

opentext™

# Open Enterprise Server 24.4 Unified Management Console (UMC)

October 2024

## **Legal Notice**

Copyright 2023 - 2024 Open Text

The only warranties for products and services of Open Text and its affiliates and licensors ("Open Text") 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. Open Text shall not be liable for technical or editorial errors or omissions contained herein. The information contained herein is subject to change without notice.

---

# Contents

<b>About This Document</b>	<b>11</b>
<b>Part I Overview</b>	<b>13</b>
<b>1 Overview of Unified Management Console</b>	<b>15</b>
<b>2 What's New or Changed in UMC</b>	<b>17</b>
What's New or Changed in UMC (OES 24.4).....	17
Enhanced UMC Health Script .....	17
Identity Console .....	17
Managing DFS Jobs .....	17
Managing Storage Management Services (SMS) .....	17
Managing Blocked user .....	17
<b>Part II Managing Clusters</b>	<b>19</b>
<b>3 Managing Clusters</b>	<b>21</b>
What tasks can be performed with clusters in UMC? .....	21
Are user-specific settings stored in UMC? .....	22
Is BCC managed through UMC?.....	22
How to access clusters? .....	22
How to list clusters? .....	23
How to access the cluster dashboard? .....	23
General .....	24
Priorities .....	24
Protocols .....	25
RME Groups .....	25
Graphical Representation .....	25
How to access cluster nodes? .....	25
What operations can be performed on cluster nodes? .....	26
How to access the node dashboard? .....	26
How to manage a cluster? .....	27
How to repair a cluster? .....	27
What actions can be performed on cluster resources? .....	28
How to create a cluster resource? .....	29
What are the configurable settings of the resource? .....	31
How to view Event logs? .....	32
How to view connected clients on a node? .....	33
What are the common tasks in each page? .....	33

<b>Part III Servers</b>	<b>35</b>
<b>4 Managing Server Settings</b>	<b>37</b>
How to view the details of all available servers? .....	37
<b>5 Log and Service Details</b>	<b>39</b>
Log Files .....	39
Service Status .....	39
<b>Part IV Storage</b>	<b>41</b>
<b>6 Managing NSS Pools</b>	<b>43</b>
What is a pool? .....	43
What are the prerequisites for creating a new pool? .....	43
How to create a new pool? .....	44
How to list pools? .....	46
How to view pool dashboard? .....	47
How to deactivate or activate pool for pool maintenance? .....	48
How to perform a pool move? .....	50
What happens when I delete a pool? .....	51
What are the prerequisites for deleting a pool? .....	51
How to delete pools? .....	51
How to rename a pool? .....	52
How to increase the size of a pool? .....	53
How to discard unused blocks in a pool? .....	55
Where are my deleted volumes? Can i restore/salvage them? .....	56
What are the prerequisites for AD users to access NSS data? .....	57
I am an AD user. How do I access NSS data? .....	58
The eDirectory pool object is corrupted. How to recover it? .....	61
<b>7 Managing Volume Jobs</b>	<b>63</b>
How to view list of DFS jobs? .....	63
What are the prerequisites to Move or Split jobs? .....	64
How to perform a Move job? .....	65
How to perform a Split job? .....	65
What happens when jobs are paused? .....	66
What happens when jobs are resumed? .....	66
How to handle files skipped by Move or Split Jobs? .....	66
How to cancel or delete jobs? .....	67
<b>8 Managing Pool Snapshots</b>	<b>69</b>
What is a pool snapshot? .....	69
What are the prerequisites for creating a pool snapshot? .....	69
How to create a pool snapshot? .....	69
How to list pool snapshots? .....	70

<b>9</b>	<b>Managing NSS Volumes</b>	<b>73</b>
	What is an NSS volume? . . . . .	73
	What are the features that can be enabled for a new volume? . . . . .	73
	What are the prerequisites for creating an encrypted volume with AES256? . . . . .	74
	How to create a new NSS volume? . . . . .	75
	How to list NSS volumes? . . . . .	77
	How to view volume dashboard? . . . . .	78
	How to deactivate and activate NSS volumes? . . . . .	78
	How to mount or dismount a volume? . . . . .	80
	How to rename a volume? . . . . .	81
	How to delete a volume? Can I restore or permanently delete it? . . . . .	82
	What is a volume object? . . . . .	83
	How to update volume objects? . . . . .	83
<b>10</b>	<b>Managing User Quota</b>	<b>85</b>
	What are user quotas? . . . . .	85
	How to add user quota? . . . . .	85
	How to list user quotas? . . . . .	86
	How to manage user quota? . . . . .	87
	How to delete user quotas? . . . . .	88
<b>11</b>	<b>Managing NSS Partitions</b>	<b>89</b>
	What is a partition? . . . . .	89
	How to list NSS partitions? . . . . .	89
	How to edit label of a partition? . . . . .	90
	How to list volumes in a partition? . . . . .	90
	What is NSS mirroring? . . . . .	91
	. . . . . How to	
	mirror a	
	partition? . . . . .	91
	How to	
	. . . . . delete	
	partitions? . . . . .	92
<b>12</b>	<b>Managing NSS Software RAID Devices</b>	<b>95</b>
	What is a software RAID? . . . . .	95
	What RAIDs do NSS support? . . . . .	95
	How to create a RAID device? . . . . .	96
	. . . . . How to list	
	the RAID	
	devices? . . . . .	98
	How to view RAID device dashboard? . . . . .	98
	How to rename a RAID device? . . . . .	99
	How to increase the size of a RAID device? . . . . .	100
	What happens when i delete a software RAID device? . . . . .	102
	What happens when i delete a RAID1 device? . . . . .	102

How to delete a software RAID device? . . . . .	102
What is disk mirroring or remirroring? . . . . .	103
How to mirror or remirror a RAID 1 device? . . . . .	103
How to deactivate or activate a RAID device? . . . . .	104
<b>13 Managing Devices</b>	<b>107</b>
What is a device? . . . . .	107
How to list devices connected to the servers? . . . . .	107
What happens when a device is initialized? . . . . .	108
What happens when a device is shared? . . . . .	108
How can i initialize a device connected to a server? . . . . .	108
Why I need to reinitialize a device? . . . . .	109
How can i reinitialize a device? . . . . .	109
How to share or unshare an initialized device? . . . . .	110
<b>Part V Files and Folders</b>	<b>113</b>
<b>14 Managing Files and Folders</b>	<b>115</b>
How to view files and folders? . . . . .	115
How to create a new folder? . . . . .	115
How to modify properties of a volume, file, or folder? . . . . .	116
Details tab . . . . .	116
Trustees tab . . . . .	118
How to modify directory quota of a volume or folder? . . . . .	118
How to modify owner of a volume, file, or folder? . . . . .	118
How to modify attributes of a volume, file, or folder? . . . . .	119
How to view deleted files and folders? . . . . .	120
How to delete files and folders? . . . . .	120
How to salvage the deleted files and folders? . . . . .	120
How to purge the files and folders? . . . . .	121
How to rename a file or folder? . . . . .	121
How to move files and folders in a volume? . . . . .	121
How to resolve file move conflicts? . . . . .	122
<b>15 Managing Rights</b>	<b>123</b>
How to add trustee(s) for a volume, file, or folder? . . . . .	123
How to modify trustee rights for users and groups? . . . . .	124
How to view trustee rights of a volume, file, or folder? . . . . .	124
How to enable all rights for users and groups? . . . . .	124
How to disable all rights for users and groups? . . . . .	125
What are the various trustee rights? . . . . .	125
What are effective rights? . . . . .	126
How to view effective rights of users and groups? . . . . .	126
What are inherited rights? . . . . .	127
How to view inherited rights of a user or group? . . . . .	127
How to use inherited rights filter? . . . . .	127
How to copy or replicate rights of a user or group to other users and groups in the context tree? . . . . .	128

How to remove trustees for a selected path? .....	128
<b>Part VI Storage Technology</b>	<b>129</b>
<b>16 Managing Replica Sites</b>	<b>131</b>
Naming Convention Changes .....	131
How to list replica sites? .....	131
Where to view details of a replica site? .....	133
How to create a Management Context? .....	133
How to add a Replica site? .....	134
How to repair the DFS Replica service? .....	134
How to configure the DFS Replica service? .....	135
How to delete a Replica Site? .....	135
What happens when a Replica Site is paused or stopped? .....	135
<b>17 Managing Junctions</b>	<b>137</b>
What are the guidelines for creating or managing junctions? .....	137
How to create a junction? .....	137
Where to view junctions? .....	138
DFS > Junctions .....	139
Files & Folders .....	139
How to configure junctions? .....	140
How to delete junctions? .....	140
How to synchronize rights between the Source and Target locations? .....	140
<b>Part VII Service configuration</b>	<b>141</b>
<b>18 Managing NCP</b>	<b>143</b>
How to manage the local code page? .....	143
How to manage NCP server caching settings? .....	143
How to manage encryption and MFA on an NCP server? .....	144
How to manage NCP server setting locks? .....	144
How to manage NCP server communications settings? .....	145
How to manage NCP server volumes settings? .....	145
How to manage NCP server logging settings? .....	146
How to manage NCP server performance tuning settings? .....	147
How to manage NCP server user ID updates? .....	147
<b>19 Managing SMDR</b>	<b>149</b>
How to configure SMDR? .....	149
<b>20 Managing TSAFS</b>	<b>151</b>
How to configure TSAFS? .....	151

**21 Managing NCP Shares 155**

What is an NCP share and how to manage it? . . . . .	155
How to list NCP shares? . . . . .	156
How to verify trustees for an NCP share? (OES 23.4). . . . .	156
How to verify the rights of an NCP share? . . . . .	157
How to resync trustees an NCP share? (OES 23.4). . . . .	157
How to resync the rights of an NCP share? . . . . .	157
How to enable or disable encryption on an NCP share? . . . . .	158
How to enable or disable MFA on an NCP share? . . . . .	158
What are accessed files and how to view them? (OES 23.4). . . . .	158
What are open files and how to view them? . . . . .	159
What are the prerequisites for adding a secondary volume? . . . . .	159
How to add secondary volume? . . . . .	159
How to view secondary volume? . . . . .	160
How to remove secondary volume? . . . . .	160
How to manage security for sub-folders on an NCP share? (OES 23.4) . . . . .	161
How to manage sub-folder security on an NCP share? . . . . .	162
How to enable or disable write permission for an NCP share? . . . . .	162
How to activate or deactivate an NCP share? . . . . .	162

**22 Managing NCP Connections 165**

How to view NCP connections? . . . . .	165
What actions can be performed on NCP connections? . . . . .	165
How to send a broadcast message to all NCP connections? . . . . .	166
How to clear unauthenticated NCP connection? . . . . .	166
How to view open files, NCP shares, and details of an NCP connection? . . . . .	166
How to send message to an NCP connection? . . . . .	167
How to clear an NCP connection? . . . . .	167

**23 Managing CIFS Shares 169**

How to create a new CIFS share? . . . . .	169
How to list CIFS shares? . . . . .	170
How to remove a CIFS share? . . . . .	170
What is encryption on a CIFS share? . . . . .	170
How to manage encryption on a CIFS share? . . . . .	170
Enable encryption while creating a new share . . . . .	171
Enable encryption on an existing share . . . . .	171
What is folder redirection on a CIFS share? . . . . .	172
What is Mac backup on a CIFS share? . . . . .	172
What is the character limit for CIFS share name and comment box? . . . . .	172
How to filter the CIFS shares? . . . . .	172
How to manage folder redirection on a CIFS share? . . . . .	173
How to manage Mac backup on a CIFS share? . . . . .	173
What are the various rights and how to manage it on CIFS shares? . . . . .	174
How to add trustees for a CIFS share? . . . . .	174

What is the CIFS share limitation that a server can host? . . . . .	175
How to modify an existing CIFS share? . . . . .	175
What are open files in a CIFS share? . . . . .	176
How to view the open files in a CIFS share? . . . . .	176
How to close open files of CIFS shares? . . . . .	176
Close all open files . . . . .	176
Close individual open file . . . . .	176
What are the various access modes for open files? . . . . .	177
<b>24 Managing CIFS Connections</b>	<b>179</b>
How to list and view the information related to CIFS connections? . . . . .	179
How to view the open files of a CIFS connection? . . . . .	180
How to view the shares associated with a CIFS connection? . . . . .	180
How to view the security equivalence of a CIFS connection? . . . . .	181
<b>25 Managing Invalid Users</b>	<b>183</b>
How to list invalid users and blocked users? . . . . .	183
Who is an invalid user? . . . . .	183
Who is a blocked user? . . . . .	184
How to add a user to the blocked user list? . . . . .	184
How to unblock an invalid user? . . . . .	184
How to unblock a blocked user? . . . . .	184
How to change an invalid user to a blocked user? . . . . .	184
<b>26 Managing User Context</b>	<b>187</b>
How to list the user contexts? . . . . .	187
How to add a user context? . . . . .	187
How to remove a user context? . . . . .	187
<b>Part IX Reports</b>	<b>189</b>
<b>27 Cluster Reports</b>	<b>191</b>
How to generate a cluster report? . . . . .	191
How to view reports? . . . . .	191
Report Failures . . . . .	192
<b>Part X Troubleshooting</b>	<b>193</b>
<b>28 Troubleshooting</b>	<b>195</b>
Known Issues . . . . .	195
UMC Health Script . . . . .	196
Autofix . . . . .	197
Missing Node Modules . . . . .	198
Unable to Connect to the Database . . . . .	198
Warning: Entered Hostname is Incorrect . . . . .	199

Volume Issues .....	199
Failing to List Pools or Volumes .....	199
Unable to Perform Storage Operations as an Admin Equivalent User .....	199
Creating Volume with AES 256 Encryption Failed .....	199
Login Failures .....	199
Action to Perform in case of Cache-Related Issues .....	200
Cluster Issues .....	200
Renaming Cluster Pool or Volume Failed .....	200
Status of Healthy Cluster is Down or Unknown .....	200

# About This Document

This document provides frequently asked questions on the tasks performed through the Unified Management Console (UMC) application.

