminology.

17.4.2 Viewing the Scripts for the Two Pool Cluster Resources

After you create two clustered pools, view the scripts. Each of the two NSS pool cluster resources has its own set of load, unload, and monitor scripts. Save the script information for the secondary pool to a text file.

- 1 In iManager, select **Clusters**, then select **My Clusters**.
- 2 Select the name link of the cluster you want to manage.
If the cluster is not in your customized list, you can add it now. Click **Add**, browse to select the cluster, then click **OK**.
- 3 On the Cluster Manager page, click the **Name** link of the primary cluster resource to go to the Cluster Pool Properties page, then click the **Scripts** tab view the load, unload, and monitor scripts.

You can view the scripts for only one server at a time in the browser. View the properties of each resource in separate browsers to compare the scripts side-by-side.

POOL1_SERVER

Policies Monitoring Preferred Nodes **Scripts** Protocols Business Continuity

Load Script | Unload Script | Monitor Script

View or edit the load script for this cluster resource. Changes other than business continuity changes made to a resource will not take affect until the resource is reloaded.

Script:

```
#!/bin/bash
. /opt/novell/ncs/lib/ncsfncs
exit_on_error nss /poolact=POOL1
exit_on_error ncpcon mount VOL1=254
exit_on_error add_secondary_ipaddress 10.10.10.38
exit_on_error ncpcon bind --ncpservname=NCS1-POOL1-SERVER
--ipaddress=10.10.10.38
exit 0
```

Timeout:

OK Cancel Apply

The following table provides sample load, unload, and monitor scripts for the POOL1-SERVER resource for the primary clustered pool named POOL1. OES CIFS can be configured as an advertising protocol when you set up the primary cluster pool.

Primary Pool Cluster Resource Scripts

Load Script

```
#!/bin/bash
. /opt/novell/ncs/lib/ncsfuncs

exit_on_error nss /poolact=POOL1

exit_on_error ncpcon mount VOL1=254

exit_on_error add_secondary_ipaddress 10.10.10.38
exit_on_error ncpcon bind --ncpsservername=NCS1-POOL1-SERVER --ipaddress=10.10.10.38

#This line is added if OES CIFS is used as an advertising protocol
#exit_on_error novcifs --add '--vserver=".cn=NCS1-POOL1-SERVER.ou=ncs.o=novell.t=AVALON_TREE."'
--ip-addr=10.10.10.38

exit 0
```

Unload Script

```
#!/bin/bash
. /opt/novell/ncs/lib/ncsfuncs

#This line is added if OES CIFS is used as an advertising protocol
#ignore_error novcifs --remove --vserver=virtualserverFDN --ip-addr=virtualserverip

ignore_error ncpcon unbind --ncpsservername=NCS1-POOL1-SERVER --ipaddress=10.10.10.38
ignore_error del_secondary_ipaddress 10.10.10.38

ignore_error nss /pooldeact=POOL1

exit 0
```

Monitor Script

```
#!/bin/bash
. /opt/novell/ncs/lib/ncsfuncs

exit_on_error status_fs /dev/pool/POOL1 /opt/novell/nss/mnt/.pools/POOL1 nsspool

exit_on_error status_secondary_ipaddress 10.10.10.38

exit_on_error ncpcon volume VOL1

exit_on_error rcnovell-cifs monitor

exit 0
```

The following are sample load and unload scripts for the ARCPool1-SERVER resource for the secondary clustered pool named ARCPool1.

Secondary Pool Cluster Resource Scripts

Load Script

```
#!/bin/bash
. /opt/novell/ncs/lib/ncsfuns

exit_on_error nss /poolact=ARCPool1

exit_on_error ncpcon mount ARCVOL1=253

exit_on_error add_secondary_ipaddress 10.10.10.48

exit_on_error ncpcon bind --ncpsservername=NCS1-ARCPool1-SERVER --ipaddress=10.10.10.48

exit 0
```

Unload Script

```
#!/bin/bash
. /opt/novell/ncs/lib/ncsfuns

ignore_error ncpcon unbind --ncpsservername=NCS1-ARCPool1-SERVER --ipaddress=10.10.10.48
ignore_error del_secondary_ipaddress 10.10.10.48

ignore_error nss /pooldeact=ARCPool1

exit 0
```

Sample Primary Monitor Script

```
#!/bin/bash
. /opt/novell/ncs/lib/ncsfuns

exit_on_error status_fs /dev/pool/ARCPool1 /opt/novell/nss/mnt/.pools/ARCPool1 nsspool

exit_on_error status_secondary_ipaddress 10.10.10.48

exit_on_error ncpcon volume ARCVOL1

exit 0
```


- 4 Copy information from the secondary resource's scripts into a text file, and save the file. You will work from this copy to add lines to the primary pool cluster resource.
- 5 At the bottom of the Scripts page, click **Cancel** to return to the Cluster Manager page.
- 6 Continue with [Section 17.4.3, "Adding Commands for the Secondary Clustered Pool and Volume to the Primary Pool Cluster Resource,"](#) on page 191.

17.4.3 Adding Commands for the Secondary Clustered Pool and Volume to the Primary Pool Cluster Resource

The clustered DST shadow volume is defined and managed in the primary pool cluster resource. You must add lines from the secondary pool cluster resource scripts and modify the mount command to define the DST shadow volume.

- 1 In iManager, select **Clusters**, then select **My Clusters**.
- 2 Select the name link of the cluster you want to manage.

- 3 Offline the primary cluster resource. The secondary cluster resource should still be offline.
 - 3a On the Cluster Manager page, select the check box next to the resource.
 - 3b Click **Offline**.
- 4 Click the name link of the primary pool cluster resource to view its Cluster Pool Properties page, then click the **Scripts** tab.

 POOL1_SERVER

Policies Monitoring Preferred Nodes **Scripts** Protocols Business Continuity

Load Script | Unload Script | Monitor Script

View or edit the load script for this cluster resource. Changes other than business continuity changes made to a resource will not take affect until the resource is reloaded.

Script:

```
#!/bin/bash
. /opt/novell/ncs/lib/ncsfncs
exit_on_error nss /poolact=POOL1
exit_on_error ncpcon mount VOL1=254
exit_on_error add_secondary_ipaddress 10.10.10.38
exit_on_error ncpcon bind --ncpservername=NCS1-POOL1-SERVER
--ipaddress=10.10.10.38
exit 0
```

Timeout:

OK Cancel Apply

- 5 On the **Scripts > Load Script** page, modify the load script for the primary cluster resource. Use the following sample load script as a guide for where to add the lines for each of the items.

Sample DST Pool Resource Load Script

```
#!/bin/bash
. /opt/novell/ncs/lib/ncsfuns

# activate the secondary pool
exit_on_error nss /poolact=ARCPPOOL1

# activate the primary pool
exit_on_error nss /poolact=POOL1

# Optional delay to allow time for pools to activate before mounting the volume
sleep 10

#comment out the original volume mount command
#exit_on_error ncpcon mount VOL1=254

# Use the ncpcon mount command to create the shadow volume on mount
exit_on_error ncpcon mount VOL1=254,shadowvolume=ARCVOL1

exit_on_error add_secondary_ipaddress 10.10.10.38

exit_on_error ncpcon bind --ncpsservername=NCS1-POOL1-SERVER --ipaddress=10.10.10.38

#This line is added if OES CIFS is used as an advertising protocol
#exit_on_error novcifs --add '--vserver=".cn=NCS1-POOL1-SERVER.ou=ncs.o=novell.t=AVALON_TREE."'
#--ip-addr=10.10.10.38

# If shadowfs is used, wait for shadowfs to start
#for (( c=1; c<=10; c++ )) do
# if [ ! -d /media/shadowfs/VOLUME/.NETWARE ]; then sleep 5; fi
#done

exit 0
```

5a Add a line to activate the secondary pool before the primary pool activation.

```
exit_on_error nss /poolact=ARCPPOOL1
```

5b (Optional) Add a `sleep` command after the pool activation commands to allow both pools time to be activated before you mount the shadow volume pair.

For example:

```
sleep 10
```

Vary the time (in seconds) according to what is needed for your system.

IMPORTANT: If wait times are added to the load script or unload script, ensure that you increase the script timeout settings accordingly. Otherwise, the script might time out while you are waiting for the action.

5c Comment out (or remove) the individual mount command for the primary NSS volume by placing a pound sign (#) at the beginning of the line.

For example:

```
#exit_on_error ncpcon mount VOL1=254
```

5d Add the shadow volume mount command to the primary load script. This line provides the primary volume, and assigns the secondary volume to shadow the primary.

```
exit_on_error ncpcon mount VOL1=254,shadowvolume=ARCVOL1
```

5e If you are using `shadowfs` to provide the merged file tree view for SMB/CIFS users or for Linux services like `rsync`, you must allow time in the load script after mounting the shadow volume to allow `shadowfs` to become active before continuing.

Use one of the following approaches to add a wait time:

- ♦ Add a `sleep 10` command after mount command, and vary it manually until it allows sufficient wait time for shadowfs to start.

```
# If shadowfs is used, wait for shadowfs to start
sleep 10
```

- ♦ Add a script that varies the wait time by checking to ensure that shadowfs is started.

For example:

```
# If shadowfs is used, wait for shadowfs to start
for (( c=1; c<=10; c++ )) do
  if [ ! -d /media/shadowfs/VOLUME/._NETWARE ]; then sleep 5; fi
done
```

IMPORTANT: If wait times are added to the load script or unload script, ensure that you increase the script timeout settings accordingly. Otherwise, the script might time out while you are waiting for the action.

- 5f** Click **Apply** to save your changes.

The changes do not take effect until the shadow volume cluster resource is brought online.

- 6** On the **Scripts > Unload Script** page, modify the unload script for the primary cluster resource. Use the following sample unload script as a guide for where to add the lines for each of the items.

Sample DST Pool Resource Unload Script

```
#!/bin/bash
. /opt/novell/ncs/lib/ncsfncs

# This line is added if OES CIFS is used as an advertising protocol
#ignore_error novcifs --remove '--vserver=".cn=NCS1-POOL1-
SERVER.ou=ncs.o=novell.t=AVALON_TREE."' --ip-addr=10.10.10.38

# If shadowfs is used, unload the volume in FUSE
#ignore_error fusermount -u /media/shadowfs/VOL1

ignore_error ncpcon unbind --ncpservname=NCS1-POOL1-SERVER --ipaddress=10.10.10.38
ignore_error del_secondary_ipaddress 10.10.10.38

# Deactivate the primary pool
ignore_error nss /pooldeact=POOL1

# Deactivate the secondary pool
ignore_error nss /pooldeact=ARCPPOOL1

exit 0
```

-
- 6a** If you use shadowfs to provide a merged file tree view to Samba users or for Linux file protocols, you must unmount the FUSE-mounted file systems that are displayed in the `/media/shadowfs/VOLUME` directory. Add the following line just before the unbind command in the unload script:

```
#unload the volume in FUSE
# Include the following line only if shadowfs is used
ignore_error fusermount -u /media/shadowfs/VOLUME
```

- 6b** Copy the pool deactivation command from the secondary pool's unload script into the primary pool's unload script after the line to deactivate the primary pool.

```
ignore_error nss /pooldeact=ARCPool1
```

IMPORTANT: Ensure that you deactivate the primary pool before deactivating the secondary pool.

- 6c** Click **Apply** to save your changes.

The changes do not take effect until the shadow volume cluster resource is brought online.

- 7** On the **Scripts > Monitor Script** page, modify the monitor script for the primary cluster resource. Use the following sample monitor script as a guide for where to add the lines for each of the items.

Sample DST Pool Resource Monitor Script

```
#!/bin/bash
. /opt/novell/ncs/lib/ncsfuncs

# Check the status of the secondary pool
exit_on_error status_fs /dev/pool/ARCPool1 /opt/novell/nss/mnt/.pools/ARCPool1 nsspool

# Check the status of the primary pool
exit_on_error status_fs /dev/pool/Pool1 /opt/novell/nss/mnt/.pools/Pool1 nsspool

exit_on_error status_secondary_ipaddress 10.10.10.38

# Check the status of the primary volume. Do not check secondary volume.
exit_on_error ncpcn volume VOL1

# This line is added if OES CIFS is used as an advertising protocol
#exit_on_error rcnovell-cifs monitor

exit 0
```

- 7a** Copy the pool status check command from the secondary pool's monitor script into the primary pool's monitor script before the line to check the status of the primary pool.

- 7b** Do not add a check for the secondary volume.

- 7c** Click **Apply** to save your changes.





The changes do not take effect until the shadow volume cluster resource is brought online.

- 8** Click **OK** to save all your changes and return to the Cluster Manager page.
- 9** Online the primary pool cluster resource. On the Cluster Manager page, select the check box next to the primary cluster resource, then click **Online**.

Leave the secondary resource offline.

- 10** Verify that the primary cluster resource is running by going to the **Cluster Manager** page.

The primary cluster resource is **Running**. The secondary cluster resource is reported as **Offline** because you are managing that cluster resource through the primary load script.

<input type="checkbox"/>		ARCPool1_SERVER		Offline	1	
<input type="checkbox"/>		Pool1_SERVER		Running	server38	1 Nov 15, 2017 11:51:50 AM

11 Verify that the shadow volume (VOL1) is mounted in NCP and is shadowed:

11a On the first node in the cluster, log in to OES Remote Manager for Linux as the `root` user.

11b Select **View File System > Dynamic Storage Technology Options**, then verify that the primary volume is listed under **Volume Information**, and that its status is **Shadowed**.

Dynamic Storage Technology Options ?

Dynamic Storage Technology allows you optimize the use of your storage by automatically moving data to storage best optimized for the data type or frequency of use. You can create one or more policies to manage, and optimize the use of your storage.

Volume Information			
Volume Name	Shadow	Status	
① SHARE	Add Shadow	Inventory	
① VOL2	Add Shadow	Inventory	
① DATA	Add Shadow	Inventory	
① VOL1	Shadowed	Inventory	View Log
① SYS	Add Shadow	Inventory	

Dynamic Storage Technology Policies			
Name	Volume	Last Executed	Total Files Moved
ProjectABC Exclude contracts	VOL1	15:58, 11/16/2017	27

[Create a new policy](#)

[Stop all running policies](#)

Duplicate File Resolution Options	
Broadcast conflict message to user:	<input checked="" type="checkbox"/>
Action to be taken:	<input type="text" value="Show duplicate shadow files"/>
Submit	

ShadowFS Configuration
<input type="checkbox"/> Load ShadowFS (Enable only to give LUM-enabled users merged-view access to DST volumes via Samba or Linux file access protocols)
Submit

11c Select **Manage NCP Services > Manage Shares**, click **NCP/NSS Bindings**, then verify that the **NCP Accessible** parameter is turned off for the secondary volume, and turned on for the primary volume.

NCP / NSS Bindings

Warning:

When a NSS Volume is changed to be not accessible via NCP, it will be dismounted immediately as a NCP share point

Available NSS volumes		
NCP Accessible	Volume Name	Mount point
Yes: <input type="radio"/> No: <input checked="" type="radio"/> Save Selection	ARCVOL1	/media/nss/ARCVOL1
Yes: <input checked="" type="radio"/> No: <input type="radio"/> Save Selection	DATA	/media/nss/DATA

[Share Management Home](#)

You can also look for the `EXCLUDE_VOLUME` entry in the `/etc/opt/novell/ncp2nss.conf` file.

11d Continue with [Section 17.7, “Configuring Shadow Volume Policies for the Clustered DST Volume Pair,”](#) on page 209.

17.5 Configuring the DST Pool Cluster Resource with a Cluster-Enabled Pool and a Shared Pool

An alternate way to set up the DST pool cluster resource is to cluster-enable the primary pool, then create a shared pool that is not cluster-enabled as the secondary location. You manually modify the primary cluster pool resource scripts with the commands needed to load and unload the secondary pool and volume. The primary pool cluster resource manages the two pools and volumes. The NCP Server object's name for the secondary volume keeps the hostname of the node where it was created.

The advantage of creating one clustered pool and one shared-but-not-cluster-enabled pool are:

- ◆ A static IP address is not consumed for the secondary pool.
- ◆ The secondary pool and volume are shared but not clustered. This can be useful in a migration scenario to move data to a secondary volume and move the volume to a different node. After you remove the shadow relationship, you mark the secondary pool's device as **Not Shareable for Clustering** to unshare the pool and volume, then use the Update eDirectory option in NSSMU to create storage objects with the new hostname. There are no Cluster objects to clean up.

There disadvantages of this approach are:

- ◆ You must manually enter the lines of code in the primary pool cluster resource's load, unload, and monitor scripts for the secondary pool.
- ◆ The option to create a pool on a shared device without cluster-enabling it is available only in NSSMU. You cannot use the Storage plug-in in iManager.
- ◆ To later use the secondary pool as an independent cluster resource requires a some extra steps. Use the Update eDirectory option in NSSMU to create storage objects with the new hostname, then cluster-enable the pool by using the Clusters plug-in in iManager.

IMPORTANT: After you modify the primary pool cluster resource, you use the resource to manage the secondary pool and volume.

Use the information in the following sections to set up the DST pool cluster resource.

- ◆ [Section 17.5.1, "Overview of the Pool Cluster Resource and Shared Pool," on page 197](#)
- ◆ [Section 17.5.2, "Viewing the Scripts for Pool Cluster Resource," on page 198](#)
- ◆ [Section 17.5.3, "Creating a Shared Pool and Volume that Are Not Cluster-Enabled," on page 200](#)
- ◆ [Section 17.5.4, "Adding Commands for the Secondary Shared Pool and Volume to the Primary Pool Cluster Resource," on page 201](#)

17.5.1 Overview of the Pool Cluster Resource and Shared Pool

For this method, you need two NSS volumes: one in a clustered-enabled pool and one in a pool that is shared, but not cluster-enabled. For instructions for creating the cluster-enabled pool and the primary NSS volume, see [Configuring and Managing Cluster Resources for Shared NSS Pools and Volumes](#) in the *OES 2018 SP3: OES Cluster Services for Linux Administration Guide*.

To create the secondary shared pool, see [Section 17.5.3, "Creating a Shared Pool and Volume that Are Not Cluster-Enabled," on page 200](#).

The cluster load scripts elsewhere in this section assume the following setup for NSS volumes in the clustered DST volume pair. Ensure that you use the actual information from your setup.

Parameter	Primary Cluster Resource	Secondary Shared Pool
Server hostname for node 1	server38	server38
Cluster server name for node 1	NCS1	NCS1
Pool name	POOL1	ARCPool1
NSS volume name	VOL1	ARCVOL1
Cluster resource virtual server name	NCS1-POOL1-SERVER	None
Cluster resource IP address	10.10.10.38	None
	You use the IP address for the primary pool's cluster resource for the shadow volume.	
Volume ID	254	Automatically assigned when you create the volume locally

When the primary volume has a state of **Shadowed**, its NCP volume ID represents the DST shadow volume pair of volumes. A separate NCP volume ID is not assigned to the secondary volume while the volume is in the shadow volume relationship. You use only the ID on the primary volume in the `ncpcon mount` command in the cluster resource load script.

IMPORTANT: In the cluster load and unload scripts, the `add_secondary_ipaddress` and `del_secondary_ipaddress` commands refer to the cluster resource's IP address that is "secondary" to the node's actual IP address. It is not related to the DST volume's terminology.

17.5.2 Viewing the Scripts for Pool Cluster Resource

After you create a clustered pool, view the scripts.

- 1 In iManager, select **Clusters**, then select **My Clusters**.
- 2 Select the name link of the cluster you want to manage.
If the cluster is not in your customized list, you can add it now. Click **Add**, browse to select the cluster, then click **OK**.
- 3 On the Cluster Manager page, click the **Name** link of the primary cluster resource to go to the Cluster Pool Properties page, then click the **Scripts** tab view the load, unload, and monitor scripts.

View or edit the load script for this cluster resource. Changes other than business continuity changes made to a resource will not take affect until the resource is reloaded.

Script:

```
#!/bin/bash
. /opt/novell/ncs/lib/ncsfuncs
exit_on_error nss /poolact=POOL1
exit_on_error ncpcon mount VOL1=254
exit_on_error add_secondary_ipaddress 10.10.10.38
exit_on_error ncpcon bind --ncpservername=NCS1-POOL1-SERVER
--ipaddress=10.10.10.38
exit 0
```

Timeout:

OK Cancel Apply

The following table provides sample load, unload, and monitor scripts for the POOL1-SERVER resource for the primary clustered pool named POOL1. OES CIFS can be configured as an advertising protocol when you set up the primary cluster pool.

Primary Pool Cluster Resource Scripts

Load Script

```
#!/bin/bash
. /opt/novell/ncs/lib/ncsfuncs

exit_on_error nss /poolact=POOL1

exit_on_error ncpcon mount VOL1=254

exit_on_error add_secondary_ipaddress 10.10.10.38
exit_on_error ncpcon bind --ncpservername=NCS1-POOL1-SERVER --ipaddress=10.10.10.38

#This line is added if OES CIFS is used as an advertising protocol
#exit_on_error novcifs --add '--vserver=".cn=NCS1-POOL1-SERVER.ou=ncs.o=novell.t=AVALON_TREE."'
--ip-addr=10.10.10.38

exit 0
```

Unload Script

```
#!/bin/bash
. /opt/novell/ncs/lib/ncsfuncs

#This line is added if OES CIFS is used as an advertising protocol
#ignore_error novcifs --remove --vserver=virtualserverFDN --ip-addr=virtualserverip
i
ignore_error ncpcon unbind --ncpservername=NCS1-POOL1-SERVER --ipaddress=10.10.10.38
ignore_error del_secondary_ipaddress 10.10.10.38

ignore_error nss /pooldeact=POOL1

exit 0
```

Primary Pool Cluster Resource Scripts

Monitor Script

```
#!/bin/bash
. /opt/novell/ncs/lib/ncsfuns

exit_on_error status_fs /dev/pool/POOL1 /opt/novell/nss/mnt/.pools/POOL1 nsspool
exit_on_error status_secondary_ipaddress 10.10.10.38

exit_on_error ncpccon volume VOL1

exit_on_error rcnovell-cifs monitor

exit 0
```

- 4 At the bottom of the Scripts page, click **Cancel** to return to the Cluster Manager page.
- 5 Continue with [Section 17.5.3, “Creating a Shared Pool and Volume that Are Not Cluster-Enabled,”](#) on page 200.

17.5.3 Creating a Shared Pool and Volume that Are Not Cluster-Enabled

- 1 On the server where the primary pool cluster resource is assigned, log in as the `root` user.
- 2 Open a terminal console, and enter `nssmu` to open the NSS Management Utility.
- 3 From the NSSMU main menu, select **Devices** to go to the Device Management page.
- 4 Ensure that the device you want to use as the secondary location is available but is not currently shared.

Do not mark the device as shareable at this time. If devices are present but not showing up for creating pools and volumes, you should initialize the disk.
- 5 From the NSSMU main menu, select **Pools** to go to the Pools Management page.
- 6 Create a pool on the device.

Because the device is not yet shared, the **Cluster Information** page is not part of the pool setup process.
- 7 From the NSSMU main menu, select **Volumes** to go to the Volumes Management page.
- 8 Create a volume on the pool.
- 9 Exit NSSMU.
- 10 In a web browser, open OES Remote Manager for the server, then log in as the `root` user.
- 11 In OES Remote Manager, disable the NCP/NSS Bindings for the NSS volume you created in [Step 8](#).

For instructions, see [Section 10.5.1, “Disabling the NCP/NSS Bindings for an NSS Volume,”](#) on page 98.

The NSS volume is removed from the list of volumes mounted in NCP. However, if you exit OES Remote Manager and check the volume in NSSMU, you can see that it is still mounted by NSS.
- 12 Exit OES Remote Manager.
- 13 At a command prompt, launch NSSMU.

```
nssmu
```
- 14 In the NSSMU main menu, select **Devices**.

- 15 Select the device that contains the secondary NSS pool and volume, then press F6 to mark the device as **Shareable for Clustering**.
This automatically changes the share status of all pools on the device to **Shareable for Clustering**.
- 16 In NSSMU, select **Pools** from the main menu.
- 17 Verify that the share status of the pool is **Shareable for Clustering**.
- 18 Exit NSSMU.
- 19 Continue with [Section 17.5.4, “Adding Commands for the Secondary Shared Pool and Volume to the Primary Pool Cluster Resource,”](#) on page 201.

17.5.4 Adding Commands for the Secondary Shared Pool and Volume to the Primary Pool Cluster Resource

Initially, you have a load, unload, and monitor script for the primary pool cluster resource. You modify these scripts to also manage the secondary shared pool and volume so that the NSS volumes in the DST shadow volume pair can fail over or be cluster migrated together.

- 1 In iManager, dismount the secondary volume and deactivate the shared pool.
 - 1a Select **Storage > Volumes**.
 - 1b Click the **Object** browser, then locate and select the server where the secondary pool is active.
 - 1c Select the secondary volume, then click **Dismount**.
 - 1d Select **Storage > Pools**.
 - 1e Select the secondary pool, then click **Deactivate**.
- 2 In iManager, select **Clusters**, then select **My Clusters**.
- 3 Select the name link of the cluster you want to manage.
- 4 Offline the primary cluster resource.
 - 4a On the Cluster Manager page, select the check box next to the resource.
 - 4b Click **Offline**.
- 5 Click the name link of the primary pool cluster resource to view its Cluster Pool Properties page, then click the **Scripts** tab.

Policies Monitoring Preferred Nodes **Scripts** Protocols Business Continuity

Load Script | Unload Script | Monitor Script

View or edit the load script for this cluster resource. Changes other than business continuity changes made to a resource will not take affect until the resource is reloaded.

Script:

```
#!/bin/bash
. /opt/novell/ncs/lib/ncsfuns
exit_on_error nss /poolact=POOL1
exit_on_error ncpcon mount VOL1=254
exit_on_error add_secondary_ipaddress 10.10.10.38
exit_on_error ncpcon bind --ncpservername=NCS1-POOL1-SERVER
--ipaddress=10.10.10.38
exit 0
```

Timeout:

OK Cancel Apply

- On the **Scripts > Load Script** page, modify the load script for the primary cluster resource. Use the following sample load script as a guide for where to add the lines for each of the items.