- ♦ [Part I, “Overview,” on page 13](#)
- ♦ [Part II, “Managing Clusters,” on page 19](#)
- ♦ [Part III, “Servers,” on page 35](#)
- ♦ [Part IV, “Storage,” on page 41](#)
- ♦ [Part V, “Files and Folders,” on page 113](#)
- ♦ [Part VI, “Storage Technology,” on page 129](#)
- ♦ [Part VII, “Service configuration,” on page 141](#)
- ♦ [Part VIII, “File Access Protocols,” on page 153](#)
- ♦ [Part IX, “Reports,” on page 189](#)
- ♦ [Part X, “Troubleshooting,” on page 193](#)

## Audience

This document is intended for UMC administrators.

## Feedback

We want to hear your comments and suggestions about this manual and the other documentation included with this product. Please use the [comment on this topic](#) link at the bottom of each page of the online documentation.

## Additional Documentation

For documentation on OES guides, see the [OES 24.4 Documentation web site](#).



# Overview

- ♦ [Chapter 1, “Overview of Unified Management Console,” on page 15](#)
- ♦ [Chapter 2, “What’s New or Changed in UMC,” on page 17](#)



# 1 Overview of Unified Management Console

Open Enterprise Server (OES) 23.4 is built on SLES and ships with eDirectory. Various consoles and command-line tools are used to manage OES services. Unified Management Console (UMC) is installed and configured through YaST.

UMC is a highly responsive, simple, and secure web-based management console for managing small and large OES deployments. UMC provides customized access to network administration utilities and content from virtually anywhere using the Internet and a web browser similar to iManager. UMC provides single point of administration for OES resources.

Because UMC is a Web-based tool, it has several advantages over client-based administrative tools:

- ◆ Changes to UMC look, feel, and functionality are immediately available to all administrative users.
- ◆ No need to open additional administrative ports for remote access. UMC leverages standard HTTPS ports (443).
- ◆ Not necessary to download and maintain an administrative client.



# 2 What's New or Changed in UMC

This section describes enhancements and changes in Unified Management Console.

- ♦ [“What's New or Changed in UMC \(OES 24.4\)” on page 17](#)

## What's New or Changed in UMC (OES 24.4)

### Enhanced UMC Health Script

The `umcServiceHealth` script has been enhanced to check the health of Redis.

For more information, see [“UMC Health Script” on page 196](#).

### Identity Console

Identity console is bundled with UMC for identity management in OES environment. The packages are installed automatically during the UMC installation and no separate installation is required.

### Managing DFS Jobs

- ♦ [List jobs](#)
- ♦ [Move job](#)
- ♦ [Split job](#)
- ♦ Reschedule job
- ♦ [Pause and Resume job](#)
- ♦ [View skipped files](#)
- ♦ [Cancel job](#)

For more information, see [Chapter 7, “Managing Volume Jobs,” on page 63](#).

### Managing Storage Management Services (SMS)

Support for managing SMS components Storage Management Data Requester (SMDR) and Target Service Agent for File System (TSAFS).

For more information, see [Chapter 19, “Managing SMDR,” on page 149](#) and [Chapter 20, “Managing TSAFS,” on page 151](#).

### Managing Blocked user

In UMC CIFS, **Permanent Invalid User** is renamed as **Blocked User**. There is no other functionality change.

For more information, see [Chapter 25, “Managing Invalid Users,”](#) on page 183.



# Managing Clusters

- ◆ [“What tasks can be performed with clusters in UMC?” on page 21](#)
- ◆ [“How to access the cluster dashboard?” on page 23](#)
- ◆ [“How to access the node dashboard?” on page 26](#)
- ◆ [“How to manage a cluster?” on page 27](#)
- ◆ [“What actions can be performed on cluster resources?” on page 28](#)
- ◆ [“How to view Event logs?” on page 32](#)



# 3 Managing Clusters

This chapter describes the procedures for managing clusters. For configuring clusters, refer to the [OES Cluster Services for Linux Administration Guide](#).

- ♦ “What tasks can be performed with clusters in UMC?” on page 21
- ♦ “Are user-specific settings stored in UMC?” on page 22
- ♦ “Is BCC managed through UMC?” on page 22
- ♦ “How to access clusters?” on page 22
- ♦ “How to list clusters?” on page 23
- ♦ “How to access the cluster dashboard?” on page 23
- ♦ “How to access cluster nodes?” on page 25
- ♦ “What operations can be performed on cluster nodes?” on page 26
- ♦ “How to access the node dashboard?” on page 26
- ♦ “How to manage a cluster?” on page 27
- ♦ “How to repair a cluster?” on page 27
- ♦ “What actions can be performed on cluster resources?” on page 28
- ♦ “How to create a cluster resource?” on page 29
- ♦ “What are the configurable settings of the resource?” on page 31
- ♦ “How to view Event logs?” on page 32
- ♦ “How to view connected clients on a node?” on page 33
- ♦ “What are the common tasks in each page?” on page 33

## What tasks can be performed with clusters in UMC?

The following tasks are available for managing clusters:

- ♦ Create resource.
- ♦ Configure resource.
- ♦ Cluster reports.
- ♦ Cluster event logs.
- ♦ Configure and repair clusters.
- ♦ Cluster dashboard page offers two views:
  - ♦ Graphical representation of the selected cluster.
  - ♦ Full-page view of nodes and resources.
- ♦ Graphical representation of nodes and quorum compliance.
- ♦ Add and remove nodes from favorites, and restart nodes.

- ◆ Node dashboard page displays:
  - ◆ **Dashboard View:** Graphical representation of server statistics.
  - ◆ **Table View:** Lists NCP and CIFS connections.
- ◆ Add and remove resources from favorites.
- ◆ List clusters.
- ◆ Cluster dashboard displays graphical representation of the clusters.
- ◆ List and node shutdowns.
- ◆ List resources, which includes actions such as bring the resources online, take them offline, and migrate resources.

## Are user-specific settings stored in UMC?

Yes, these settings are stored in the PostgreSQL database and are user-specific and persistent across logins, browsers and devices. This applies to the primary filter settings.

For example, if you have selected two clusters for management with specific columns to be displayed, the user-specific settings are available during subsequent logins.

## Is BCC managed through UMC?

BCC will be supported in upcoming releases. You can continue to manage BCC through iManager.

## How to access clusters?

- 1 Log in to UMC with your admin credentials.
- 2 Click **Clusters**.

During the initial login, the cluster listing page is empty. However, as you browse, only cluster objects are listed due to enhanced and context aware filter functionality. The selected clusters are listed on the **Clusters** page.

**Figure 3-1** Cluster Listing

Total: 10 item(s)								
<input type="checkbox"/>	Status	Name	Master node	Master IP address	Location	Node availability	Resources	Epoch
<input type="checkbox"/>	●	Aadhaar	private-12-3	192.168.12.5	Delhi	<div style="width: 100%; background-color: green;"></div> 2/2	1	5
<input type="checkbox"/>	●	MyGov	private-14-3	192.168.14.5	Jaipur	<div style="width: 100%; background-color: green;"></div> 2/2	1	5
<input type="checkbox"/>	●	GSTN	private-16-3	192.168.16.5	Kochi	<div style="width: 100%; background-color: green;"></div> 2/2	1	5
<input type="checkbox"/>	●	eNAM	private-18-3	192.168.18.5	Kolkata	<div style="width: 100%; background-color: green;"></div> 2/2	1	5
<input type="checkbox"/>	●	DigLocker	private-8-6	192.168.8.8	Chennai	<div style="width: 60%; background-color: orange;"></div> 3/5	1	19
<input type="checkbox"/>	●	GeM	private-20-3	192.168.20.5	Madurai	<div style="width: 100%; background-color: green;"></div> 2/2	1	5
<input type="checkbox"/>	●	PassportSeva	private-22-3	192.168.22.5	Mumbai	<div style="width: 100%; background-color: green;"></div> 2/2	1	5
<input type="checkbox"/>	●	Parivahan	private-24-3	192.168.24.5	Pune	<div style="width: 50%; background-color: orange;"></div> 1/2	1	0
<input type="checkbox"/>	●	UMANG	private-10-4	192.168.10.6	Hyderabad	<div style="width: 100%; background-color: green;"></div> 3/3	1	8
<input type="checkbox"/>	●	GSTSeva	private-26-3	192.168.26.5	Surat	<div style="width: 100%; background-color: green;"></div> 2/2	1	5

# How to list clusters?

Log in to UMC with your admin credentials, and then follow these steps:

- 1 Browse and select the cluster objects you want to view.
- 2 The following information is displayed for each cluster object.

Column Name	Description
Status (Colour Coding)	Status
Green	<b>Running:</b> The cluster is up and running.
Blue	<b>Maintenance:</b> The cluster is temporarily suspended by the admin for maintenance.
Gray	<b>Down:</b> The cluster is stopped, and admin intervention is required.
Red	<b>Failed:</b> One or more nodes in the cluster have failed, and admin intervention is required.
White	<b>Unknown:</b> UMC cannot determine the status of the cluster.
Name	The name assigned to the cluster.
Master node	The name of the currently assigned master node for the cluster.
Node availability	Number of available nodes out of the total nodes.
Resources	The number of resources running in this cluster.
Epoch	The number of times the cluster state has changed. The cluster state changes every time a server joins or leaves the cluster.

The columns listed above are the default ones. You can select  to add additional columns such as **Type**, **Master IP address**, and **Location**.

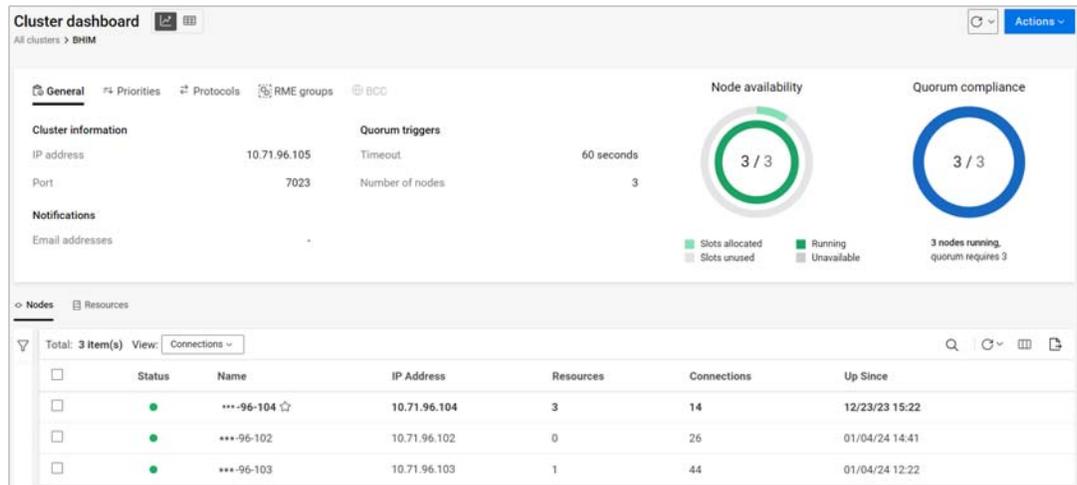
- 3 Select a refresh frequency that allows you to comfortably view all items in the list.

**NOTE:** If the status of a healthy cluster is Down or Unknown, then increase the timeout value `CLUSTER_LISTING_FAILURE_TIMEOUT = 2000` in the `/opt/novell/umc/apps/umc-server/prod.env` file. The default value is 2000 ms, and due to network latency, it might not be able to retrieve the correct status of the cluster. Additionally, if this parameter is missing in the `prod.env` file, ensure to add it so that the cluster listing timeout occurs after the specified time.

# How to access the cluster dashboard?

The cluster dashboard provides a graphical representation of a cluster. To view details:

- 1 Browse and select the cluster objects you want to view.
- 2 Select a cluster, then select **Dashboard**.
- 3 The cluster dashboard offers two views:
  - ◆ **Dashboard view** : Displays the Dashboard, Nodes and Resources.



- ◆ **Table view** : Displays a comprehensive view of Nodes and Resources, which is useful when dealing with a long list of nodes and resources.



The cluster dashboard displays the following information:

- ◆ “General” on page 24
- ◆ “Priorities” on page 24
- ◆ “Protocols” on page 25
- ◆ “RME Groups” on page 25
- ◆ “Graphical Representation” on page 25

## General

- ◆ **Cluster Information:** Displays the IP address bound to the master node and remains associated with the master node, regardless of any server changes. The default cluster port number is 7023.
- ◆ **Quorum Triggers:** Displays the number of nodes required in the quorum and the time the cluster should wait before ignoring the quorum.
- ◆ **Notifications:** Email messages are sent for specific cluster events, such as changes in cluster and resource state or nodes joining or leaving the cluster.

## Priorities

Displays the load priorities of individual cluster resources on a node during the cluster startup, failover, or failback. Resource priority determines the order in which resources load.

## Protocols

Displays details on transmit frequency and tolerance settings for all nodes in the cluster, including the master node. The master node is typically the first node brought online, but in case of failure, any other node can become the master. For more information, see the [Configuring Cluster Protocols](#) in the [OES Cluster Services for Linux Administration Guide](#).

## RME Groups

Displays a combination of resources available to the cluster. Resources within the same group cannot run concurrently on a node, and a resource can belong to multiple groups. There are four fixed RME groups (Group A, Group B, Group C, and Group D), and their names cannot be customized.

## Graphical Representation

On the right side of the dashboard there is a graphical representation of the cluster nodes.

- ♦ **Node availability:** The outer graph indicates the number of nodes in use out of a total of 32, while the inner graph displays the available nodes and unhealthy nodes.
- ♦ **Quorum compliance:** Displays the number of nodes required for the quorum to be met and number of nodes running.



## How to access cluster nodes?

- 1 Select a cluster, then select **Dashboard**.
- 2 The **Nodes** tab displays all nodes for the selected cluster. You can view details in two different modes: **Connections** and **Performance**.
  - ♦ **Connections:** This is the default view, displaying a list of connections with other common columns.
  - ♦ **Performance:** Displays CPU utilization and core information, in addition to the other common columns.
- 3 In the **Name** column, the master node is identified by a star symbol at the end of its name. The following status is displayed for each node:

Color	State	Description
Green	Running (LIVE)	The node is running.
White	Non-member (LEFT)	The node is no longer part of the cluster. The cluster migrates any resources running on this node to another eligible live node before the node leaves the cluster.
Red	Unavailable (DEAD)	The node is not running properly and requires admin intervention.
White with red ring	Failed to start (GASP)	The node is waiting for the quorum to be established so that it can begin loading.
Gray	Banned (PILL)	The cluster has intentionally triggered an immediate node shutdown.

- 4 The operations that can be performed on the nodes are shutdown, restart, add a node to favorites, and dashboard.

Select a refresh frequency long enough to allow the task to be completed.

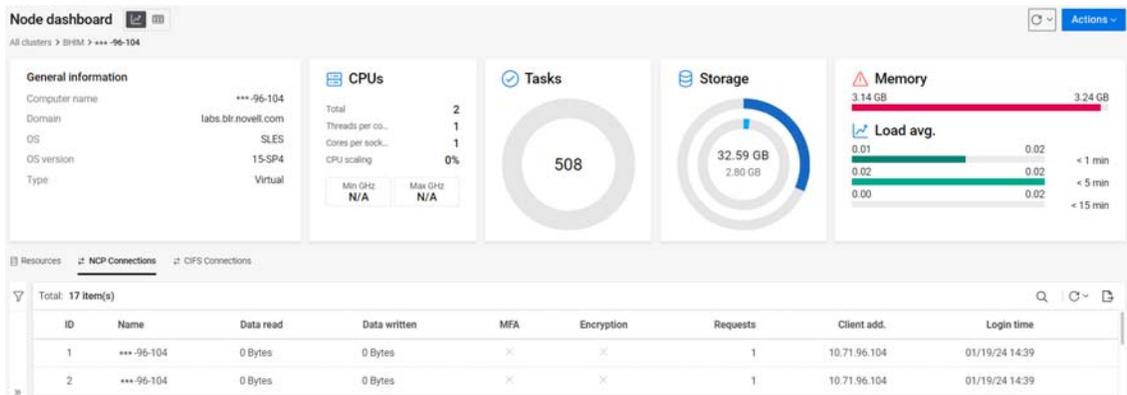
## What operations can be performed on cluster nodes?

The following operations can be performed on nodes:

- ◆ To shut down a node, select the desired node and select **Shut Down** from the menu. This action brings down the selected node, making it no longer available to the clients.
- ◆ To restart a node, select the desired node and select **Restart** from the menu. If the resource running a service was only on this node due to an RME group or preferred node setting, that service becomes unavailable.
- ◆ To add a node to favorites, select the desired node and select **Add to my nodes**. To view these nodes, select **Show my nodes only** in the **Advanced Filters**.
- ◆ To access the node dashboard, select the desired node and select **Dashboard**.

## How to access the node dashboard?

- 1 Select a cluster, then select **Dashboard**.
- 2 The **Nodes** tab displays all nodes for the selected cluster.
- 3 Select a node, then select **Dashboard**.



The node dashboard displays server statistics such as general information, CPU utilization, tasks, storage, and memory details.

The **Actions** menu provides options to shut down or restart the selected node.

**NOTE:** For a virtual machine, the minimum and maximum values of the CPU are displayed as N/A.

## How to manage a cluster?

- 1 Select a cluster, then select **Configure**.

Alternatively, you can access this option from the dashboard by clicking **Actions > Configure**.

- 2 A configuration wizard is available to modify the required configuration settings.

**2a Configuration:** This is a view-only page displaying details of the Master IP address and port.

**2b Policies:** View or modify **Quorum triggers**, **Notifications**, and **Log level** details.

**2c Priorities:** Choose one of these methods to change the load order (from highest priority to lowest priority) of a resource compared to other cluster resources on the same node:

- ♦ **Arrows:** Click the up-arrow or down-arrow adjacent to each resource.
- ♦ **Drag:** Drag the resource to modify the load order.

**2d Protocols:** View or modify protocol settings such as **Heartbeat**, **Master watchdog**, **Maximum re-transmits**, **Tolerance**, and **Slave watchdog**.

**2e RME groups:** Select the resources that must not be assigned to the same node simultaneously.

**2f Summary:** Displays a summary of the modified configuration. Review it and click **Finish**.

The dashboard refreshes, and the updated data is displayed.

## How to repair a cluster?

Perform a repair when there might be a discrepancy of resources between the cluster and eDirectory.

- 1 Select a cluster, then select **Repair**. This action triggers a restart of the cluster, which might modify the node ids.

After a successful repair, the additional resources that are not part of eDirectory are removed from the cluster.

- 2 To verify, view **Resources** in the dashboard. After a successful repair, the additional resources are cleaned up from the list of resources.

## What actions can be performed on cluster resources?

- 1 Select a cluster, then select **Dashboard**.
- 2 On the dashboard, navigate to the **Resources** tab, you can perform the following tasks: create resource, add to favorites, configure a resource, bring online, take offline, and migrate.
  - ♦ To create resource, click **Create resource**. Alternatively, you can access this option from the dashboard by clicking **Actions > Create resource**.
  - ♦ To add a resource to favorites, select the desired resource and select **Add to my resources**. To view these resources, select **Show my resources only** in the **Advanced Filters**.
  - ♦ To configure a resource, select the desired resource and select **Configure**.
  - ♦ To bring a resource online, select the desired resource and select **Take online**. This action runs the load script, loading the resource on its primary preferred node or on an alternate preferred node.
  - ♦ To take a resource offline, select the desired resource and select **Take offline**. This action runs the unload script, removing the resource from the server. The resource cannot be loaded on any other server in the cluster and remains unloaded until you load it again.
  - ♦ To migrate a resource, select the desired resource and select **Migrate**. When a resource is migrated, it moves from the node where it is currently running to another node. You can select node from the Preferred Nodes list or other unassigned nodes.
- 3 In the list of resources, the master resource (MASTER\_IP\_ADDRESS\_RESOURCE) cannot be selected, as no actions can be performed on it.

The following status is displayed for each resource.

Color	State	Description
Green	Online	The resource is online.
Orange	Alert	The resource is waiting for the admin to take an action, such as starting, failing over, or failing back the resource on the specified server.
Red	Comatose	The resource is not running properly and requires admin intervention.
White with red ring	Quorum Wait	The resource is waiting for the quorum to be established so it can begin loading.
Blue	Loading	The resource is being loaded on the server.
White with blue ring	Unloading	The resource is being unloaded from the server it was running on.
Gray	Offline	The resource is shut down or is in a dormant or inactive state.
White	Unassigned	No node has been assigned for loading the resource.

Color	State	Description
	NDS Sync	The properties of the resource have changed and the changes are still being synchronized in the eDirectory.

If any resources are in an intermediate state, such as loading or unloading, click **Refresh** to get the updated status of the resources, or adjust the refresh frequency to be long enough to allow the task to be completed.

## How to create a cluster resource?

Cluster resources should be created for every shared file system or any server-based applications or services you want to make available to users at all times.

- 1 Select a cluster, then select **Dashboard**.
- 2 Navigate to the **Resources** tab and click **Create resource**.

Alternatively, you can access this option from the dashboard by clicking **Actions > Create resource**.

---

**NOTE:** A pool resource is automatically created when an NSS pool is created.

---

- 3 A wizard is displayed to create a new resource.

### 3a Configuration

**3a1** Specify the name of the resource you want to create.

**3a2** In **Type**, select one of the available templates. The cluster resource templates can be used on physical servers, virtualization host servers, and virtual machine (VM) guest servers.

Cluster Resource Template	Use
Generic	An empty template.
Generic_IP_Service	This template is auto-populated with commands or variables and is used to create cluster resources for a certain server applications that run on your cluster.
Generic_FS	This template is auto-populated with commands or variables and is used to configure resource for Linux Logical Volume Manager (LVM).
DNS	This template is auto-populated with commands or variables and is used to configure resource for the DNS service.
DHCP	This template is auto-populated with commands or variables and used to configure resource for the DHCP service.

**3a3** If you want the resource to be immediately available after creation, enable **Initialize after creation**.

**3a4** Click **Next**.

### 3b Policies

**3b1** If you want to ensure that the resource runs only on the master node in the cluster, select **Resource follows master**.

If the master node in the cluster fails, the resource fails over to whichever node becomes the master.

**3b2** If you don't want the cluster-wide timeout period and node number limit enforced, select **Ignore quorum**.

This ensures that the resource is launched immediately on any server in the Preferred Nodes list as soon as any server in the list is brought online.

**3b3** Specify the **Failover mode**. When the mode is enabled, the resource automatically starts on the next server in the Preferred Nodes list if there is a hardware or software failure. If the mode is disabled, you can intervene after a failure occurs and before the resource is moved to another node.

**3b4** Specify the **Start mode**. When the mode is enabled, the resource automatically starts on a server when the cluster is first brought up. If the mode is disabled, you can manually start the resource on a server when you want, instead of having it automatically start when servers in the cluster are brought up.

**3b5** Specify the **Failback mode**. When the mode is set to **Disabled**, the resource does not fall back to its most preferred node when the most preferred node rejoins the cluster. If the mode is set to **Auto**, the resource automatically falls back to its most preferred node when the most preferred node rejoins the cluster. Set the mode to **Manual** to prevent the resource from moving back to its preferred node when that node is brought back online, until you are ready to allow it to happen.

**3b6** Click **Next**.

**3c Assigned nodes:** Allows you to assign nodes to use for the resource. You also sequence the list of nodes to specify the preferred order that the nodes will be tried when a resource is brought online after its current node fails.