Sample DST Pool Resource Load Script

```
#!/bin/bash
. /opt/novell/ncs/lib/ncsfuns

# activate the secondary pool
exit_on_error nss /poolact=ARCPPOOL1

# activate the primary pool
exit_on_error nss /poolact=POOL1

# Optional delay to allow time for pools to activate before mounting the volume
sleep 10

#comment out the original volume mount command
#exit_on_error ncpcon mount VOL1=254

# Use the ncpcon mount command to create the shadow volume on mount
exit_on_error ncpcon mount VOL1=254,shadowvolume=ARCVOL1

exit_on_error add_secondary_ipaddress 10.10.10.38

exit_on_error ncpcon bind --ncpservername=NCS1-POOL1-SERVER --ipaddress=10.10.10.38

#This line is added if OES CIFS is used as an advertising protocol
#exit_on_error novcifs --add '--vserver=".cn=NCS1-POOL1-SERVER.ou=ncs.o=novell.t=AVALON_TREE."'
--ip-addr=10.10.10.38

# If shadowfs is used, wait for shadowfs to start
#for (( c=1; c<=10; c++ )) do
# if [ ! -d /media/shadowfs/VOLUME/.NETWARE ]; then sleep 5; fi
#done

exit 0
```

- Add a line before the primary pool activation that will activate the secondary pool.

```
exit_on_error nss /poolact=ARCPool1
```

- 6b** (Optional) Add a `sleep` command after the pool activation commands to allow both pools time to be activated before you mount the shadow volume pair.

For example:

```
sleep 10
```

Vary the time (in seconds) according to what is needed for your system.

IMPORTANT: If wait times are added to the load script or unload script, ensure that you increase the script timeout settings accordingly. Otherwise, the script might time out while you are waiting for the action.

- 6c** Comment out (or remove) the individual mount command for the primary NSS volume by placing a pound sign (#) at the beginning of the line.

For example:

```
#exit_on_error ncpcon mount VOL1=254
```

- 6d** Add the mount command for the shadow volume to the primary load script.

```
exit_on_error ncpcon mount VOL1=254,shadowvolume=ARCVOL1
```

- 6e** (Optional) If you are using `shadowfs` to provide the merged file tree view for SMB/CIFS users or for Linux services like `rsync`, you must allow time in the load script after mounting the shadow volume to allow `shadowfs` to become active before continuing.

IMPORTANT: Do not perform this step if you are using OES CIFS to provide access to CIFS users.

Use one of the following approaches to add a wait time:

- ◆ Add a `sleep 10` command after mount command, and vary it manually until it allows sufficient wait time for `shadowfs` to start.

```
# If shadowfs is used, wait for shadowfs to start
sleep 10
```

- ◆ Add a script that varies the wait time by checking to ensure that `shadowfs` is started.

For example:

```
# If shadowfs is used, wait for shadowfs to start
for (( c=1; c<=10; c++ )) do
  if [ ! -d /media/shadowfs/VOLUME/._NETWARE ]; then sleep 5; fi
done
```

IMPORTANT: If wait times are added to the load script or unload script, ensure that you increase the script timeout settings accordingly. Otherwise, the script might time out while you are waiting for the action.

- 6f** Click **Apply** to save your changes.

The changes to the script do not take effect until the cluster resource is taken offline and brought online.

- 7** On the **Scripts > Unload Script** page, modify the unload script for the primary cluster resource. Use the following sample unload script as a guide for where to add the lines for each of the items.

Sample DST Pool Resource Unload Script

```
#!/bin/bash
. /opt/novell/ncs/lib/ncsfuncs

# This line is added if OES CIFS is used as an advertising protocol
#ignore_error novcifs --remove '--vserver=".cn=NCS1-POOL1-
SERVER.Ou=ncs.o=novell.t=AVALON_TREE."' --ip-addr=10.10.10.38

# If shadowfs is used, unload the volume in FUSE
#ignore_error fusermount -u /media/shadowfs/VOL1

ignore_error ncpcon unbind --ncpsservername=NCS1-POOL1-SERVER --ipaddress=10.10.10.38

ignore_error del_secondary_ipaddress 10.10.10.38

# Deactivate the primary pool
ignore_error nss /pooldeact=POOL1

# Deactivate the secondary pool
ignore_error nss /pooldeact=ARCPPOOL1

exit 0
```

- 7a** If you use `shadowfs` to provide a merged file tree view to Samba users or for Linux file protocols, you must unmount the FUSE-mounted file systems that are displayed in the `/media/shadowfs/VOLUME` directory.

Add the following line before the `unbind` command in the unload script:

```
#unload the volume in FUSE
ignore_error fusermount -u /media/shadowfs/VOLUME
```

- 7b** Add the following command to dismount the secondary pool after the command to dismount the primary pool.

```
ignore_error nss /pooldeact=ARCPPOOL1
```

- 7c** Click **Apply** to save your changes.

The changes to the script do not take effect until the cluster resource is taken offline and brought online.

- 8 On the **Scripts > Monitor Script** page, modify the monitor script for the primary cluster resource. Use the following sample monitor script as a guide for where to add the lines for each of the items.

Sample DST Pool Resource Monitor Script

```
#!/bin/bash
. /opt/novell/ncs/lib/ncsfuns

# Check the status of the secondary pool
exit_on_error status_fs /dev/pool/ARCPool1 /opt/novell/nss/mnt/.pools/ARCPool1 nsspool

# Check the status of the primary pool
exit_on_error status_fs /dev/pool/POOL1 /opt/novell/nss/mnt/.pools/POOL1 nsspool

exit_on_error status_secondary_ipaddress 10.10.10.38

# Check the status of the primary volume. Do not check secondary volume.
exit_on_error ncpcon volume VOL1

# This line is added if OES CIFS is used as an advertising protocol
#exit_on_error rcnovell-cifs monitor

exit 0
```

- 8a Add a command to check the status of the secondary pool.

```
# Check the status of the secondary pool
exit_on_error status_fs /dev/pool/ARCPool1 /opt/novell/nss/mnt/.pools/ARCPool1
nsspool
```

- 8b Do not add a check for the secondary volume.

- 8c Click **Apply** to save your changes.

The changes do not take effect until the shadow volume cluster resource is brought online.

- 9 Click **OK** to save all your changes and return to the Cluster Manager page.
- 10 Online the primary pool cluster resource. On the Cluster Manager page, select the check box next to the primary cluster resource, then click **Online**.
- 11 Verify that the primary cluster resource is running by going to the **Cluster Manager** page. The primary cluster resource is **Running**.

12 Verify that the shadow volume (VOL1) is mounted in NCP and is shadowed:

12a On the first node in the cluster, log in to OES Remote Manager for Linux as the `root` user.

12b Select **View File System**, then verify that the secondary pool `ARCPPOOL1` and the NSS volume `ARCVOL1` are listed under **File Systems**, but the secondary NSS volume is not listed under **NCP Volumes**.

File System Management ?

File Systems		
Mounted Device	Mount Location	
<i>i</i> rootfs	/	(95% free)
<i>i</i> udev	/dev	(99% free)
/dev/disk/by-id/scsi-36001c230c175cf000e70368e60a6e6fe-part2 /		
proc	/proc	
sysfs	/sys	
debugfs	/sys/kernel/debug	
devpts	/dev/pts	
securityfs	/sys/kernel/security	
adminfs	/admin	
<i>i</i> admin	/_admin	(100% free)
/dev/pool/POOL1 /opt/novell/nss/mnt/.pools/POOL1		
/dev/pool/ARCPPOOL1 /opt/novell/nss/mnt/.pools/ARCPPOOL1		
<i>i</i> ARCVOL1	/media/nss/ARCVOL1	(99% free)
<i>i</i> VOL1	/media/nss/VOL1	(99% free)

NCP Volumes	
<i>i</i> SYS	/usr/novell/sys
<i>i</i> ADMIN	/_admin
<i>i</i> VOL1	/media/nss/VOL1

12c Select **View File System > Dynamic Storage Technology Options**, then verify that the primary volume is listed under **Volume Information**, and that its status is **Shadowed**.

Dynamic Storage Technology Options ?

Dynamic Storage Technology allows you optimize the use of your storage by automatically moving data to storage best optimized for the data type or frequency of use. You can create one or more policies to manage, and optimize the use of your storage.

Volume Information			
Volume Name	Shadow	Status	
<i>i</i> SHARE	Add Shadow	Inventory	
<i>i</i> VOL2	Add Shadow	Inventory	
<i>i</i> DATA	Add Shadow	Inventory	
<i>i</i> VOL1	Shadowed	Inventory	View Log
<i>i</i> SYS	Add Shadow	Inventory	

Dynamic Storage Technology Policies			
Name	Volume	Last Executed	Total Files Moved
ProjectABC Exclude contracts	VOL1	15:58, 11/16/2017	27

[Create a new policy](#)

[Stop all running policies](#)

Duplicate File Resolution Options	
Broadcast conflict message to user:	<input checked="" type="checkbox"/>
Action to be taken:	Show duplicate shadow files
Submit	

ShadowFS Configuration
<input type="checkbox"/> Load ShadowFS (Enable only to give LUM-enabled users merged-view access to DST volumes via Samba or Linux file access protocols)
Submit

- 12d** Select **Manage NCP Services > Manage Shares**, click NCP/NSS bindings, then verify that the NCP Accessible parameter is turned off for the secondary volume, and turned on for the primary volume.

NCP / NSS Bindings

Warning:

When a NSS Volume is changed to be not accessible via NCP, it will be dismounted immediately as a NCP share point

Available NSS volumes		
NCP Accessible	Volume Name	Mount point
Yes: <input type="radio"/> No: <input checked="" type="radio"/> <input type="button" value="Save Selection"/>	ARCVOL1	/media/nss/ARCVOL1
Yes: <input checked="" type="radio"/> No: <input type="radio"/> <input type="button" value="Save Selection"/>	DATA	/media/nss/DATA

You can also look for the EXCLUDE_VOLUME entry in the /etc/opt/novell/ncp2nss.conf file.

- 13** Continue with [Section 17.7, “Configuring Shadow Volume Policies for the Clustered DST Volume Pair,”](#) on page 209.

17.6 Sample Scripts for a DST Pool Cluster Resource

The cluster scripts in this section assume the following setup for NSS volumes in the clustered DST volume pair. Ensure that you use the actual information from your setup.

Setup	Primary NSS Volume	Secondary NSS Volume
Server name for node 1	server38	server38
Cluster name for node 1	NCS1	NCS1
Cluster pool name	POOL1	ARCPPOOL1
NSS volume name	VOL1	ARCVOL1
Cluster resource virtual server name	NCS1-POOL1-SERVER	
Cluster resource IP address	10.10.10.38	
Volume ID on the cluster node	254	

- ◆ [Section 17.6.1, “Sample Load Script for a DST Shadow Volume,”](#) on page 208
- ◆ [Section 17.6.2, “Sample Unload Script,”](#) on page 208
- ◆ [Section 17.6.3, “Sample Monitor Script for a DST Volume,”](#) on page 209

17.6.1 Sample Load Script for a DST Shadow Volume

The following is a sample load script for a DST pool cluster resource.

```
#!/bin/bash
. /opt/novell/ncs/lib/ncsfuncs

# activate the secondary pool
exit_on_error nss /poolact=ARCPPOOL1

# activate the primary pool
exit_on_error nss /poolact=POOL1

# Optional delay to allow time for pools to activate before mounting the volume
sleep 10

#comment out the original volume mount command
#exit_on_error ncpcon mount VOL1=254

# Use the ncpcon mount command to create the shadow volume on mount
exit_on_error ncpcon mount VOL1=254,shadowvolume=ARCVOL1

exit_on_error add_secondary_ipaddress 10.10.10.38

exit_on_error ncpcon bind --ncpsservername=NCS1-POOL1-SERVER --ipaddress=10.10.10.38

#This line is added if OES CIFS is used as an advertising protocol
#exit_on_error novcifs --add '--vserver=".cn=NCS1-POOL1-
SERVER.ou=ncs.o=novell.t=AVALON_TREE."'
--ip-addr=10.10.10.38

# If shadowfs is used, wait for shadowfs to start
#for (( c=1; c<=10; c++ )) do
# if [ ! -d /media/shadowfs/VOLUME/._NETWARE ]; then sleep 5; fi
#done

exit 0
```

17.6.2 Sample Unload Script

The following is a sample unload script for a DST pool cluster resource.

```
#!/bin/bash
. /opt/novell/ncs/lib/ncsfuncs

# This line is added if OES CIFS is used as an advertising protocol
#ignore_error novcifs --remove '--vserver=".cn=NCS1-POOL1-
SERVER.ou=ncs.o=novell.t=AVALON_TREE."' --ip-addr=10.10.10.38

# If shadowfs is used, unload the volume in FUSE
#ignore_error fusermount -u /media/shadowfs/VOL1

ignore_error ncpcon unbind --ncpsservername=NCS1-POOL1-SERVER --ipaddress=10.10.10.38

ignore_error del_secondary_ipaddress 10.10.10.38

# Deactivate the primary pool
ignore_error nss /pooldeact=POOL1

# Deactivate the secondary pool
ignore_error nss /pooldeact=ARCPPOOL1

exit 0
```


17.6.3 Sample Monitor Script for a DST Volume

The following is a sample unload script for a DST pool cluster resource.

```
#!/bin/bash
. /opt/novell/ncs/lib/ncsfuncs

# Check the status of the secondary pool
exit_on_error status_fs /dev/pool/ARCPPOOL1 /opt/novell/nss/mnt/.pools/ARCPPOOL1 nsspool

# Check the status of the primary pool
exit_on_error status_fs /dev/pool/POOL1 /opt/novell/nss/mnt/.pools/POOL1 nsspool

exit_on_error status_secondary_ipaddress 10.10.10.38

# Check the status of the primary volume. Do not check secondary volume.
exit_on_error ncpcon volume VOL1

# This line is added if OES CIFS is used as an advertising protocol
#exit_on_error rcnovell-cifs monitor

exit 0
```

17.7 Configuring Shadow Volume Policies for the Clustered DST Volume Pair

After you have modified the scripts for the primary pool cluster resource to make it a DST pool cluster resource, you are ready to bring the resource online and create policies for it. For planning information, see [“Policies for DST Nodes and Volumes in a Cluster” on page 186](#).

- 1 Create shadow volume policies for the clustered DST volume pair as described in [Chapter 11, “Creating and Managing Policies for Shadow Volumes,” on page 113](#).
- 2 If a shadow volume policy uses the **All Shadow Volumes** option, you must copy the policy information from the `/usr/novell/sys/._NETWARE/shadow_policy.xml` file to the same file on each node.

Copy only the entries for policies that you intend to apply for all DST shadow volume pairs across all nodes.

17.8 Renaming a Shared Pool in a DST Cluster Resource

The DST cluster resource is a pool cluster resource for the primary pool that has been modified to manage both the primary pool and secondary pool. When you use an NSS management tool to rename a cluster-enabled pool, it automatically changes the Pool object name in eDirectory; updates the pool name for its Volume objects; and updates the pool name that appears in the load, unload, and monitor scripts for the pool’s cluster resource. However, renaming a secondary pool cannot also update its name in the scripts for the primary pool’s resource, because the secondary pool’s default Pool object and Volume object do not contain information about the primary pool. You must manually modify the scripts for the DST cluster resource after you rename a secondary pool.

Renaming the shared primary pool or secondary pool in the DST cluster resource does not affect the settings that define the DST shadow relationship.

The resource objects for a DST cluster resource are typically named when you cluster-enable the primary pool. Renaming a primary pool does not modify the names of the resource objects. To rename a primary pool name in a way that also modifies names of resource objects, you can disable clustering for the primary pool, rename the pool, and then enable clustering for the renamed pool. This process is described in [“Renaming a Shared Pool and Its Resource Objects”](#).

Afterwards, reconfigure the primary pool cluster resource to include the necessary commands for the secondary pool and the DST shadow volume pair, just as you did when you first set up the DST cluster resource. Because the primary volume name and secondary volume name are not affected by the pool rename, you do not need to modify the related commands in the `/etc/opt/novell/ncp2nss.conf` file and `/etc/opt/novell/ncpserv.conf` file.

To rename a shared pool in a DST cluster resource:

- 1 Log in to the master node as the `root` user, then launch a terminal console.
- 2 Ensure that the pool cluster resource is running on the master node. Enter

```
cluster status
```

If you need to migrate the resource, enter the following at the command prompt:

```
cluster migrate <resource_name> <master_node_name>
```

- 3 Take the pool cluster resource offline. At the command prompt, enter

```
cluster offline <resource_name>
```

You can verify the offline state by entering

```
cluster status
```

- 4 Activate the pool to be renamed locally on the master node:

- 4a At the command prompt, enter

```
nssmu
```

- 4b Select **Pools** to view a list of pools. The pool state is **Deactive**.

- 4c Select the pool, then press F7 **Activate**.

- 5 Rename the pool:

- 5a On the NSSMU Pools page, select the pool.

- 5b Press F6 **Rename**.

- 5c Read the warning message that the pool’s volumes will be dismounted, then click **OK** to confirm that you want to rename the pool.

- 5d Specify the new name, then press Enter.

- 5e After the rename, the pool is in a deactive state.

- 6 Press Esc twice to exit NSSMU.

- 7 Verify that the pool name was modified in the pool resource scripts:

- 7a In iManager, select **Clusters > My Clusters**, then select the cluster that you want to manage.

- 7b On the Cluster Manager page, click the name link of the pool resource.

Remember that the pool resource name has not been modified by the pool rename.

- 7c On the Resource Properties page, click the **Scripts** tab.

7d Verify that the new pool name is used in the load, unload, and monitor scripts.

If you modified the secondary pool name, you must manually modify that line in the script.

Remember that the resource name and virtual server name are not modified by a pool rename. Do not modify those names in scripts.

7d1 On the Load Script page, verify that the pool name was changed in the `nss` command. If you modify the load script, click **Apply**.

```
exit_on_error nss /poolact=<new_pool_name>
```

For example:

```
exit_on_error nss /poolact=PUSERS
```

7d2 Click **Unload Script**, then verify that the pool name was changed in the `nss` command. Also check any custom commands where the volume name appears. If you modify the unload script, click **Apply**.

```
exit_on_error nss /pooldeact=<new_pool_name>
```

For example:

```
exit_on_error nss /pooldeact=PUSERS
```

7d3 Click **Monitor Script**, then verify that the pool name was changed in the `status_fs` command. Also check any custom commands where the pool name appears. If you modify the monitor script, click **Apply**.

```
exit_on_error status_fs /dev/pool/<new_pool_name> /opt/novell/nss/mnt/.pools/<new_pool_name>
```

For example:

```
exit_on_error status_fs /dev/pool/PUSERS /opt/novell/nss/mnt/.pools/PUSERS
```

7e Click **Cancel** to return to the Cluster Manager page, or click **OK** to apply any changes you made in the scripts.

8 Bring the pool cluster resource online. At the command prompt, enter

```
cluster online <resource_name>
```

You can verify the online state by entering

```
cluster status
```

17.9 Renaming a Shared Volume in a DST Cluster Resource

The DST cluster resource is a pool cluster resource for the primary pool that has been modified to manage both the primary pool and secondary pool. Renaming a volume changes the Volume object name, updates the volume information for the pool, and updates the volume name in its pool's resource scripts. However, a volume rename for a secondary volume cannot also update the scripts

for the primary pool's resource because its default Pool object and Volume object do not contain information about the primary pool. You must manually modify the scripts for the DST cluster resource after you rename a secondary volume.

When you rename a shared volume in a DST cluster resource, you also affect the settings that define the DST shadow relationship. The DST settings are not automatically updated when you rename the primary volume or secondary volume.

The DST relationship between the primary volume and secondary volume is defined by the `ncpcon mount` command in the DST cluster resource. DST settings that use the volume names are found in the following files that exist on each cluster node. You must manually remove the information on each node.

- ♦ `/etc/opt/novell/ncp2nss.conf`

This file contains the `EXCLUDE_VOLUME <secondary_volume_name>` entry that prevents the secondary volume from being mounted in NCP. For example:

```
EXCLUDE_VOLUME VOLD
```

The rename function does not remove this entry from the file. Cluster Services also does not remove the entry from the file. After the volume is renamed, when you bring the resource online on a node, a new entry is made for the secondary volume with the new name. You must manually remove the entry for the old volume name from the file on each node as part of the rename process.

- ♦ `/etc/opt/novell/ncpserv.conf`

This file contains the shadow volume entry. For example:

```
SHADOW_VOLUME HOME /media/nss/VOLD
```

After you rename the primary volume or secondary volume, a new entry is made in the file with the new primary volume name or secondary volume path. You must remove the entry for the DST volume from the file as part of the rename process.

After you rename the volume, new entries are added to the `ncp2nss.conf` file and `ncpserv.conf` file on the node where you bring the modified resource online.

To rename a volume in a shared DST shadow volume pair:

- 1 Log in to the master node as the `root` user, then launch a terminal console.
- 2 Ensure that the DST cluster resource is running on the master node. Enter

```
cluster status
```

If you need to migrate the resource, enter the following at the command prompt:

```
cluster migrate <resource_name> <master_node_name>
```

- 3 Take the DST cluster resource offline. At the command prompt, enter

```
cluster offline <resource_name>
```

You can verify the offline state by entering

```
cluster status
```

- 4 Open the `/etc/opt/novell/ncpserv.conf` file in a text editor, remove the entry for the DST shadow volume pair, then save the file. Repeat this task for every node in the cluster.

```
SHADOW_VOLUME <primary_volume_name> /media/nss/<secondary_volume_name>
```

For example:

```
SHADOW_VOLUME HOME /media/nss/VOLD
```

- 5 If you are renaming the secondary volume, open the `/etc/opt/novell/ncp2nss.conf` file in a text editor, remove the `EXCLUDE_VOLUME` entry for the volume, then save the file. Repeat this task for every node in the cluster.

```
EXCLUDE_VOLUME <secondary_volume_name>
```

For example:

```
EXCLUDE_VOLUME VOLD
```

- 6 Activate the pool locally on the master node:

- 6a At the command prompt, enter

```
nssmu
```

- 6b On the NSSMU **Main Menu**, select **Pools**, then press Enter.

- 6c On the Pools page, select the pool that contains the volume you want to rename, then press F7 **Activate**.

- 6d Press Esc to return to the NSSMU **Main Menu**.

- 7 Mount the volume to be renamed locally on the master node:

- 7a On the NSSMU **Main Menu**, select **Volumes**, then press Enter.

- 7b On the Volumes page, select the volume, then press F7 **Mount**.

- 8 On the Volumes page, view the volume properties to ensure that the attribute **Allow renaming the mount point** has been enabled. If it is disabled, select the attribute value, press **y** (Yes) to enable it, then select **Apply** and press Enter.

- 9 Rename the volume:

- 9a On the Volumes page, select the mounted NSS volume that you want to rename.

- 9b Press F3 **Rename**, specify the new name, then press Enter.

- 9c Wait for the volume to be automatically dismounted and remounted. When the page refreshes, the volume is mounted in the default mount point path `/media/nss/<new_volume_name>`.

If you try to force the page to refresh before the eDirectory changes have been synchronized, you might get an eDirectory error. The error condition is temporary and should not prevent the transaction from completing successfully.

- 10 Dismount the volume:

- 10a On the NSSMU Volumes page, select the volume, then press F7 **Dismount**.

- 10b Read the warning message that open files will be closed by the action, then press **y** (Yes) to confirm the dismount.

- 10c Press Esc to return to the NSSMU **Main Menu**.

- 11 Deactivate the pool that you activated in [Step 6](#):
 - 11a On the NSSMU **Main Menu**, select **Pools**, then press Enter.
 - 11b On the Pools page, select the pool, then press F7 **Deactivate**.
 - 11c Read the warning message that the pool's volumes will be dismounted, then press **y** (Yes) to confirm the deactivation.
 - 11d When the page refreshes, verify that the pool state is **Deactive**.
 - 11e Press Esc twice to exit NSSMU.
- 12 Modify the volume name in the scripts for its DST cluster resource:
 - 12a In iManager, select **Clusters > My Clusters**, then select the cluster that you want to manage.
 - 12b On the Cluster Manager page or Cluster Options page, click the name of the pool cluster resource to open the resource properties.
 - 12c Click the **Scripts** tab.
 - 12d On the Load Script page, modify the volume name in the `ncpcon mount` command. Also check any custom commands where the volume name appears. If you modify the load script, click **Apply**.

```
exit_on_error ncpcon mount
<primary_volume_name>=<volume_id>,shadowvolume=<secondary_volume_name>
```

For example, if you rename the primary volume from `HOME` to `USERS`, and the secondary volume is `VOLD`, change this line

```
exit_on_error ncpcon mount HOME=252,shadowvolume=VOLD
```

to this:

```
exit_on_error ncpcon mount USERS=252,shadowvolume=VOLD
```

If you rename the secondary volume name from `VOLD` to `HOME_SH`, and the primary volume is named `HOME`, change this line

```
exit_on_error ncpcon mount HOME=252,shadowvolume=VOLD
```

to this:

```
exit_on_error ncpcon mount HOME=252,shadowvolume=HOME_SH
```

- 12e Click **Unload Script**, then check any custom commands where the volume name appears. If you modify the unload script, click **Apply**.

The unload script does not contain any default volume commands.

- 12f If you renamed the primary volume, click **Monitor Script**, then modify the primary volume name in the `ncpcon volume` command. The secondary volume is not monitored because it is hidden from NCP. If you modify the monitor script, click **Apply**.

```
exit_on_error ncpcon volume <new_primary_volume_name>
```

For example, if you rename the primary volume from `HOME` to `USERS`, change this line

```
exit_on_error ncpcon volume HOME
```

to this:

```
exit_on_error ncpcon volume USERS
```

12g Click **Cancel** to return to the Cluster Manager page, or click **OK** to apply any changes you made in the scripts.

13 Bring the pool cluster resource online. At the command prompt, enter

```
cluster online <resource_name>
```

You can verify the online state by entering

```
cluster status
```

The first time that the modified DST cluster resource is brought online on a node, the new entries are added to the `ncp2nss.conf` file and `ncpserv.conf` file.

17.10 Removing the Shadow Relationship for a Clustered DST Volume Pair

Removing a clustered DST volume pair removes the shadow relationship between the primary and secondary storage area. You remove commands in the primary pool's cluster scripts that you added to manage the secondary pool and volume. Removing the shadow relationship does not remove the underlying volumes themselves. The files remain on whichever storage area they are on at the time when you remove the shadow relationship.