**3c1** In the **Unassigned** area, select a node that the resource can use, and then click the arrow button to move the selected node to the **Assigned** nodes area.

Repeat this step for all of the cluster nodes you want to assign to the resource.

**3c2** In the **Assigned** area, select a node that you want to unassign from the resource, then click the arrow button to move the selected node to the **Unassigned** nodes area.

**3c3** Click **Next**.

**3d Scripts:** You can add an unload script to specify how the application or resource should terminate. Resource monitoring allows cluster to detect a the resource failure independently of its ability to detect node failures.

**3d1** A load script is required for each resource, service, disk, or pool in your cluster. The load script specifies the commands to start the resource or service on a server.

**3d1a** Edit or add the necessary commands to the script to load the resource on a server.

**3d1b** Specify the **Timeout** value. The timeout value determines how much time the script is given to complete. If the script does not complete within the specified time, the resource becomes comatose. The timeout value is applied only when the resource is migrated to another node. It is not used during resource online/offline procedures.

**3d2** An unload script is not required by all resources, but is required for cluster-enabled Linux partitions. You can add an unload script to specify how the application or resource should terminate. Programs should be unloaded in the reverse order of how they were loaded. This ensures that supporting programs are not unloaded before programs that rely on them in order to function properly.

**3d2a** Edit or add the necessary commands to the script to unload the resource on a server.

**3d2b** Specify the **Timeout** value. The timeout value determines how much time the script is given to complete. If the script does not complete within the specified time, the resource becomes comatose. The timeout value is applied only when the resource is migrated to another node. It is not used during resource online/offline procedures.

**3d3** The monitor script is used to monitor the status of service or storage objects.

**3d3a** Edit or add the necessary commands to the script to monitor the resource on the server.

**3d3b** Specify the **Timeout** value. The timeout value determines how much time the script is given to complete. If the script does not complete within the specified time, the failure action you have chosen gets initiated.

**3d4** Click **Next**.

**3e Monitoring:** Allows you to monitor the health of the specified resource by using a script that you create or customize. By default, resource monitoring is disabled. To enable or modify the settings, you must configure the resource.

**3f Summary:** Displays a summary of the resource. Review it and click **Finish**.

## What are the configurable settings of the resource?

- ♦ **Configuration:** Displays name and type of a resource and these fields are not editable. If the resource is pool resource, then additional fields like IP address and advertising protocols is displayed and can be modified.

**IP address change:** When the IP address of a pool cluster resource is modified and saved, the load, unload, and monitor scripts are automatically updated with the new IP address. It also automatically updates the resource's IP address that is stored in its NCP Virtual Server object.

- ♦ **Policies:** View or modify the default policies set for the cluster resource.
- ♦ **Assigned Nodes:** View or modify the preferred nodes used for the cluster resource.
- ♦ **Scripts:** View or modify the load, unload, and monitor scripts for the cluster resource.
- ♦ **Monitoring:** Allows you to monitor the health of the specified resource by using a script that you create or customize. On enabling resource monitoring, you need to set the interval to poll the resource's health, and the action to take if the resource fails to load on the maximum number of local restarts.
  - ♦ In the **Polling interval**, specify how often you want the resource monitor script for this resource to run.
  - ♦ The **Failure rate** is the maximum number of failures (**Maximum local failures**) detected by the monitor script during a specified amount of time (**Time interval**).

A failure action is initiated when the resource monitor detects that the resource fails more times than the maximum number of local failures allowed to occur during the specified time interval. For failures that occur before it exceeds the maximum, cluster services automatically attempts to unload and load the resource.

- ◆ The **Failure action** indicates what action to take on the resource when a failure occurs.
  - ◆ **Set resource as comatose:** (Default) The resource is placed in a comatose state when the failure action initiates. Administrator action is required to take the resource offline, resolve the issue, and bring it online again on the same or different node.
  - ◆ **Migrate the resource based on the preferred nodes list:** Each time a failure action triggers a failover, the resource migrates to a different node, according to the order in its Preferred Nodes list and availability of the nodes. The resource is not automatically failed back to the original node. Administrator action is required to cluster migrate the resource to the node, as desired.
  - ◆ **Reboot the hosting node without syncing or unmounting disks:** If the failure action initiates, all the resources on the hosting node will fail over to the next available node in its Preferred Nodes list because of the reboot. This is a hard reboot, not a graceful one. The reboot option is normally used only for a mission-critical cluster resource that must remain available. The resources are not automatically failed back to the original node. Administrator action is required to cluster migrate them back to the node, as desired.
- ◆ **Summary:** Displays a summary of the modified settings of the resource.

To configure the resource, do the following:

- 1 Select a cluster, then select **Dashboard**.
- 2 On the dashboard, navigate to the **Resources** tab.
- 3 Select a resource and click **Configure**. A configuration wizard is available to modify the resource settings.

## How to view Event logs?

The Event logs displays the events logged by the cluster. Events can be node specific or they can be resource specific.

You can use the **Advanced Filters** to filter out events according to the following categories:

- ◆ Severity (Error, Warning, Information)
- ◆ Event types (Failed, Comatose, Quorum Wait, Running)
- ◆ Node (by node name)
- ◆ Resource (by resource name)
- ◆ Timestamp (by specified time range)

To view the event logs:

- 1 Select a cluster, then select **Dashboard**.
- 2 Click **Actions > View event logs**.
- 3 The cluster events are displayed. Using **Advanced Filter**, you can filter the logs and save the logs to a .csv file.

# How to view connected clients on a node?

The clients are connected through either NCP or CIFS to a node.

To view NCP or CIFS connections:

- 1 Select a cluster, then select **Dashboard**.
- 2 Navigate to the **Nodes** tab, which displays all nodes for the selected cluster. Select a node, then select **Dashboard**.
- 3 Click on the **NCP Connections** or **CIFS Connections** tab to view details such as data read or written, encryption status, the number of requests from that connection, and so on.

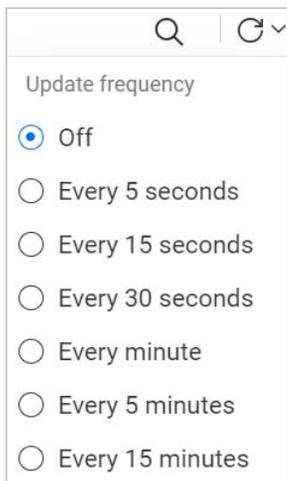
# What are the common tasks in each page?

Figure 3-2 Common Tasks



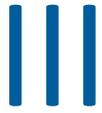
Some of the common tasks available on each page are:

- ◆ **Search:** Displays the list of objects that matches the specified criteria.
- ◆ **Refresh:** Reloads the page with the latest status of the object. If no frequency is set, then you must manually refresh the page for displaying the updated change.



- ◆ **Choose Column:** Displays available columns.
- ◆ **Export:** Downloads the data on the page in .csv format.





# Servers

- ◆ [Chapter 4, “Managing Server Settings,” on page 37](#)
- ◆ [Chapter 4, “Log and Service Details,” on page 25](#)



# 4 Managing Server Settings

This chapter describes the procedures for managing server settings through the Unified Management Console (UMC). For more information on NCP server settings, see [NCP Server for Linux Administration Guide](#).

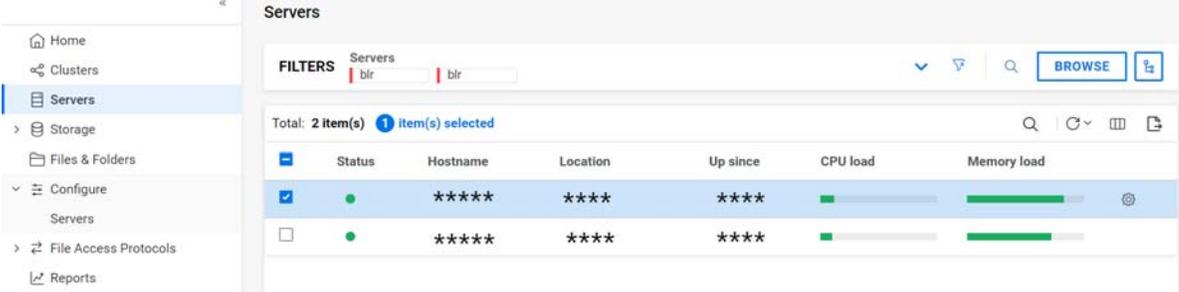
- ♦ “How to view the details of all available servers?” on page 37

## How to view the details of all available servers?

You can view the details of all the available servers in the **Servers** tab.

- 1 In UMC, select the **Servers** tab.
- 2 Search or browse to select the servers and click **APPLY**.

This displays the list of selected servers with related information such as **Status**, **Hostname**, **Location**, **Up since**, **CPU load**, and **Memory load**.



The screenshot shows the UMC interface with the 'Servers' tab selected. A sidebar on the left contains navigation options: Home, Clusters, Servers (selected), Storage, Files & Folders, Configure (with sub-items Servers and File Access Protocols), and Reports. The main content area is titled 'Servers' and features a search filter with the text 'blr'. Below the filter, it indicates 'Total: 2 item(s) 1 item(s) selected'. A table displays the following data:

	Status	Hostname	Location	Up since	CPU load	Memory load	
<input checked="" type="checkbox"/>	●	*****	****	****	<div style="width: 50%;"></div>	<div style="width: 70%;"></div>	⚙️
<input type="checkbox"/>	●	*****	****	****	<div style="width: 20%;"></div>	<div style="width: 40%;"></div>	

You can use the settings  icon to configure the selected server settings.



# 5 Log and Service Details

This section provides some additional information on UMC.

- ♦ [“Log Files” on page 25](#)
- ♦ [“Service Status” on page 25](#)

## Log Files

Check the below-mentioned logs for UMC issues related to debugging.

- ♦ UMC Server Details:

```
/var/opt/novell/log/umc/apps/umc-server/server.log
```

```
/var/opt/novell/log/umc/apps/umc-server/error.log
```

- ♦ OES-REST related messages:

```
/var/log/messages
```

- ♦ UMC Service Health Details:

```
/var/opt/novell/log/umc/apps/umc-server/health.log
```

## Service Status

To view the status of the services, use the following commands:

- ♦ To check details of edirapi service - `systemctl status docker-edirapi.service`
- ♦ To check details of PostgreSQL database server - `systemctl status postgresql.service`
- ♦ To check details of UMC REST API backend services - `systemctl status microfocus-umc-backend.service`
- ♦ To check details of UMC REST API Server service - `systemctl status microfocus-umc-server.service`
- ♦ To check details of Apache Webserver - `systemctl status apache2.service`
- ♦ To check details of Tomcat Servlet Container for OES services - `systemctl status novell-tomcat.service`



# IV Storage

- ♦ Chapter 6, “Managing NSS Pools,” on page 43
- ♦ Chapter 7, “Managing Volume Jobs,” on page 63
- ♦ Chapter 8, “Managing Pool Snapshots,” on page 69
- ♦ Chapter 9, “Managing NSS Volumes,” on page 73
- ♦ Chapter 10, “Managing User Quota,” on page 85
- ♦ Chapter 11, “Managing NSS Partitions,” on page 89
- ♦ Chapter 12, “Managing NSS Software RAID Devices,” on page 95
- ♦ Chapter 13, “Managing Devices,” on page 107



# 6 Managing NSS Pools

This chapter describes the procedures for creating and managing NSS pools on a server.

- ♦ [“What is a pool?” on page 43](#)
- ♦ [“What are the prerequisites for creating a new pool?” on page 43](#)
- ♦ [“How to create a new pool?” on page 44](#)
- ♦ [“How to list pools?” on page 46](#)
- ♦ [“How to view pool dashboard?” on page 47](#)
- ♦ [“How to deactivate or activate pool for pool maintenance?” on page 48](#)
- ♦ [“How to perform a pool move?” on page 50](#)
- ♦ [“What happens when I delete a pool?” on page 51](#)
- ♦ [“What are the prerequisites for deleting a pool?” on page 51](#)
- ♦ [“How to delete pools?” on page 51](#)
- ♦ [“How to rename a pool?” on page 52](#)
- ♦ [“How to increase the size of a pool?” on page 53](#)
- ♦ [“How to discard unused blocks in a pool?” on page 55](#)
- ♦ [“Where are my deleted volumes? Can i restore/salvage them?” on page 56](#)
- ♦ [“What are the prerequisites for AD users to access NSS data?” on page 57](#)
- ♦ [“I am an AD user. How do I access NSS data?” on page 58](#)
- ♦ [“The eDirectory pool object is corrupted. How to recover it?” on page 61](#)

## What is a pool?

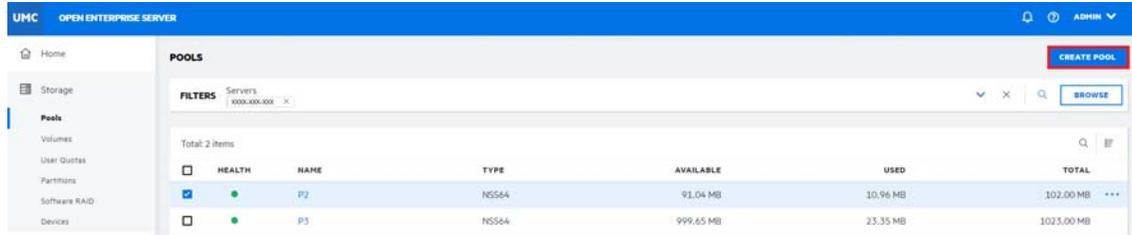
A pool is an area of storage that consists of space, called a partition, obtained from one or more of the storage devices available on a server. The amount of space that each storage device contributes varies. NSS uses storage pools to efficiently acquire and use all free space available on devices.