- ◆ [Section 17.10.1, "Planning to Remove the Shadow Relationship for a Clustered DST Volume Pair," on page 215](#)
- ◆ [Section 17.10.2, "Preparing to Remove a Shadow Relationship," on page 216](#)
- ◆ [Section 17.10.3, "Removing the Shadow Definition and NCP/NSS Bindings Exclusion on All Nodes," on page 217](#)
- ◆ [Section 17.10.4, "Preparing the Primary Pool Cluster Resource for Independent Use," on page 218](#)
- ◆ [Section 17.10.5, "Preparing the Secondary Pool and Volume for Independent Use," on page 219](#)

17.10.1 Planning to Remove the Shadow Relationship for a Clustered DST Volume Pair

As you plan to remove the shadow relationship for a clustered DST volume pair, consider the following short service outages that are involved:

- ◆ The data on the DST volume pair will be unavailable to users while you remove the shadow relationship as described in [Section 17.10.3, "Removing the Shadow Definition and NCP/NSS Bindings Exclusion on All Nodes," on page 217](#), and until you have performed the tasks necessary for the two volumes to function independently.
- ◆ An NDSD restart is required after you remove the shadow volume and NCP/NSS bindings information on each node in turn. This results in a brief service outage for all NSS volumes and NCP volumes that are mounted on the node at that time.

- ♦ The primary volume is available to users as an independent volume after you perform the procedure described in [Section 17.10.4, “Preparing the Primary Pool Cluster Resource for Independent Use,”](#) on page 218.
- ♦ The secondary volume is available to users as an independent volume after you perform one of the procedures described in [Section 17.10.5, “Preparing the Secondary Pool and Volume for Independent Use,”](#) on page 219.

17.10.2 Preparing to Remove a Shadow Relationship

Removing a shadow relationship does not automatically move files in either direction between the two volumes. The files remain undisturbed. The volumes function independently after the shadow relationship is successfully removed. Ensure that the files are distributed as desired before you remove the shadow relationship.

To move files between the two volumes to achieve a desired distribution of files:

- 1 In OES Remote Manager for Linux, log in as the `root` user.
- 2 Select **View File System > Dynamic Storage Technology Options**, locate the volume in the list, then click the **Inventory** link next to it.

View the volume inventory for the shadow volume to determine the space in use and the available space for both the primary and the secondary areas of the shadow volume. Ensure that there is sufficient free space available in either location for the data that you plan to move to that location.

- 3 Use any combination of the following techniques to move data between the two areas:
 - ♦ **Shadow Volume Policies:** Run an existing shadow volume policy by using the **Execute Now** option in the **Frequency** area of the policy. You can also create a new shadow volume policy that moves specific data, and run the policy by using the **One Time** and **Execute Now** options in the **Frequency** area of the policy.

For information about configuring policies to move data between the primary and secondary areas, see [Chapter 11, “Creating and Managing Policies for Shadow Volumes,”](#) on page 113.
 - ♦ **Inventories:** Use the detailed inventory reports or customized inventories to move specific files to either area.

For information about using the volume customized inventory options to move data between the primary and secondary areas, see [Section 14.6, “Generating a Custom Inventory Report,”](#) on page 167.
- 4 (Optional) While the DST pool cluster resource is online, you can delete volume-specific policies as described in [Section 11.8, “Deleting a Shadow Volume Policy,”](#) on page 127.

Shadow volume policies that you configured for the DST volume do not run after you remove the shadow volume relationship. The policy information is stored in the `/media/nss/<primary_volumename>/._NETWARE/shadow_policy.xml` file. Policy information is not automatically removed from the file when you remove the shadow relationship. If you later define a new shadow volume relationship for the primary volume, the policies apply to it.
- 5 Continue with [Section 17.10.3, “Removing the Shadow Definition and NCP/NSS Bindings Exclusion on All Nodes,”](#) on page 217.

17.10.3 Removing the Shadow Definition and NCP/NSS Bindings Exclusion on All Nodes

You must remove the shadow definition for the DST shadow volume pair and the NCP/NSS bindings exclusion for the secondary volume on each node in turn. This requires a restart of NDS and NCP2NSS on each node, which creates a short service outage for all NSS volumes on the node. You can minimize the impact by cluster migrating the pool cluster resources for the other NSS volumes to other nodes while you are modifying the configuration files on a given node.

- 1 Log in as the `root` user to the node where the primary pool cluster resource is online, then open a terminal console.
- 2 Offline the DST pool cluster resource that is managing the clustered shadow volume.

```
cluster offline resource_name
```

This unloads the cluster resource and deactivates the cluster pools and their volumes so that the cluster is not controlling them. Do not bring the primary resource or secondary resource online, and do not locally mount the volumes on any node at this time.

- 3 Remove the shadow volume and NCP/NSS bindings exclusion information from each node in the cluster:
 - 3a Log in to the node as the `root` user.

- 3b In a text editor, open the `/etc/opt/novell/ncp2nss.conf` file, remove the `EXCLUDE_VOLUME` line for the secondary volume from the file, then save the file.

```
EXCLUDE_VOLUME secondary_volumename
```

For example:

```
EXCLUDE_VOLUME ARCVOL1
```

- 3c In a text editor, open the `/etc/opt/novell/ncpserv.conf` file, remove the `SHADOW_VOLUME` line for the shadow volume from the file, then save the file.

```
SHADOW_VOLUME primary_volumename secondary_volume_path
```

For example:

```
SHADOW_VOLUME VOL1 /media/nss/ARCVOL1
```

- 3d Restart the eDirectory daemon by entering the following commands:

```
rndsd stop (or) systemctl stop ndsd.service
```

```
rndsd start (or) systemctl start ndsd.service
```

- 3e Restart the NCP/NSS IPC daemon to synchronize the changes you made to the `/etc/opt/novell/ncp2nss.conf` file. At the terminal console prompt, enter

```
systemctl restart ncp2nss.service
```

- 3f Repeat these steps for each node in the cluster.

- 4 After the shadow and bindings information has been removed from all nodes, continue with [Section 17.10.4, "Preparing the Primary Pool Cluster Resource for Independent Use,"](#) on page 218.

17.10.4 Preparing the Primary Pool Cluster Resource for Independent Use

In the primary pool cluster resource scripts, remove (or comment out) the lines for the management of the secondary pool, secondary volume, and `shadowfs`. This allows the pool cluster resource to function independently.

- 1 In iManager, select **Clusters**, then select **My Clusters**.
- 2 Select the name link of the cluster you want to manage.
- 3 On the Cluster Manager page, click the name link of the primary cluster resource to view its Cluster Pool Properties page, then click the **Scripts** tab.
- 4 On the **Scripts > Load Script** page, modify the load script of the primary pool cluster resource:

- 4a Remove or comment out the activation command for the secondary pool and the sleep command you added for the pool activation:

```
#exit_on_error nss /poolact=ARCPPOOL1
#sleep 10
```

- 4b Remove or comment out the `ncpcon mount` command for the shadow volume:

```
#exit on error ncpcon mount VOL1=254,shadowvolume=ARCVOL1
```

- 4c Add (or uncomment) a command to mount the NSS volume:

```
exit_on_error ncpcon mount <volume_name>=<volume_id>
```

Replace *volume_name* with the primary NSS volume name, such as VOL1.

Replace *volume_id* with a number that is unique across all nodes in the cluster, such as 254.

For example:

```
exit_on_error ncpcon mount VOL1=254
```

- 4d If `shadowfs` was used, remove or comment out the wait time for `shadowfs` to start.

```
# If shadowfs is used, wait for shadowfs to start
#for (( c=1; c<=10; c++ )) do
# if [ ! -d /media/shadowfs/VOLUME/._NETWARE ]; then sleep 5; fi
#done
```

- 4e Click **Apply** to save your changes.

The changes do not take effect until the cluster resource is brought online.

- 5 On the **Scripts > Unload Script** page, modify the unload script of the primary pool cluster resource:

- 5a Remove or comment out the deactivation command for the secondary pool:

```
#ignore_error nss /pooldeact=ARCPPOOL1
```

- 5b If `shadowfs` was used, remove or comment out the `fusermount -u` command.

```
# If shadowfs is used, unload the volume in FUSE
#ignore_error fusermount -u /media/shadowfs/VOL1
```

5c Click **Apply** to save your changes.

The changes do not take effect until the cluster resource is brought online.

6 On the **Scripts > Monitor Script** page, modify the monitor script of the primary pool cluster resource:

6a Remove or comment out the status command for the secondary pool:

```
# Check the status of the secondary pool
#exit_on_error status_fs /dev/pool/ARCPPOOL1 /opt/novell/nss/mnt/.pools/ARCPPOOL1
nsspools
```

6b Click **Apply** to save your changes.

The changes do not take effect until the cluster resource is brought online.

7 Click **OK** to return to the Cluster Manager page.

8 Online the revised pool cluster resource. On the Cluster Manager page, select the check box next to the pool cluster resource, then click **Online**.

The resource comes online as an independent pool and volume on a node in the resource's preferred nodes list.

If the resource goes comatose instead of coming online, take the resource offline, check the scripts, then try again.

9 Continue with [Section 17.10.5, "Preparing the Secondary Pool and Volume for Independent Use,"](#) on page 219.

17.10.5 Preparing the Secondary Pool and Volume for Independent Use

When you defined the clustered DST shadow volume pair, you might have used a clustered pool or a shared-but-not-cluster-enabled pool.

Apply one of the following methods to use the pool and volume independently:

- ♦ ["Modifying the Secondary Pool Cluster Resource" on page 219](#)
- ♦ ["Cluster-Enabling a Shared Secondary Pool" on page 220](#)
- ♦ ["Unsharing the Secondary Pool to Use It Locally on the Node" on page 222](#)

Modifying the Secondary Pool Cluster Resource

If you used a clustered secondary pool cluster resource, ensure that the volume ID is unique across all nodes in the cluster before you bring the pool resource online as an independent pool.

IMPORTANT: If you deleted the secondary pool cluster resource after you merged its information in the primary pool cluster resource scripts, the secondary resource no longer exists. You can cluster-enable the shared-but-not-cluster-enabled pool as described in ["Cluster-Enabling a Shared Secondary Pool" on page 220](#), or you can unshare the pool as described in ["Unsharing the Secondary Pool to Use It Locally on the Node" on page 222](#).

- 1 In iManager, select **Clusters**, then select **My Clusters**.
- 2 Select the name link of the cluster you want to manage.
- 3 Go to the secondary resource's load script and verify that the volume ID is unique for the secondary volume.

If you have assigned the volume ID to another clustered volume while the secondary resource was unused, the duplicate volume ID will cause the secondary resource to go comatose when you try to bring it online.

3a On the Cluster Manager page, click the name link of the secondary cluster resource to view its Cluster Pool Properties page, then click the **Scripts** tab.

3b On the **Scripts > Load Script** page, check the volume ID to ensure that it is unique:

```
exit_on_error ncpcon mount ARCVOL1=253
```

3c Click **OK** to save your changes and return to the Cluster Manager page.

The changes do not take effect until the cluster resource is brought online.

4 Online the secondary pool cluster resource. On the Cluster Manager page, select the check box next to the primary pool cluster resource, then click **Online**.

The resource comes online as an independent pool and volume on a node in the resource's preferred nodes list.

If the resource goes comatose instead of coming online, take the resource offline, check the scripts, then try again.

If the resource goes online successfully, you are finished.

Cluster-Enabling a Shared Secondary Pool

You can cluster-enable the shared pool and volume as an independent pool cluster resource under the following conditions:

- ◆ If you used a shared-but-not-clustered pool as the secondary pool.

In this case, the Pool object name and Volume object name contain the name of the node where they were originally created.

- ◆ If you used a cluster-enabled pool as the secondary and deleted the secondary pool cluster resource after you copied its commands into the DST pool cluster resource scripts.

In this case, the Pool object name and Volume object name contain the cluster name, because the objects were recreated when you cluster-enabled them.

Before you attempt to cluster-enable the shared pool, you must update the Pool object and Volume object in eDirectory to use the hostname of the server where you took the primary pool cluster resource offline.

1 In NSSMU, update the eDirectory object for the shared pool and volume.

You can alternatively use the Storage plug-in for iManager to update the eDirectory objects. Select the server where you took the clustered DST pool resource offline. Ensure that you dismount the volume and deactivate the pool after you have updated their objects.

1a Log in as the `root` user on the node where you took the primary pool cluster resource offline, then open a terminal console.

1b Launch NSSMU. At the command prompt, enter

```
nssmu
```

- 1c** Activate the pool and update its eDirectory object to create a Pool object that is named based on the hostname of the current node.
 - 1c1** In the NSSMU menu, select **Pools**, then press Enter.
 - 1c2** Select the secondary pool (ARCPool1), then press F7 to activate it.
 - 1c3** Select the secondary pool, press F4 (**Update NDS**), then press y (Yes) to confirm that you want to delete the old Pool object and add a new Pool object.
 - 1c4** Press Esc to return to the NSSMU menu.
- 1d** Mount the volume, update its eDirectory object to create a Volume object that is named based on the hostname of the current node, then dismount the volume.
 - 1d1** In the NSSMU menu, select **Volumes**, then press Enter.
 - 1d2** Select the secondary volume (ARCVOL1), then press F7 to mount it.
 - 1d3** Select the secondary volume, press F4 (**Update NDS**), then press y (Yes) to confirm that you want to delete the old Volume object and add a new Volume object.
 - 1d4** Select the secondary volume, then press F7 to dismount it.
 - 1d5** Press Esc to return to the NSSMU menu.
- 1e** Deactivate the pool.
 - 1e1** In the NSSMU menu, select **Pools**, then press Enter.
 - 1e2** Select the secondary pool (ARCPool1), then press F7 to deactivate it.
- 1f** Press Esc twice to exit NSSMU.
- 2** In iManager, select **Clusters > My Clusters**.
- 3** Select the name link of the cluster you want to manage.
- 4** Cluster-enable the shared pool.

For detailed instructions, see “[Cluster-Enabling an Existing NSS Pool and Its Volumes](#)” in the *OES 2018 SP3: OES Cluster Services for Linux Administration Guide*.

- 4a** Click the **Cluster Options** tab, then click **New**.
- 4b** On the Resource Type page, select **Pool**, then click **Next**.
- 4c** On the Cluster Pool Information page:
 - 4c1** Browse to select the secondary pool, such as `<hostname>_ARCPool1_POOL`.
 - 4c2** Specify a unique IP address.
 - 4c3** Select the NCP, AFP, or CIFS check boxes for the advertising protocols that you want to enable for the volume.

NCP is selected by default and is required to support authenticated access to data via the OES Trustee Model. If OES CIFS or OES AFP is not installed, selecting its check box has no effect.
 - 4c4** If you enable CIFS, verify the default name in the **CIFS Server Name** field.

You can modify this name. The name must be unique and can be up to 15 characters, which is a restriction of the CIFS protocol.
 - 4c5** **Online Resource After Creation** is disabled by default. This allows you to review the settings and scripts before you bring the resource online for the first time.

4c6 **Define Additional Properties** is enabled by default. This allows you to set resource policies and preferred nodes before the resource is brought online.

4c7 Click **Next**.

4d On the Resource Policies page, configure the policies for the start, failover, and failback mode, then click **Next**.

4e On the Resource Preferred Nodes page, assign and rank order the preferred nodes to use for the resource, then click **Finish**.

5 (Optional) Enable monitoring for the pool cluster resource.

5a On the Cluster Options page, select the name link for the resource to open its Properties page.

5b Click the **Monitoring** tab.

5c Select **Enable Resource Monitoring**, set the **Polling Interval**, **Failure Rate**, and **Failure Action**, then click **Apply**.

5d Click the **Scripts** tab, then click Monitor Script.

5e View the script settings and verify that they are as desired.

5f If you modify the script, click **Apply**.

5g Click **OK**.

6 Bring the pool cluster resource online. Click the **Cluster Manager** tab, select the resource check box, then click **Online**.

If the resource goes online successfully, you are finished.

Unsharing the Secondary Pool to Use It Locally on the Node

You can unshare the shared pool and volume and use them locally on the node under the following conditions:

- ◆ If you used a shared-but-not-clustered pool as the secondary pool.

In this case, the Pool object name and Volume object name contain the name of the node where they were originally created.

- ◆ If you used a cluster-enabled pool as the secondary and deleted the secondary pool cluster resource after you copied its commands into the DST pool cluster resource scripts.

In this case, the Pool object name and Volume object name contain the cluster name, because the objects were recreated when you cluster-enabled them.

Before you attempt to use pool and volume locally, you must update the Pool object and Volume object in eDirectory to use the hostname of the server where you took the primary pool cluster resource offline.

To mount the secondary volume as an independent local volume:

- 1** Log in as the `root` user on the node where you took the primary pool cluster resource offline, then open a terminal console.
- 2** Launch NSSMU. At the command prompt, enter

```
nssmu
```

- 3 In NSSMU, update the eDirectory object for the shared pool and volume.

You can alternatively use the Storage plug-in for iManager to update the eDirectory objects. Select the server where you took the clustered DST pool resource offline.

- 3a Activate the pool and update its eDirectory object to create a Pool object that is named based on the hostname of the current node.

- 3a1 In the NSSMU menu, select **Pools**, then press Enter.

- 3a2 Select the secondary pool (ARCPOOL1), then press F7 to activate it.

- 3a3 Select the secondary pool, press F4 (**Update NDS**), then press y (Yes) to confirm that you want to delete the old Pool object and add a new Pool object.

- 3a4 Press Esc to return to the NSSMU menu.

- 3b Mount the volume, then update its eDirectory object to create a Volume object that is named based on the hostname of the current node.

- 3b1 In the NSSMU menu, select **Volumes**, then press Enter.

- 3b2 Select the secondary volume (ARCVOL1), then press F7 to mount it.

- 3b3 Select the secondary volume, press F4 (**Update NDS**), then press y (Yes) to confirm that you want to delete the old Volume object and add a new Volume object.

- 3b4 Press Esc to return to the NSSMU menu.

- 4 In the NSSMU menu, select **Devices**, then press Enter.

- 5 Disable sharing for the device. Select the device, press F6 to unshare the device, then press y (Yes) to confirm.

If NSSMU does not allow you to unshare the device, you can use the SAN management software to ensure that the device is allocated only to the current server, and then try again.

Before you unshare the device, ensure that the device contains only the pool that you are changing to local use. It should not contain other shared pools or SBD partitions.

- 6 In the NSSMU menu, select **Pools**, then press Enter.

- 7 Select the pool and verify that it is unshared.

- 8 Press Esc twice to exit NSSMU.

17.11 Upgrading a Cluster with DST Resources from OES 2 SP3 to OES 2015 or Later

You can upgrade a cluster from OES 2 SP3 to OES 2015 or later. The same release version of OES 2015 or later must be installed and running on each new or upgraded node in the cluster.

The procedure in this section refers to OES 2015 or later.

If you replace an OES 2 SP3 node with OES 2015 or later node instead of using an in-place upgrade to OES 2015 or later, you must manually configure the global policies and move the **All Shadowed Volumes** policies data to the OES 2015 or later node.

- 1 On the replacement OES 2015 or later node, set the same global configuration policies that are set on the OES 2 SP3 nodes.
- 2 If you have **All Shadowed Volumes** policies set on the OES 2 SP3 nodes, copy the information in the `/usr/novell/sys/._NETWARE/shadow_policy.xml` file to the same path on the replacement OES 2015 or later nodes.
- 3 In iManager, use the Clusters plug-in to modify the preferred nodes list for each DST pool cluster resource to specify only OES 2015 or later nodes as fail-over candidates.
- 4 Launch a terminal console, then cluster migrate the DST pool cluster resource to an OES 2015 or later node:

```
cluster migrate <dst_volume_cluster_resource> [oes2015_node_name]
```

If the node name is not specified, the resource automatically is brought online on an available node in its preferred nodes list. To bring it online on a specific node, you can specify a node in its preferred nodes list. The resource will be mounted there if it is possible to do so; that is, if the node is available and there are no resource conflicts.

A

Commands and Utilities for Dynamic Storage Technology

This section describes commands and utilities for Dynamic Storage Technology (DST) for Open Enterprise Server (OES) for Linux.

- ♦ [Section A.1, “Using NCPCON Commands for DST,” on page 225](#)
- ♦ [Section A.2, “NCPCON Commands for Managing DST,” on page 226](#)
- ♦ [Section A.3, “NCPCON Commands for DST in a OES Cluster Services Cluster,” on page 231](#)
- ♦ [Section A.4, “Configuring Global DST Policies by Using the SET Command,” on page 233](#)
- ♦ [Section A.5, “DST Commands for /etc/opt/novell/ncpserv.conf,” on page 236](#)
- ♦ [Section A.6, “DST Commands for /etc/opt/novell/shadowfs.conf,” on page 237](#)
- ♦ [Section A.7, “DST EXCLUDE_VOLUME Command for /etc/opt/novell/ncp2nss.conf,” on page 238](#)
- ♦ [Section A.8, “DST Shadow Volume Information in /etc/NCPVolumes,” on page 238](#)
- ♦ [Section A.9, “DST ShadowFS Volume Information in /etc/mtab.shadowfs,” on page 239](#)
- ♦ [Section A.10, “Verifying and Syncing ACLs \(Inherited Rights Filter \(IRF\) and Trustees\) from DST Primary Volume to Shadow Volume,” on page 239](#)

A.1 Using NCPCON Commands for DST

The NetWare Core Protocol (NCP) Console Command (NCPCON) utility provides an interface for issuing NetWare commands in a Linux environment. You can issue commands via the NCPCON in three modes:

- ♦ [Section A.1.1, “Interactive Mode,” on page 225](#)
- ♦ [Section A.1.2, “Command Line Mode,” on page 226](#)
- ♦ [Section A.1.3, “Scripting Mode,” on page 226](#)

A.1.1 Interactive Mode

Open a terminal console, log in as the `root` user, then enter

```
ncpcon
```

This opens the NCPCON interactive console in the terminal console, so you can enter the NCP Server console commands. Enter `exit` to stop interactive mode.

Escaping the quotation mark character (") is not required when you enter the command from the `ncpcon` prompt.

For example, enter the following commands from the `ncpcon` prompt:

```
mount sys
```

```
shift VOL1:"path\file name with spaces.txt" shadow
send "hello world" to all
```

A.1.2 Command Line Mode

For command line mode, issue an NCP Server command at a terminal console prompt by preceding the command with `ncpcon`:

```
ncpcon [command]
```

When you use `ncpcon` to issue commands directly from the console command prompt, you must escape the quotation mark character (") by preceding the character with a backslash (\), such as \".

For example, enter the following commands from the terminal console prompt:

```
ncpcon mount sys
ncpcon shift VOL1:\"path\file name with spaces.txt\" shadow
ncpcon send \"hello world\" to all
```

A.1.3 Scripting Mode

For scripting mode, issue the NCP Server command in the script by preceding the command with `ncpcon`, then placing quotation marks (") around the NCP Server command:

```
ncpcon "[command]"
```

If the command includes a field that must be contained in quotation marks (such as a file name), you must escape each internal quotation mark character (") with a backslash (\) character, such as \".

For example, place the following commands in a script file:

```
ncpcon "mount sys"
ncpcon "shift VOL1:\"path\file name with spaces.txt\" shadow"
ncpcon "send \"hello world\" to all"
```

A.2 NCPCON Commands for Managing DST

The commands in this section can be used only with the NCP Console Command utility. You can issue the commands from the NCP Console interactive mode, or precede the command with `ncpcon` when issuing it from a script or at a terminal console prompt as the `root` user. For information, see [Section A.1, “Using NCPCON Commands for DST,” on page 225](#).

- ◆ [Section A.2.1, “Creating a DST Shadow Volume Pair,” on page 227](#)
- ◆ [Section A.2.2, “Removing the Shadow Relationship, or Unlinking the Volumes,” on page 227](#)
- ◆ [Section A.2.3, “Listing or Moving Files that Match Search Criteria,” on page 228](#)
- ◆ [Section A.2.4, “Listing or Moving a File, or Shifting a File between Volumes,” on page 230](#)

A.2.1 Creating a DST Shadow Volume Pair

ncpcon create shadow_volume <primary_volumename> <shadow_path>

Creates a non-clustered shadow association between a primary NSS volume and secondary NSS volume, and adds the SHADOW_VOLUME mount information to the /etc/opt/novell/ncpserv.conf file.

When you issue the command from the NCP Console, you do not need to restart ndsd in order for the changes to take effect. When you issue the command from a Linux prompt, you must restart ndsd in order for the changes to take effect.

OPTIONS

primary_volumename

Specifies the volume name for the primary NSS volume, such as VOL1.

shadow_path

Specifies the Linux path of the mount location for the secondary NSS volume, such as /media/nss/ARCVOL1.

EXAMPLES

create shadow_volume VOL1 /home/shadows/VOL1

Creates a shadow volume where VOL1 is the primary storage area and /home/shadows/VOL1 is its mount point as a shadow volume.

A.2.2 Removing the Shadow Relationship, or Unlinking the Volumes

ncpcon remove shadow_volume [/l] [/i] [/f] <primary_volumename>

Removes the non-clustered shadow relationship between a primary NSS volume and a secondary NSS volume, and removes the SHADOW_VOLUME command from the /etc/opt/novell/ncpserv.conf file. You must unmount the volume before you issue the command.

IMPORTANT: You can use this command as part of the process to unlink the primary and secondary volumes of a non-clustered DST shadow volume. For information, see [Section 10.12, “Removing the Shadow Relationship for a Non-Clustered DST Shadow Volume,”](#) on page 110.

Typically, you specify the /l option, which leaves the files in place on the primary volume and secondary volume, and removes the shadow relationship. This is equivalent to the **Volume Tasks > Remove Shadow Action Options > Remove Shadow** option in OES Remote Manager.

When the /l option is not used, the command attempts to move all files on the secondary volume to the primary volume, and then removes the shadow relationship between the two volumes. Ensure that the primary volume has sufficient space to accommodate the files before you unmount the volume and issue the remove command. Moving the files can take some time, depending on how much data must be moved. If a file move fails, the unlinking of the shadow relationship also fails. You can use the /i option to ignore file move errors and allow the unlinking to succeed. After the files on the secondary volume have been moved to the primary volume, the shadow relationship is removed, and a summary report is created and displayed.

OPTIONS

primary_volumename

Specifies the volume name for the primary NSS volume, such as VOL1.

/l

Leaves the files in place on the two volumes and removes the shadow relationship.

/i

Ignores any file move errors that might occur if you issue the command without the **/l** option, and allows the unlinking of the shadow relationship to succeed.

For example, if there are duplicate files on the volumes, the duplicate instance on the secondary volume cannot be moved to the primary volume, and the shadow relationship cannot be unlinked. Using the **/i** option ignores the file move error and allows the relationship to be unlinked.

/f

Provides a full detail report of actions taken. Use this option to understand which file moves might be failing.