## What are the prerequisites for creating a new pool?

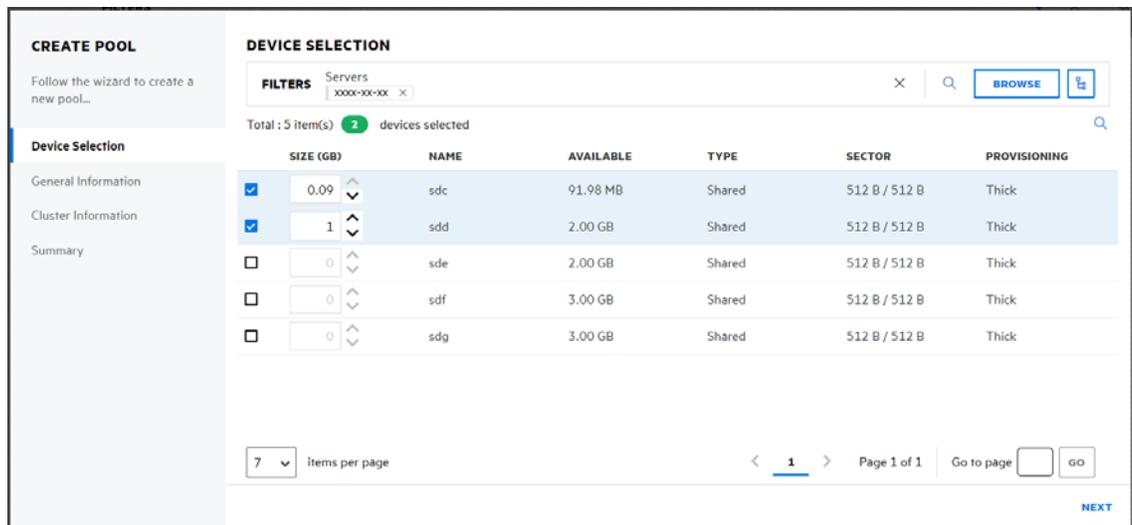
- ♦ Devices must be initialized, so that the available space is displayed for creating a pool.
- ♦ OES CIFS must be installed, configured, and running before you can select the CIFS option when cluster enabling an NSS pool.

# How to create a new pool?

- 1 In UMC, click  **Storage > Pools**.
- 2 Click **CREATE POOL**.



- 3 On the **DEVICE SELECTION** page, search or browse to select the server, and select the required device(s).



- 4 Specify the device space up to the free space available on the devices for the pool, and then click **NEXT**.  
Only initialized devices that have free space are listed on the device selection page. If no devices are listed, cancel the wizard, add more devices to the server or free up space on the existing devices.
- 5 On the **INFORMATION** page, specify a name for the new pool, then click **NEXT**.

The description is an optional field. All the NSS64-bit pools are AD media upgraded by default.

- 6 If the selected device type is shared, on the **CLUSTER INFORMATION** page, specify the details as required, and then click **NEXT**.

The Enable Cluster toggle is automatically turned on. Turn it off to create a non-clustered pool with shared devices.

---

**NOTE:** This page is not available if the selected device type is local on the **DEVICE SELECTION** page.

---

### Parameters Required for Cluster Enabled Pools:

- ◆ **Virtual Server Name:** This name is assigned to the virtual server that represents the shared pool in the cluster. When you cluster-enable a pool, a virtual server object is automatically created in eDirectory and given the name of the cluster object plus the name of the cluster-enabled pool. For example, if the cluster name is cluster1 and the cluster-enabled pool name is pool1, then the default virtual server name will be cluster1\_pool1\_server. You can edit the field to change the default virtual server name. The virtual server name used for NCP and CIFS servers will be the same.

- ◆ **IP Address:** The IP address that you want to assign to the virtual server. Each cluster-enabled NSS pool requires its own IP address. The IP address is used to provide access and fail-over capability to the cluster-enabled pool (virtual server). The IP address you assign to the pool remains assigned to the pool regardless of which server in the cluster is accessing the pool.

---

**IMPORTANT:** The IP address for the virtual server must be in the same IP subnet as the server nodes in the cluster where you plan to use it.

---

- ◆ **Advertising Protocols:** Protocols that give users native file access to data.  
Specify one or more advertising protocols by using the toggle button of the protocols you want to enable for data requests to this shared pool.
- ◆ **OES NCP:** NCP is the networking protocol used by the Client for Open Enterprise Server. It is selected by default. Selecting NCP causes commands to be added to the pool-resource load and unload scripts to activate the NCP protocol on the cluster.
- ◆ **CIFS:** CIFS is a Windows networking protocol. Selecting CIFS causes commands to be added to the pool-resource load and unload scripts to activate the CIFS protocol on the cluster.

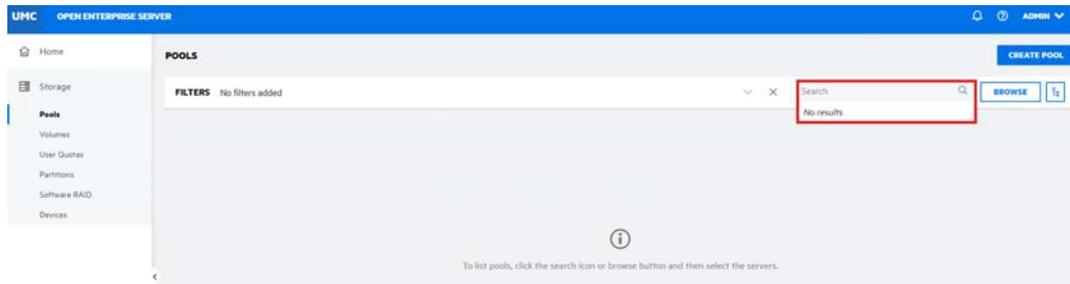
7 Review the pool details, and click **FINISH** to create the pool.

CREATE POOL		SUMMARY		ALLOCATED DEVICES			
Follow the wizard to create a new pool.		<b>GENERAL INFORMATION</b>		<b>NAME</b>	<b>ALLOCATED</b>	<b>PROVISIONING</b>	<b>TYPE</b>
Device Selection	✓	Selected Server	xxxx-xxx-xxx	sdc	3.50 GB	Thick	Shared
General Information	✓	Pool Name	example_pool1				
Cluster Information	✓	Pool Description					
<b>Summary</b>		<b>CLUSTER INFORMATION</b>					
		Clustering enabled	✓				
		Virtual Server	CLUSTER123				
		IP Address	xxxx-xxx-xxx				
		NCP	✓				
		CIFS	✗				
				PREVIOUS FINISH			

## How to list pools?

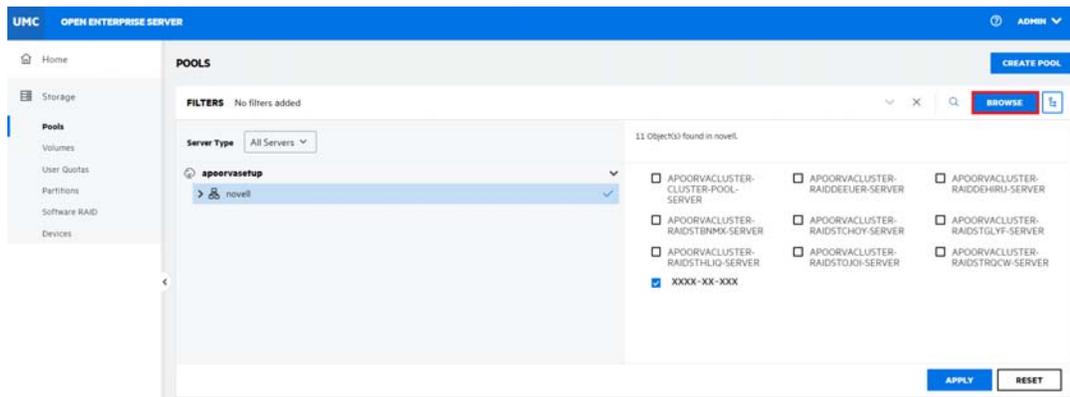
You can view the list of pools and their related information available on the server. The pool's list also includes pool snapshots if you have previously created a pool snapshot.

- 1 In UMC, click  **Storage > Pools**.
- 2 Click the Search Icon and specify the server name.



or

Click **Browse**, select **Server Type** to list their associated servers. Select the required servers from the list, and then click **APPLY**.




---

**NOTE:** When the **BROWSE** or tree view icon is clicked, you cannot perform other actions outside the browse area. Click the same button again to close the browse or tree view area.

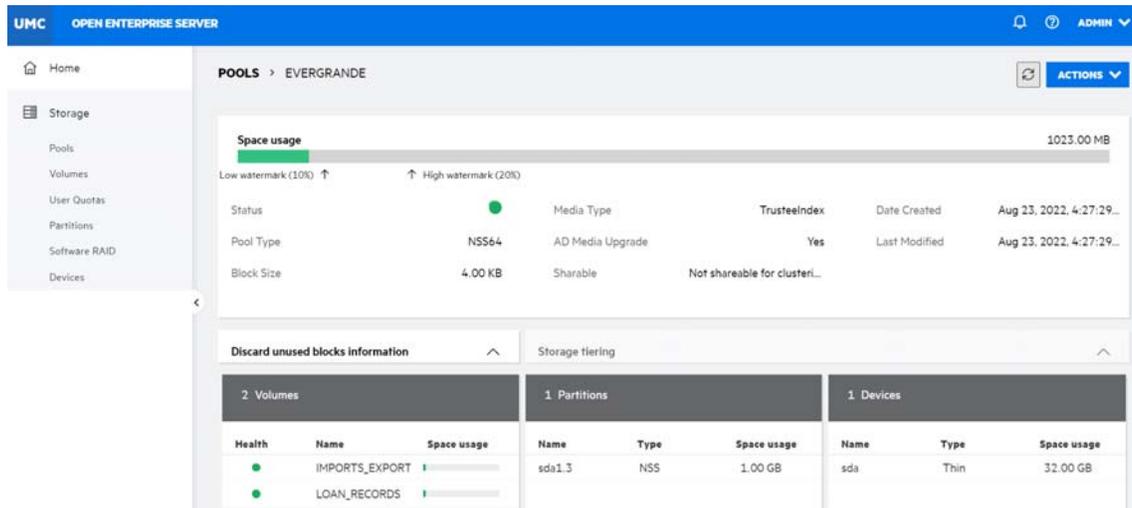
---

The list of pools available on the selected servers is displayed.

## How to view pool dashboard?

You can view the details of a pool like space usage, volumes, partitions, and devices on the pool dashboard page.

- 1 In UMC, click  **Storage > Pools**.
- 2 Search or browse the servers to list the pools associated with them.
- 3 Click on the pool name to view the pool dashboard page.



You can use **ACTIONS** to perform various pool operations like rename, increase size, manage snapshot, create snapshot, update pool object, discard unused blocks, activate, deactivate, and delete.

## How to deactivate or activate pool for pool maintenance?

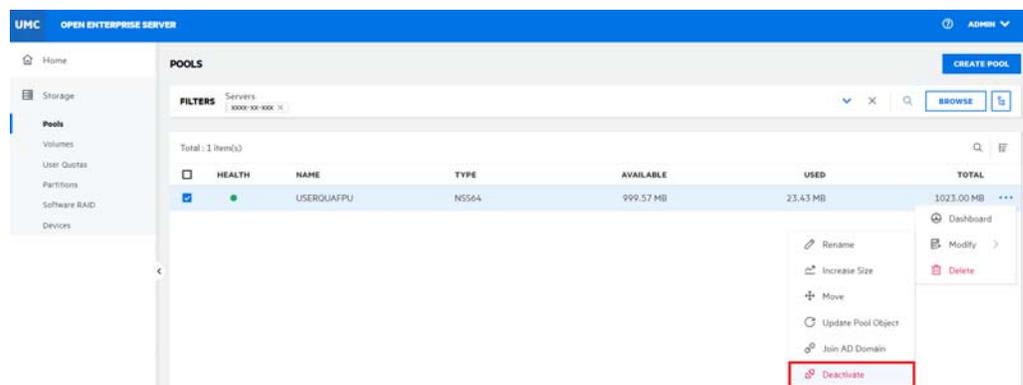
You can perform pool maintenance without shutting down the server. Access to a pool can temporarily be restricted by deactivating them.