EXAMPLES

Issue the following commands from the NCP Console, or add `ncpcon` at the front of the command when issuing it from a script or at a terminal console prompt.

```
ncpcon remove shadow_volume /i /f VOL1
```

Removes the shadow relationship for shadow volume `VOL1`, and moves all files from the secondary storage area to the primary storage area. You must dismount `VOL1` before you issue this command. File move errors are ignored. Full details of the actions taken are reported.

```
remove shadow_volume /l VOL1
```

Removes the shadow relationship for shadow volume `VOL1`, and leaves files where they currently are on the secondary storage area and the primary storage area. You must dismount `VOL1` before you issue this command.

A.2.3 Listing or Moving Files that Match Search Criteria

```
ncpcon shadow <primary_volumename> operation=<lp | ls | mp | ms> [options]
```

Allows you to list files on the shadow volume, or to move files between the primary storage area and the secondary storage area based on specified search criteria. All files on the selected shadow volume that match the criteria are moved. Use the command from within `cron` jobs to automate data partitioning.

Replace *primary_volumename* with the volume name for the primary NSS volume, such as `VOL1`.

OPERATION OPTIONS

lp

Lists primary files. Lists all files currently residing on the primary storage area.

ls

Lists shadow files. Lists all files currently residing on the secondary storage area.

mp

Moves files to primary. Moves files that match the specified criteria to the primary storage area from the secondary storage area.

ms

Moves files to shadow. Moves files that match the specified criteria to the secondary storage area from the primary storage area.

OPTIONS

pattern="searchPattern"

Specifies the file pattern to match against.

owner="username.context"

Specifies the eDirectory user name and context of the owner of the files to match against.

uid=uidValue

Specifies the Linux user ID to match against.

time=[time_field]

Specifies which time field to match against, where the *time_field* is:

[m] [a] [c]

- ♦ **m**: Last time modified (content)
- ♦ **a**: Last time accessed
- ♦ **c**: Last time changed (metadata)

range=[time_period]

Specifies which time period to match against, where the *time_period* is:

[a] [b] [c] [d] [e] [f] [g] [h] [i] [j]

- ♦ **a**: Within last day
- ♦ **b**: 1 day to 1 week
- ♦ **c**: 1 week to 2 weeks
- ♦ **d**: 2 weeks to 1 month
- ♦ **e**: 1 month to 2 months
- ♦ **f**: 2 months to 4 months
- ♦ **g**: 4 months to 6 months
- ♦ **h**: 6 months to 1 year
- ♦ **i**: 1 year to 2 years
- ♦ **j**: More than 2 years

size=[size_differential]

Specifies the size differential to match against, where the *size_differential* is:

[a] [b] [c] [d] [e] [f] [g] [h] [i] [j] [k]

- ♦ **a**: Less than 1 KB
- ♦ **b**: 1 KB to 4 KB
- ♦ **c**: 4 KB to 16 KB
- ♦ **d**: 16 KB to 64 KB
- ♦ **e**: 64 KB to 256 KB

- ♦ **f:** 256 KB to 1 MB
- ♦ **g:** 1 MB to 4 MB
- ♦ **h:** 4 MB to 16 MB
- ♦ **i:** 16 MB to 64 MB
- ♦ **j:** 64 MB to 256 MB
- ♦ **k:** More than 256 MB

output="filename"

Outputs the search results to the specified file.

EXAMPLES

shadow VOL1 operation=ls pattern="*.exe"

Lists all files of type EXE that currently reside on the secondary storage area for the shadow volume VOL1.

shadow VOL1 operation=lp size=g

Lists all files of sizes between 1 MB to 4 MB that currently reside on the primary storage area for the shadow volume VOL1.

shadow VOL1 operation=ms time=m range=j

Moves all files on the primary storage area that have not been modified in more than two years from the primary storage area to the secondary storage area for the shadow volume VOL1.

A.2.4 Listing or Moving a File, or Shifting a File between Volumes

ncpcon shift "primary_volumename:\path\filename" [primary | shadow]

Returns the specified file's location as being on the primary storage area or secondary storage area. Specify the primary or secondary options to move the specified file from its current location to the specified storage area.

IMPORTANT: The `shift` command works only at the command line, and not in `ncpcon` interactive mode. Enter the command as the `root` user at a terminal console prompt.

OPTIONS

primary

Moves the specified file from the secondary storage area to the primary storage area. The file must be closed when you issue the command; otherwise, the command fails.

shadow

Moves the specified file from the primary storage area to the secondary storage area. The file must be closed when you issue the command; otherwise, the command fails.

EXAMPLES

Enter the commands as the `root` user at a terminal console prompt.

ncpcon shift VOL1:"path\textfile.txt"

Shows the specified file's storage area location in the shadow volume as primary (the primary storage area) or shadow (the secondary storage area) for the shadow volume `sys`.

```
ncpcon shift VOL1:"path\textfile.txt" primary
```

Moves the specified file's storage area location from the secondary storage area to the primary storage area for the shadow volume `sys`.

```
ncpcon shift VOL1:"path\textfile.txt" shadow
```

Moves the specified file's storage area location from the primary storage area to the secondary storage area for the shadow volume `sys`.

A.3 NCPCON Commands for DST in a OES Cluster Services Cluster

NCPCON supports the commands in this section for use with Dynamic Storage Technology in combination with OES Cluster Services for Linux clusters.

Use the syntax examples in this section in cluster load scripts to mount the volume in a cluster. With clustering, no changes are needed to the `ncpserv.conf` file for shadowing. The primary volume information should not be manually added to the `ncpserv.conf` file.

When the primary volume has a state of **Shadowed**, the volume ID that you assign as its NCP volume ID represents the DST shadow volume pair of volumes. The secondary volume does not have a separate volume ID when it is in the shadow relationship.

- ♦ [Section A.3.1, "Scenario 1: Primary NSS and Shadow NSS," on page 231](#)
- ♦ [Section A.3.2, "Scenario 2: Primary Non-NSS and Shadow Non-NSS \(Not supported\)," on page 232](#)
- ♦ [Section A.3.3, "Scenario 3: Primary Non-NSS and Shadow NSS \(Not supported\)," on page 232](#)
- ♦ [Section A.3.4, "Scenario 4: Primary NSS and Shadow Non-NSS \(Not Supported\)," on page 232](#)

A.3.1 Scenario 1: Primary NSS and Shadow NSS

```
ncpcon mount volumename=volID,SHADOWVOLUME=shadow_volumename
```

Use this command in a cluster load script when the primary volume is an NSS volume and the secondary volume is an NSS volume. Both NSS volumes must already exist. The secondary NSS volume must already exist on the system and be mountable in NSS, but not available to NCP (that is, its NCP/NSS bindings are disabled by using OES Remote Manager). The primary pool must be cluster-enabled. The secondary pool must be shared. The secondary pool can be clustered, but its Cluster object and resource IP address are not used while its volume is used in a clustered DST volume.

Replace `volID` with a value from 0 to 254 as the server volume ID to ensure that the volume has the same ID on all servers when it is mounted in a cluster resource.

EXAMPLE

```
ncpcon mount VOL1=254,SHADOWVOLUME=ARCVOL1
```

Mounts the NSS volume named `VOL1` with a volume ID of 254. The primary volume is an existing NSS volume named `VOL1 (/media/nss/VOL1)`. The secondary volume is an existing NSS volume named `ARCVOL1 (/media/nss/ARCVOL1)`.

A.3.2 Scenario 2: Primary Non-NSS and Shadow Non-NSS (Not supported)

`ncpcon mount volumename=volID,SHADOWPATH=shadowpath,path=primarypath`

Use this command when the primary volume is a non-NSS volume and the secondary volume is a non-NSS volume.

Replace *volID* with a value from 0 to 254 as the server volume ID to ensure that the volume has the same ID on all servers when it is mounted in a cluster resource.

EXAMPLE

`ncpcon mount VOL1=254,SHADOWPATH=/media/ncpvolumes/ARCVOL1,path=/media/ncpvolumes/VOL1`

Mounts the NCP volume named `VOL1` with a volume ID of 254. The primary volume's path is `/media/ncpvolumes/VOL1`. The secondary volume's path is `/media/ncpvolumes/ARCVOL1`.

A.3.3 Scenario 3: Primary Non-NSS and Shadow NSS (Not supported)

`ncpcon mount volumename=volID,SHADOWVOLUME=shadow_volumename,path=primarypath`

Use this command when the primary volume is a non-NSS volume and the secondary volume is an NSS volume. The secondary NSS volume must already exist on the system and be mounted in NSS, but not available to NCP (that is, its NCP/NSSNCP/NSS bindings are disabled by using OES Remote Manager).

Replace *volID* with a value from 0 to 254 as the server volume ID to ensure that the volume has the same ID on all servers when it is mounted in a cluster resource.

EXAMPLE

`ncpcon mount VOL1=254,SHADOWVOLUME=ARCVOL1,path=/media/ncpvolumes/VOL1`

Mounts the NCP volume named `VOL1` with a volume ID of 254. The primary volume's path is `/media/ncpvolumes/VOL1`. The secondary volume is an existing NSS volume named `ARCVOL1` (mounted at `/media/nss/ARCVOL1`).

A.3.4 Scenario 4: Primary NSS and Shadow Non-NSS (Not Supported)

`ncpcon mount volumename=volID,SHADOWPATH=shadowpath`

Use this command when the primary volume is an NSS volume and the secondary volume is a non-NSS volume. The NSS volume must already exist on the system. The remote server must reside in the same eDirectory partition and tree as the DST server.

Replace *volID* with a value from 0 to 254 as the server volume ID to ensure that the volume has the same ID on all servers when it is mounted in a cluster resource.

EXAMPLE

`ncpcon mount VOL1=254,SHADOWPATH=/media/ncpvolumes/ARCVOL1`

Mounts an NSS volume named `VOL1` with a volume ID of 254. The primary volume is an existing NSS volume named `VOL1` (`/media/nss/VOL1`). The secondary volume is an NCP volume named `ARCVOL1` that is mounted at `/media/ncpvolumes/ARCVOL1`.

A.4 Configuring Global DST Policies by Using the SET Command

DST provides several global parameters for the `SET` command that can be used to customize DST for a given server. These settings control how DST behaves for all shadow volumes on the server. Initially, the parameters and default settings are in force, but the parameters are not explicitly added to the `/etc/opt/novell/ncpserv.conf` file. After you modify its default setting, an entry for the parameter and its new setting are added to the file. The parameter entry remains in the file even if you modify the setting back to the default.

IMPORTANT: If you use DST shadow volumes in a cluster, ensure that you set the same global policies on each OES node in the cluster where you plan to fail over the shared volumes.

- ◆ [Section A.4.1, “Understanding DST Parameters for the SET Command,” on page 233](#)
- ◆ [Section A.4.2, “Using OES Remote Manager to Configure DST Parameters for the SET Command,” on page 234](#)
- ◆ [Section A.4.3, “Using the `ncpcon set` Command to Configure DST Parameters,” on page 236](#)

A.4.1 Understanding DST Parameters for the SET Command

[Table A-1](#) lists the DST parameters for the `SET` command with their default values and valid options.

Table A-1 Manage NCP Services > Manage Server > Server Parameter Information

Parameter Name and Description	Default Value	Valid Values
<code>DUPLICATE_SHADOW_FILE_ACTION</code> Controls how duplicate files conflicts are handled. For information, see Section 7.3.1, “Understanding Conflict Resolution for Duplicate Files,” on page 55.	0	0 - Show duplicate shadow files (default) 1 - Hide duplicate shadow files 2 - Rename duplicate shadow files 3 - Delete duplicate files from shadow area 4 - Move duplicate shadow files to <code>/. _DUPLICATE_FILES</code>
<code>DUPLICATE_SHADOW_FILE_BROADCAST</code> Controls whether broadcast messages are sent to NCP users whenever duplicate files conflicts occur. For information, see Section 7.3.1, “Understanding Conflict Resolution for Duplicate Files,” on page 55.	1	0 - Disable 1 - Allow

Parameter Name and Description	Default Value	Valid Values
<p>REPLICATE_PRIMARY_TREE_TO_SHADOW</p> <p>By default, the REPLICATE_PRIMARY_TREE_TO_SHADOW option is set to 1. Ensure this option is always set to 1, so when the existing NSS volumes are configured for DST, the primary volume directory tree structure will get replicated to the shadow volume directory tree.</p> <p>Beginning with OES 2018 SP3, the NSS filesystem is DST aware for new volumes, so the primary volume directory tree structure will always be replicated on the shadow volume directory tree.</p> <p>For information, see Section 7.1, “Replicating Branches of the Primary File Tree in the Secondary File Tree,” on page 49.</p>	1	<p>0 - Disable</p> <p>1 - Allow</p>
<p>SHIFT_MODIFIED_SHADOW_FILES</p> <p>Controls whether a file is moved from the secondary file tree to the primary file tree based on its modification time.</p> <p>For information, see “Shift Modified Shadow Files” on page 51.</p>	1	<p>0 - Disable</p> <p>1 - Allow</p>
<p>SHIFT_ACCESSED_SHADOW_FILES</p> <p>Controls whether a file is moved from the secondary file tree to the primary file tree if it is accessed twice during a specific period of time. Use with SHIFT_DAYS_SINCE_LAST_ACCESS to specify the period of time.</p> <p>For information, see “Shift Accessed Shadow Files” on page 51.</p>	0	<p>0 - Disable</p> <p>1 - Allow</p>
<p>SHIFT_DAYS_SINCE_LAST_ACCESS</p> <p>Specifies the number of days to use when determining if a file should be moved back to the primary storage area. When it is used with SHIFT_ACCESSED_SHADOW_FILES, the parameter sets the time when files are migrated back to the primary storage area after the second access within the specified elapsed time.</p>	1	<p>0 - Disable</p> <p>1 to 365 (in days)</p>

A.4.2 Using OES Remote Manager to Configure DST Parameters for the SET Command

You can configure the DST parameters for the SET command by using OES Remote Manager for Linux.