After deactivating a pool, perform the pool maintenance. The pool and its volumes are temporarily unavailable to the users. Deactivating a pool does not delete the volumes or their data.

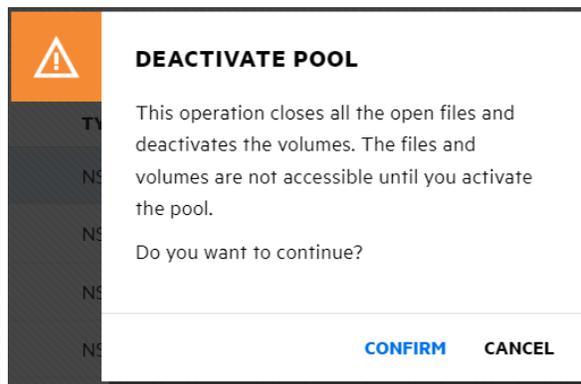
- 1 In UMC, click  **Storage > Pools**.
  - 2 Search or browse the servers to list the pools associated with them.
- 
- 3 **NOTE:** If you select multiple pools, the More Options  icon is available at the top right corner of the table.

### 3a To deactivate a pool:

- 3a1 Select a pool, click More Options  icon, click **Modify**, and then select **Deactivate**.



- 3a2 Click **CONFIRM** to deactivate the selected pools.

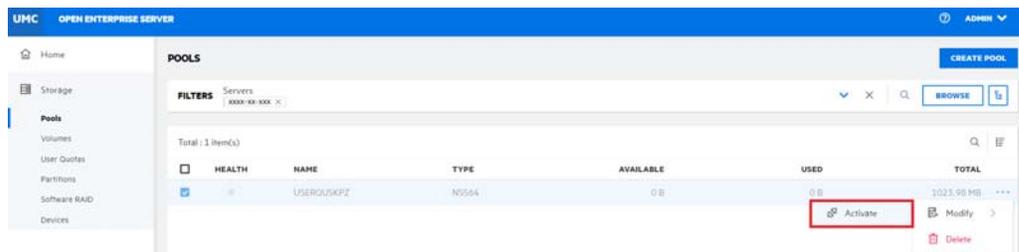


The deactivated pool details are not displayed on the **POOLS** page.

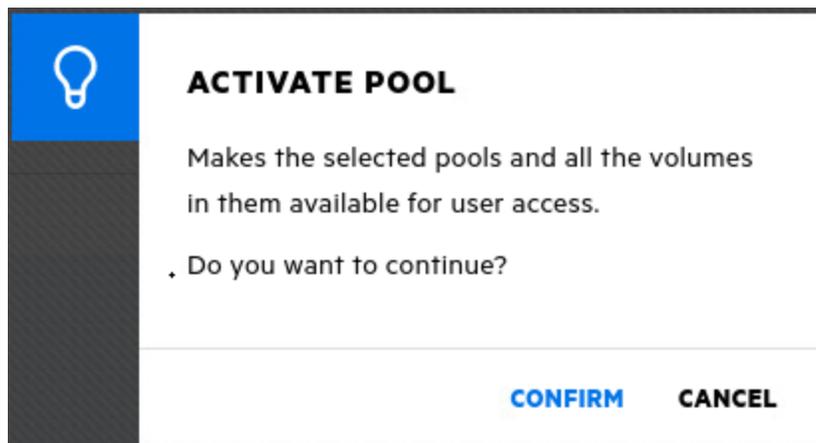
OR

**3b To activate a pool:**

**3b1** Select a pool, click More Options **...** icon, click **Modify**, and then select **Activate**.



**3b2** Click **CONFIRM** to activate the selected pool.

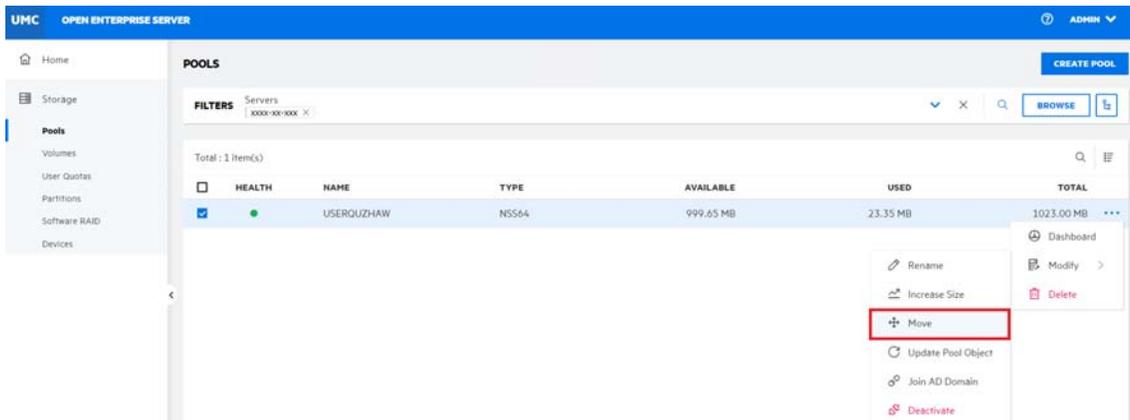


You must manually activate the volumes after the pool is active. For more information on activating volumes, see [“How to deactivate and activate NSS volumes?”](#) on page 78.

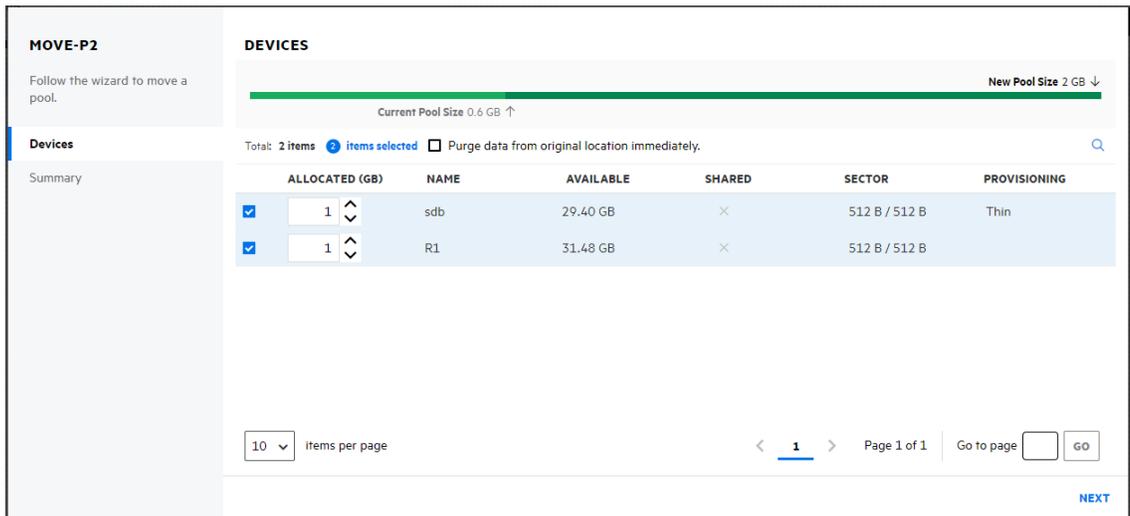
# How to perform a pool move?

You can move an NSS pool from one location to another on the same system. The pool remains active during this process. All the segments in the pool are consolidated and moved to the specified device(s). If a specified device is larger than the original device, the pool is automatically expanded on completion of the move job.

- 1 In UMC, click  **Storage > Pools**.
- 2 Search or browse the servers to list the pools associated with them.
- 3 Select the pool you want to move, click More Options  icon, click **Modify**, and then select **Move**.



- 4 Select the device(s), specify the required allocated size for the selected device(s), and then click **NEXT**.



Select the **Purge data from original location immediately** checkbox to permanently delete the moved pool from the original location after the pool move.

- 5 Review the details and click **FINISH**.

**MOVE-P2**

Follow the wizard to move a pool.

Devices ✓

**Summary**

**SUMMARY**

Purge data from original location immediately.

**POOL ALLOCATIONS**

NAME	REMAINING	ALLOCATED
sdb	28.40 GB	1.00 GB <span style="color: blue;">↑</span>
R1	30.48 GB	1.00 GB <span style="color: blue;">↑</span>

PREVIOUS FINISH

The pool is moved to the selected devices after the process is successfully completed.

## What happens when I delete a pool?

Deleting a pool removes the ownership of the space it occupied, freeing the space for reassignment. The **Delete** option on the **Pools** page removes the selected pools from the server, including all member partitions and the data on them.

NSS pools can be deleted to create free space for other pools.

---

**WARNING:** ♦Deleting a pool deletes all the volumes and data in it. These volumes cannot be restored.

- ♦ If the pool is created on a RAID1 device, deleting the pool deletes the RAID1 device.
- 

## What are the prerequisites for deleting a pool?

- ♦ If the pool is shared in an OES Cluster, you must offline the cluster resource before you attempt to delete the clustered pool or its cluster resource.
- ♦ If the pool has pool snapshots, you must delete the pool snapshots before deleting the pool.

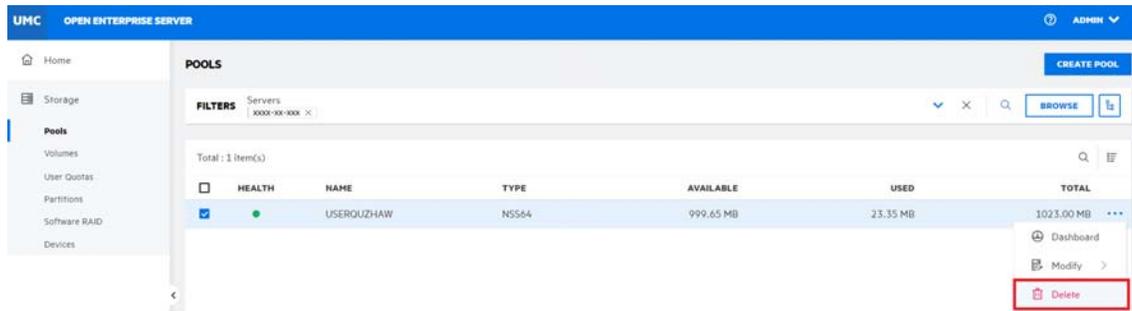
## How to delete pools?

- 1 In UMC, click **Storage > Pools**.
- 2 Search or browse the servers to list the pools associated with them.
- 3 Select the pool, click More Options , and then select **Delete**.

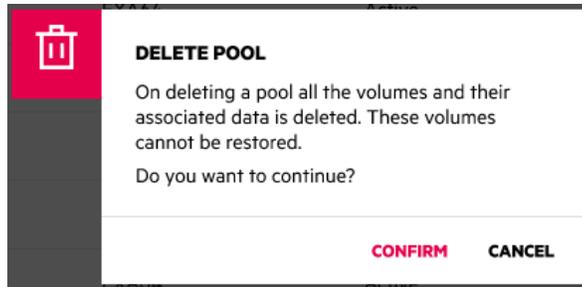
---

**NOTE:** If you select multiple pools, the More Options icon is available at the top right corner of the table.

---



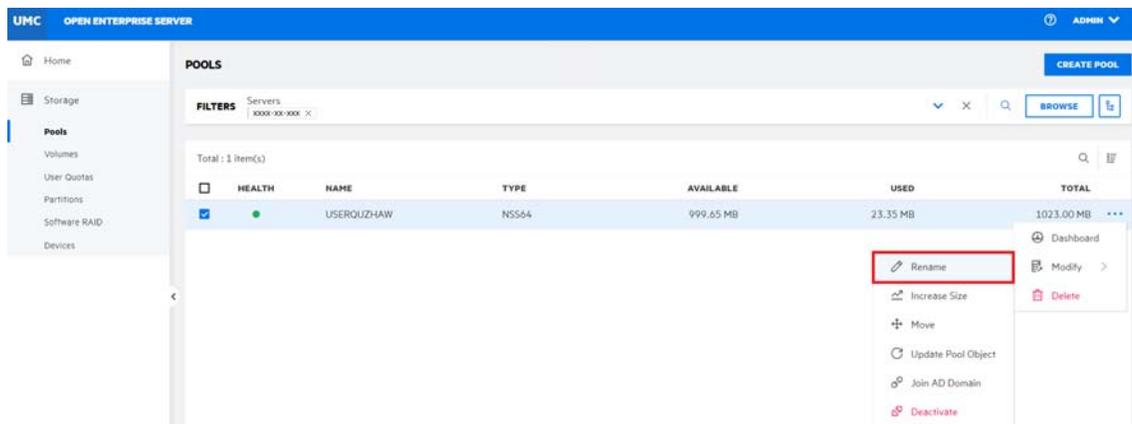
4 Click **CONFIRM** to delete the selected pool.



## How to rename a pool?

The **Rename** option on the **Pools** page allows you to modify the name of a pool. You can change the name of a pool that corresponds to a department name change. When you rename a pool, it must be in the active state so that the eDirectory gets updated.

- 1 In UMC, click **Storage > Pools**.
- 2 Search or browse the servers to list the pools associated with them.
- 3 Select the pool, click More Options **...** icon, click **Modify**, and then select **Rename**.



4 Specify the new pool name and click **CONFIRM**.

**RENAME POOL**

Name POOL\_6

New Name\*

[CONFIRM](#) [CANCEL](#)

**Table 6-1** Actions Required After Renaming a Pool

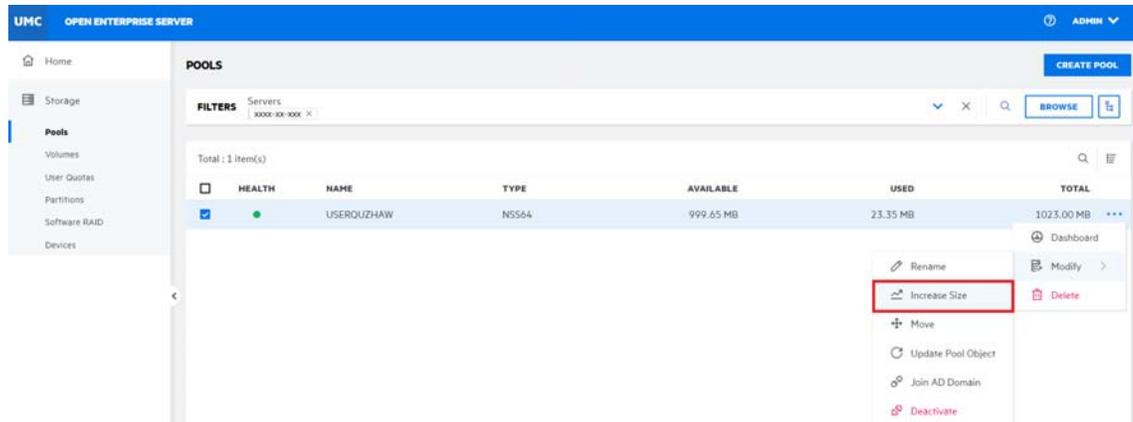
Pool Share State	Pool Load-Time State	Pool State After a Rename	Action Required
Unshared	Autoloaded	Active with volumes dismounted	Mount the pool's volumes
Unshared	Not autoloaded	Deactive	Activate the pool, then mount its volumes
Shared	Load and unload are controlled by OES Cluster Services. Before you rename a cluster-enabled pool, make sure to offline the pool resource, activate the pool by using UMC or NSSMU instead of using the load script, then you can rename the pool by using UMC or NSSMU.	Deactive	Online the pool resource to activate the pool and its volumes. OES Cluster Services automatically updates the pool resource load and unload scripts to reflect the name change. Also, NSS automatically changes the Pool Resource object name in eDirectory.

## How to increase the size of a pool?

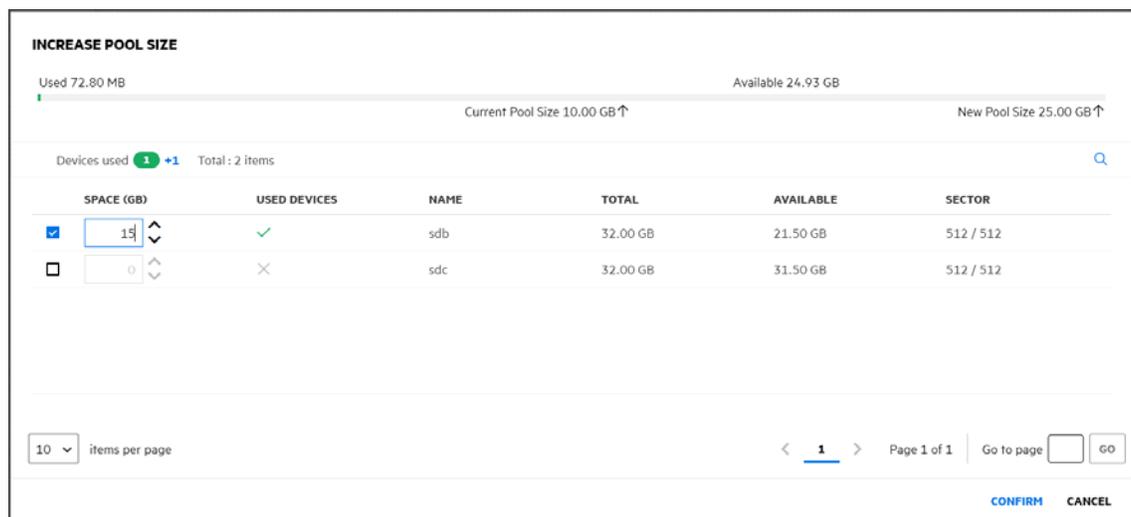
The **Increase Size** option on the **Pools** page allows you to expand the storage capacity of a selected pool by adding new partitions. You can increase the size of your storage pools, but you cannot reduce their size.

- 1 In UMC, click  **Storage > Pools**.
- 2 Search or browse the servers to list the pools associated with them.

- 3 Select the pool, click More Options **⋮** icon, click **Modify**, and then select **Increase Size**.



- 4 Select the devices and specify the space to be used from each device.  
Only the devices that have free space are listed. If no devices are listed, there is no space available to increase the size of the pool. Click **Cancel**, add more devices to the server or free up space on the existing devices, then return to the **POOLS** page to increase the size of the pool.



- 5 Click **CONFIRM** to expand the size of the selected pool.

# How to discard unused blocks in a pool?

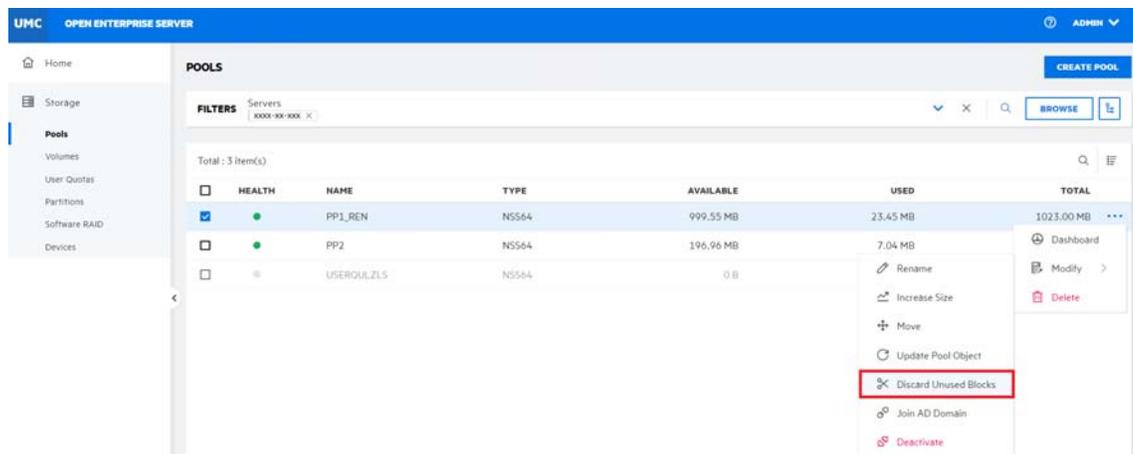
The unused blocks in the selected pool can be freed by using the **Discard** option, making them available for usage. This functionality is supported only on thin-provisioned SCSI devices with VMware ESXi on a linear target.

**Table 6-2** Support Matrix

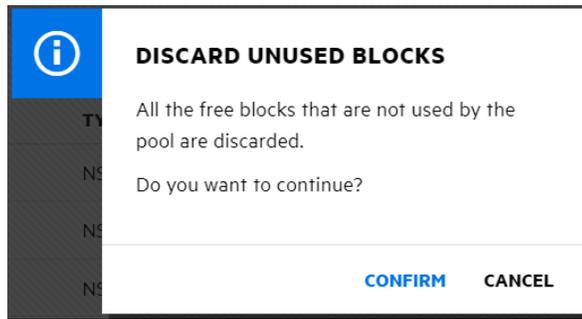
Device Type	Provisioning Type		Support on Pool
SCSI Devices with Vmware ESXi	Thin		Supported
SCSI Devices with Vmware ESXi		Thick	Not Supported
SCSI Devices with Vmware ESXi	Thin	Thick	Not Supported
RAID Devices	Any type	Any type	Not Supported
Pools that contain Snapshots	Any type	Any type	Not Supported

- 1 In UMC, click  **Storage > Pools**.
- 2 Search or browse the servers to list the pools associated with them.
- 3 Select the pool, click More Options  icon, click **Modify**, and then select **Discard Unused Blocks**.

**NOTE:** If you select multiple pools, the More Options  icon is available at the top right corner of the table.



- 4 Click **CONFIRM** to discard the unused blocks in the selected pool.



The process is executed in the background and discards unused blocks on the selected pool.

## Where are my deleted volumes? Can i restore/salvage them?

On deleting a volume, NSS removes it from the pool. The **Deleted Volumes** option on the **Pools** page displays a separate **Deleted Volumes** page where you can purge or salvage the deleted volumes for the pool. This option is only available if the selected pool has deleted volumes in it.

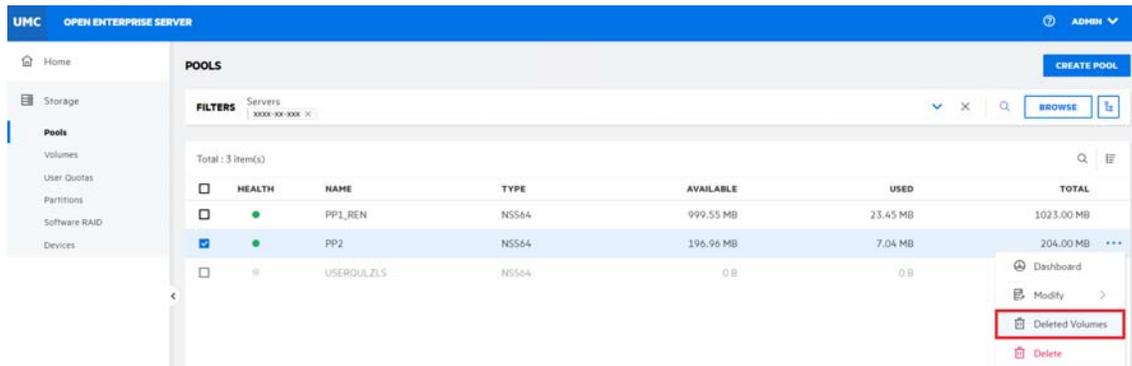
During the Purge Delay (by default, four days after a volume is deleted), you can manually purge deleted volumes, view the volume contents, transfer files from the deleted volume to other volumes, or salvage the entire volume. When you salvage a volume, the data and metadata are the same as they were at delete time, with no changes. After the Purge Delay time elapses, NSS automatically purges the deleted volume from the system, and it can no longer be accessed.

---

**WARNING:** If you delete an entire pool, all the volumes are deleted with it. You cannot restore a deleted pool or volumes in it.

---

- 1 In UMC, click  **Storage** > **Pools**.
- 2 Search or browse the servers to list the pools associated with them.
- 3 Select the pool, click More Options  icon, and then select **Deleted Volumes**.



- 4 Select the deleted volume, click the (...) option, and then select **Salvage/Purge**.



**Salvage:** You can restore and assign a new name to the deleted volume or reuse the old name if no other volume is using that name.

**NOTE:** If you salvage an encrypted volume, you are prompted for the related password.

**SALVAGE VOLUME**

Existing Volume Name: VOLUME\_6

New Volume Name\*:

[CONFIRM](#) [CANCEL](#)

**Purge:** You can manually delete one or more deleted volumes and can no longer be salvaged or recovered.

**PURGE VOLUME**

Purging permanently deletes the volume and its contents.