- 1 In OES Remote Manager for Linux, select **Manage NCP Services**, then select **Manage Server**.
- 2 In the **Set Parameter Information** table, locate the DST parameter you want to configure.

The following server parameters are available. The settings shown are the default values. For information, see [Section A.4.1, “Understanding DST Parameters for the SET Command,”](#) on page 233.

DUPLICATE_SHADOW_FILE_ACTION	0
DUPLICATE_SHADOW_FILE_BROADCAST	1
REPLICATE_PRIMARY_TREE_TO_SHADOW	1
SHIFT_ACCESSED_SHADOW_FILES	0
SHIFT_MODIFIED_SHADOW_FILES	1
SHIFT_DAYS_SINCE_LAST_ACCESS	1

- 3 Modify settings by clicking the link for the value in the **Parameter Value** column to open a page where you can change the value.
- 4 In **New Value**, type the value for the parameter, then click **Change** to save and apply the setting.

Current Value	New Value
1	<input type="text" value="0"/> <input type="button" value="Change"/>
<input type="button" value="Back"/>	

- 5 If you enabled DUPLICATE_SHADOW_FILE_BROADCAST, ensure that NCP Server is configured to support broadcast messages by verifying that the Disable Broadcast (DISABLE_BROADCAST) parameter for the SET command is disabled:
 - 5a In OES Remote Manager for Linux, select **Manage NCP Services**, then select **Manage Server**.
 - 5b In the **Set Parameter Information** table, locate the DISABLE_BROADCAST parameter, then view the current value of the parameter. By default, the parameter is disabled (set to 0), which means that NCP Server supports broadcast messages.

DISABLE_BROADCAST	0
-------------------	-------------------

- 5c If the DISABLE_BROADCAST parameter is enabled (set to 1), click the link for the value in the **Parameter Value** column to open a page where you can change the value.

DISABLE_BROADCAST	1
-------------------	-------------------

- 5d In **New Value**, type 0, then click **Change** to save and apply the settings that disable the DISABLE_BROADCAST parameter, which enables broadcasting for NCP Server.

IMPORTANT: Messages are received only by logged-in users who are using Client for Open Enterprise Server versions that are capable of receiving broadcast messages, and that are configured to receive them.

DISABLE_BROADCAST	
Current Value	New Value
1	<input type="text" value="0"/> <input type="button" value="Change"/>
<input type="button" value="Back"/>	

A.4.3 Using the `ncpcon set` Command to Configure DST Parameters

- 1 Open a terminal console on the Linux server, then log in as the `root` user.
- 2 At the terminal console prompt, enter

```
ncpcon set parameter_name=value
```

Replace *parameter_name* and *value* with the settings you want to change.

IMPORTANT: Ensure that you enter the commands in lowercase.

For example, the following commands set the DST parameters to their default values.

```
ncpcon set duplicate_shadow_file_action=0
ncpcon set duplicate_shadow_file_broadcast=1
ncpcon set replicate_primary_tree_to_shadow=1
ncpcon set shift_modified_shadow_files=1
ncpcon set shift_accessed_shadow_files=0
ncpcon set shift_days_since_last_access=1
```

If the `DUPLICATE_SHADOW_FILE_BROADCAST` parameter is enabled, ensure that the `DISABLE_BROADCAST` parameter is disabled in order to allow broadcasting for NCP Server. For example, enter

```
ncpcon set disable_broadcast=0
```

A.5 DST Commands for `/etc/opt/novell/ncpserv.conf`

Use the commands in this section for the NCP Server configuration file (`/etc/opt/novell/ncpserv.conf`). The `ncpserv.conf` file is read only at eDirectory startup time. If you modify this file directly, you must restart `ndsd` in order for the changes to take effect.

SHADOW_VOLUME *volume_name shadow_area_path*

Identifies a volume as having a secondary storage area and specifies the path to that secondary volume. Any NCP volume can have a shadow. The root directory for the shadow area needs to already exist; the rest of the directories in the secondary file tree is automatically created as needed. The volume shadow area is available the next time the volume is mounted.

SHIFT_MODIFIED_SHADOW_FILES *value*

Enables a modified file to be moved from the secondary storage area to the primary storage area. The value can be either 0 (Disabled) or 1 (Allow). The default value is 1. When this parameter is on, and a file that is located in the secondary storage area is modified, the file is automatically moved back to the primary storage area when the file is closed.

SHIFT_ACCESSED_SHADOW_FILES *value*

Enables a file to be moved from the secondary storage area to the primary storage area if it is accessed as read-only a second time during a specified period of time. The value can be either 0 (Disabled) or 1 (Allow). The default value is 0. When this parameter is on, and a file that is

located in the shadow area is accessed, if this is the second access within the configured `SHIFT_DAYS_SINCE_LAST_ACCESS`, the file is automatically moved back to the primary area when the file is closed.

`SHIFT_DAYS_SINCE_LAST_ACCESS` value

Specifies the number of days to use when determining if a file should be moved back to the primary storage area. The value may be 0 (Disable), or between 1 and 365 (in days). The default is 1. When it is used with `SHIFT_ACCESSED_SHADOW_FILES`, the parameter sets the time when files are migrated back to the primary storage area after the second access within the specified elapsed time.

`DUPLICATE_SHADOW_FILE_ACTION` value

Controls how duplicate files conflicts are handled. The default is 0.

- 0 - Show duplicate shadow files (default)
- 1 - Hide duplicate shadow files
- 2 - Rename duplicate shadow files
- 3 - Delete duplicate files from shadow area
- 4 - Move duplicate shadow files to `/. _DUPLICATE_FILES`

`DUPLICATE_SHADOW_FILE_BROADCAST` value

Enables a message to be broadcast to an NCP user when a duplicate copy of a file is located on both the primary volume and the secondary volume. Valid settings are 0 (Disabled) and 1 (Allow). The default is Allow. The Client for Open Enterprise Server version in use must support receiving broadcast messages in order for the user to receive the message.

`REPLICATE_PRIMARY_TREE_TO_SHADOW` value

By default, the `REPLICATE_PRIMARY_TREE_TO_SHADOW` option is set to 1. Ensure this option is always set to 1, so when the existing NSS volumes are configured for DST, the primary volume directory tree structure will get replicated to the shadow volume directory tree.

The NSS filesystem is DST aware for new volumes, so the primary volume directory tree structure will always be replicated on the shadow volume directory tree. location.

A.6 DST Commands for `/etc/opt/novell/shadowfs.conf`

Use the commands in this section for the Shadow File System configuration file (`/etc/opt/novell/shadowfs.conf`).

`SHADOW` *root_path primary_area_path shadow_area_path*

Defines a shadow volume for ShadowFS. A shadow volume that is defined by the NCP engine is automatically mounted by ShadowFS and does not need to be defined in this configuration file.

`SHIFT_ON_MODIFY` value

Enables a modified file to be moved from the secondary storage area to the primary storage area. The value can be either 0 (Off) or 1 (On). The default value is 1. When this parameter is on, and a file that is located in the secondary storage area is modified, the file is automatically moved back to the primary area when the file is closed.

SHIFT_ON_ACCESS value

Enables a file to be moved from the secondary storage area to the primary storage area if it is accessed a second time during a specified time period. The value can be either 0 (Off) or 1 (On). The default value is 0. When this parameter is on, and a file that is located in the shadow area is accessed, if this is the second access within the configured `SHIFT_DAYS_SINCE_LAST_ACCESS`, the file is automatically moved back to the primary storage area when the file is closed.

SHIFT_DAYS_SINCE_LAST_ACCESS value

Specifies the number of days to use when determining if a file should be moved back to the primary storage area. The value may be 0 (Disable), or between 1 and 365 (in days). The default is 1. When it is used with `SHIFT_ON_ACCESS`, the parameter sets the time when files are migrated back to the primary storage area after the second access within the specified elapsed time.

A.7 DST EXCLUDE_VOLUME Command for /etc/opt/novell/ncp2nss.conf

Use the command in this section for the `/etc/opt/novell/ncp2nss.conf` file.

EXCLUDE_VOLUME nss_volumename

Prevents the named NSS volume from mounting in NCP Server. This command is added when you are using a specified NSS volume as the secondary storage area of a DST shadow volume.

An entry is automatically created in the `/etc/opt/novell/ncp2nss.conf` file by using OES Remote Manager for Linux to set the **Manage NCP Services > Manage Shares > NCP/NSS Bindings > NCP Accessible** option to No for a given NSS volume that you want to use as a secondary storage location in a DST shadow volume. For instructions, see [Section 10.5, “Configuring the NCP/NSS Bindings for an NSS Volume,” on page 97](#).

In a cluster, you must manually copy the line to the `/etc/opt/novell/ncp2nss.conf` file on each node.

A.8 DST Shadow Volume Information in /etc/NCPVolumes

The `/etc/NCPVolumes` file is an XML file that contains an entry for each mounted volume. It lists the volume’s name and the path for the volume’s primary file tree (`PRIMARY_ROOT`). If the volume is a shadow volume, it also shows the path for the secondary file tree (`SHADOW_ROOT`). Using this data file, a backup utility can easily locate each mounted NCP volume and find its primary and secondary file trees.

For example, the following XML entry defines the DST shadow volume named `VOL1`:

```
<VOLUME>
  <NAME>VOL1</NAME>
  <PRIMARY_ROOT>/media/nss/VOL1</PRIMARY_ROOT>
  <SHADOW_ROOT>/media/nss/ARCVOL</SHADOW_ROOT>
</VOLUME>
```

A.9 DST ShadowFS Volume Information in /etc/mtab.shadowfs

The `/etc/mtab.shadowfs` file is an XML file that contains an entry for each shadow volume mounted by ShadowFS. It lists the mount point, the path for the primary file tree, and the path for the secondary file tree.

For example, the following XML entry defines the DST shadow volume for ShadowFS named `VOL1`:

```
<SHADOWFS_MOUNTPOINTS>
  <MOUNTPOINT>
    <PATH>/media/shadowfs/VOL1</PATH>
    <PRIMARY_TREE>/media/nss/VOL1</PRIMARY_TREE>
    <SHADOW_TREE>/media/nss/ARCVOL</SHADOW_TREE>
  </MOUNTPOINT>
</SHADOWFS_MOUNTPOINTS>
```

A.10 Verifying and Syncing ACLs (Inherited Rights Filter (IRF) and Trustees) from DST Primary Volume to Shadow Volume

`python dst_verify_resync_trustee.py <OPTIONS>`

Verifies and syncs the ACLs only from the DST primary volume to shadow volume. Any difference in ACLs on DST volume pair is captured in the log file. This utility is located in the `/opt/novell/ncpserv/sbin` directory. It can be used with both eDirectory and AD trustees.

OPTIONS

-v, --verify

Verifies the difference in ACLs on the DST shadow volume compared to primary volume. If there is a difference in ACLs, the information is captured in the log file.

-s, --sync

Syncs the trustee rights from DST primary volume to shadow volume.

-n, --ncpcon

Obtains all DST volume pairs available on the server. This option is used only with the verify, sync or both verify and sync options.

--volumes=

Specify the DST primary volume name. For multiple volumes, specify the primary volume names with space separated. For example:

```
--volumes=VOL1 VOL2 VOL3
```

This option is used only with the verify, sync or both verify and sync options.

NOTE: If both `-n` and `--volumes` are used in a command, the `-n` option takes the priority.

-L, --LONG

Ignores case when comparing the file names on DST primary and shadow volumes.

--logfile=

Specify the path of log file along with the file name. If this option is not specified, the log file is generated in the current working directory.

-d, --debug

Prints the debug logs.

-h, --help

Displays the help information.

EXAMPLES

```
python dst_verify_resync_trustee.py --verify --logfile=var/log/dst.log  
--volumes=VOL1 VOL2 VOL3
```

Verifies the ACLs from DST primary volumes VOL1, VOL2, VOL3 to their corresponding shadow volumes and the difference in ACLs are captured in the `var/log/dst.log` file.

```
python dst_verify_resync_trustee.py --verify --sync --volumes=VOL1 VOL2  
VOL3
```

Verifies and syncs the ACLs from DST primary volumes VOL1, VOL2, VOL3 to their corresponding shadow volumes.

```
python dst_verify_resync_trustee.py --verify --sync --volumes=VOL1
```

Verifies and syncs the ACLs from DST primary volume VOL1 to its shadow volume.

In this example, the rights of user1 on the DST primary volume (VOL1) are RWFMA and the rights of user1 on shadow volume (VOL1_SHADOW) are RWF.

After successful execution of the command, the rights of user1 on the DST primary volume (VOL1) are RWFMA and the rights of user1 on shadow volume (VOL1_SHADOW) are RWFMA.

```
python dst_verify_resync_trustee.py --verify --sync --ncpcon
```

Verifies and syncs the ACLs on all DST volume pairs on the server.