Do you want to continue?

[CONFIRM](#) [CANCEL](#)

5 Click **CONFIRM** to complete the selected process.

## What are the prerequisites for AD users to access NSS data?

- ◆ The pool must contain at least one active volume.
- ◆ The pool must support AD Media.
- ◆ CIFS service must be configured and operational on the pool.
- ◆ CIFS service must be configured and operational on the OES server.
- ◆ Server must be added to AD domain.

# I am an AD user. How do I access NSS data?

Active Directory (AD) users are Windows users who use the CIFS protocol to access NSS volumes on OES servers and administer them. AD users and groups are not required to be moved to eDirectory as NSS resources can be accessed by both AD and eDirectory users at the same time.

---

**NOTE:** Ensure that the logged-in user has sufficient rights to create the object in the particular container in AD before joining a pool to an AD domain.

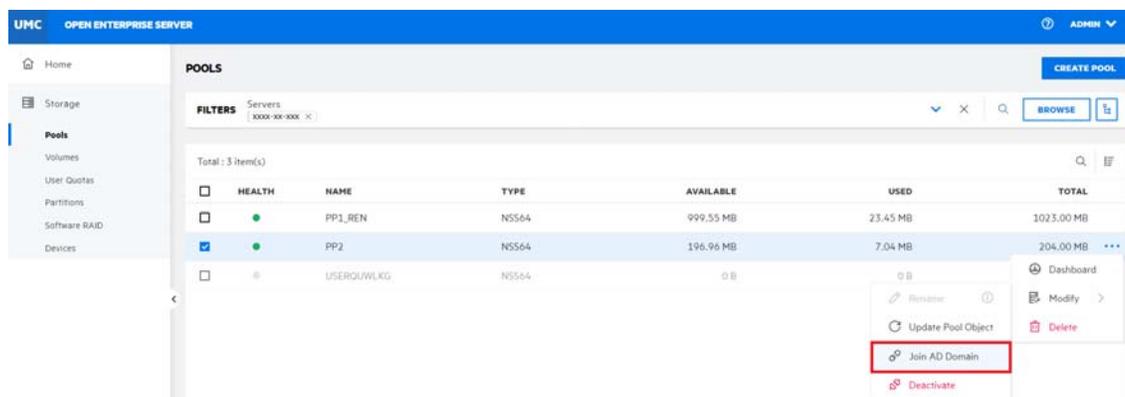
---

- 1 In UMC, click  **Storage > Pools**.
- 2 Search or browse the servers to list the pools associated with them.
- 3 Select the pool, click More Options  icon, and then select **Join AD Domain**.

---

**NOTE:** The **Join AD Domain** option is available only if the selected pool is AD enabled.

---



---

**NOTE:** If any of the prerequisites for joining the AD Domain is not met, the **CANNOT JOIN AD DOMAIN** page is displayed to indicate it. Click **CANCEL**, ensure to meet the prerequisites, then perform Join AD Domain from the **POOLS** page. See [“What are the prerequisites for AD users to access NSS data?”](#) on page 57.

## CANNOT JOIN AD DOMAIN

Ensure that the following prerequisites are met:



### Volume

The pool must have one active volume.



### AD Media Support

The pool must support AD Media.



### CIFS

CIFS should be enabled for the pool.



### CIFS

CIFS should be enabled for the server.



### Server

Server should be added to AD domain.



**CANCEL**

- 4 In the **AUTHENTICATION** page, specify the **Username** and **Password** of the AD user, and click **TEST CONNECTION**.

**JOIN AD DOMAIN**  
Follow the wizard to join AD Domain

**Authentication**

General Information

**AUTHENTICATION**

Domain Name MFDOMAIN.COM

NETBIOS Name clst75-tp1-w

Username\*

Password\*

**TEST CONNECTION**

NEXT

The existence of the user in the AD database is verified. After successful verification of the domain, click **NEXT**.

**5** Follow the steps for selecting or creating an object.

**5a** For selecting a pre-existing object in the active directory:

If you already have a computer object created in the active directory for the server, follow the steps to select the object.

**5a1** Select the **Use pre-created computer object** checkbox.

**5a2** Specify the **Container** name.

**5a3** Specify the description details, and then click **FINISH**.

**JOIN AD DOMAIN**  
Follow the wizard to join AD Domain

Authentication

**General Information**

**INFORMATION**

Domain Name WIN2022IMD.COM

NETBIOS Name ster123-cp123-w

Use pre-created computer object

Container\*

Description

PREVIOUS **FINISH**

OR

**5b** For creating a new object in the active directory:

If you have no computer object created in the active directory for the server, follow the steps to create an object.

---

**NOTE:** : Make sure to uncheck the “Use pre-created computer object” checkbox.

---

**5b1** Specify the container name.



The screenshot shows the 'JOIN AD DOMAIN' wizard in the 'General Information' tab. The 'Domain Name' is 'WIN2022IMD.COM' and the 'NETBIOS Name' is 'ster123-cp123-w'. The 'Container' field is set to 'CN=Computers'. The 'Use pre-created computer object' checkbox is unchecked. The 'Description' field is empty. At the bottom right, there are 'PREVIOUS' and 'FINISH' buttons.

**5b2** Specify the description details, and then click **FINISH**.

The AD users have access to NSS volumes after the process is successfully completed.

## The eDirectory pool object is corrupted. How to recover it?

The **Update Pool Object** option on the **Pools** page allows you to add or update the eDirectory pool object. If the pool object already exists, NSS prompts with two options: Delete and replace the existing object, or Retain the existing object.

---

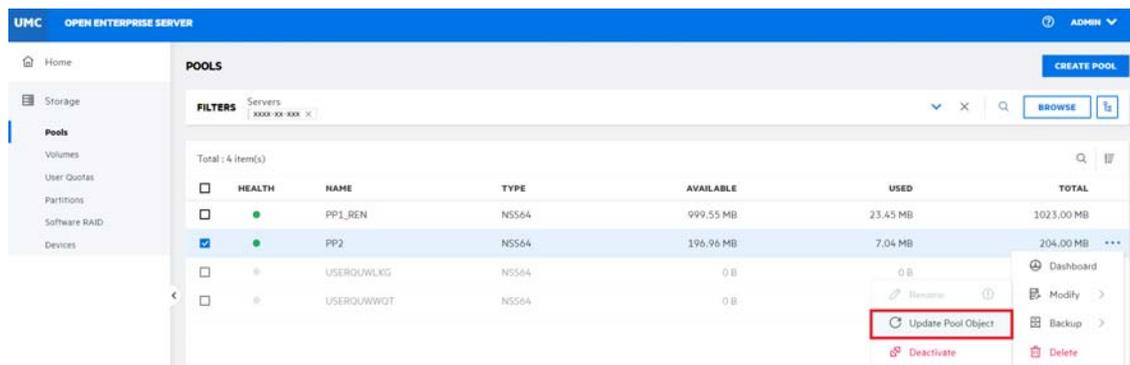
**NOTE:** Updating eDirectory pool object is a recovery process and is required only when the pool object is lost, corrupted, or deleted.

---

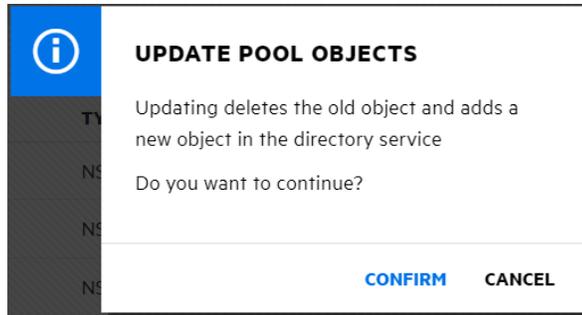
Update eDirectory object after modifying pool parameters or renaming it.

For more information on eDirectory, refer to the [eDirectory 9.2 documentation](#).

- 1 In UMC, click  **Storage > Pools**.
- 2 Search or browse the servers to list the pools associated with them.
- 3 Select the pool, click More Options  icon, click **Modify**, and then select **Update Pool Object**.



- 4 Click **CONFIRM** to update the pool objects for the selected pool.



If the pool object does not exist, NSS adds it at the same context level as the server.

# 7 Managing Volume Jobs

This chapter outlines the procedures for monitoring and managing DFS jobs.

You can monitor the status of all the active Move and Split jobs and recently completed jobs that are initiated for a selected server. Additionally, you have options to pause, resume, start now, reschedule, finish, or cancel a job, depending on its state.

---

**NOTE:** To access DFS jobs, log in to UMC with admin credentials, then click **Storage > Volumes > Jobs**.

---

## How to view list of DFS jobs?

To view Move and Split jobs, click **Storage > Volumes > Jobs**, browse and select a server. The following information is displayed:

*Table 7-1 DFS jobs*

Column Name	Description
Source	Displays the name of the source volume, such as VOL1 : or path to folder on the source volume for Split jobs, such as VOL2 :beta/dev.
Status (color coding)	Displays the current status of the job.
Green	<b>Completed:</b> The jobs are complete.  Completed jobs remain in the status report for seven days. When a job reaches the <b>Completed</b> state, no files are left in the source volume or below the DFS junction point in the source volume.
Blue	<b>In Progress:</b> Data is actively being transferred from the source volume to the destination, represented in form of percentage.
Orange	<b>Suspended:</b> The job is not responding, and administrator intervention is required.
Red	<b>Failed:</b> The job has failed.  <b>Paused:</b> The job is manually paused. It can be resumed or deleted.  <b>Scheduled:</b> The job is set to run at a scheduled date and time.

Column Name	Description
	<p><b>Cleanup failed:</b> DFS was unable to delete the files from the source volume after transferring the data to the destination volume. One of the reason is the files were in use when cleanup was initiated.</p> <ul style="list-style-type: none"> <li>◆ To retry the cleanup, click <b>Retry</b>.</li> </ul> <p>If undeleted files remain in use, DFS might return to this state. This option can be repeated as needed.</p> <ul style="list-style-type: none"> <li>◆ To complete the job and leave undeleted files in the source volume, click <b>Finish</b>. It is recommended to delete the duplicate files from the source volume.</li> </ul> <p><b>Files skipped:</b> Files that are in use at the time of data transfer and could not be copied to the destination.</p> <p><b>Canceled:</b> Deletes the job.</p>
Type	Indicates whether the job is a Move job or a Split job.
ID	An auto-generated unique ID assigned to the job.
Server	The servers displaying the Move and Split jobs.
Date scheduled	The date and time when the job is set to run.

## What are the prerequisites to Move or Split jobs?

- ◆ **DFS Management Context:** Make sure that a DFS management context is set up, containing both source and destination servers.
- ◆ **DFS Replica service:** Verify the DFS Replica service for the management context is synchronized and running.
- ◆ **NCP and SMS services:** Make sure that NCP and SMS services are installed and running on both the source and destination servers.
- ◆ **Deleted files:** If there are deleted files on the source volume that need to be transferred to the destination, salvage those deleted files before starting the process.
- ◆ **Destination volume space:** Make sure the destination volume has sufficient free space to accommodate the data being transferred.
- ◆ **Administrative rights:** If moving the volume to a pool on a different server, ensure administrative rights are available on the destination server.
- ◆ **SLP registration:** Verify that the destination server is registered with SLP for `smrd` service and that SLP is operational.

## How to perform a Move job?

A Move job transfers the file structure, data, and file system trustee rights from a source NSS volume to a destination NSS volume within the same DFS management context.

Before you begin, make sure you have met the [prerequisites](#).

- 1 Log in to UMC with admin credentials.
- 2 Click **Storage > Volumes**.
- 3 Browse and select the server that contains the NSS volume you want to move.
- 4 From the **Volumes** list, select the volumes you want to move, and click **Move**.
- 5 Select the destination volume to transfer the data.  
The trustee rights from the source volume are automatically applied to the destination volume.
- 6 Select **Start now** to begin the data transfer. The job may take a few seconds to initiate, changing the status to Scheduled. No action is required from your end during this process.

or

Alternatively, specify the date and time to schedule the transfer.

---

**NOTE:** Make sure that the volumes are active at the scheduled time.

---

- 7 Displays a summary of the Move job. Review it and click **Finish**.  
The data transfer might take from a few minutes to several hours, depending on the volume of data being moved.

## How to perform a Split job?

A Split job transfers a portion of the file structure, data, and the file system trustee rights from a source NSS volume to a destination NSS volume within the same DFS management context.

Before you begin, make sure you have met the [prerequisites](#).

- 1 Log in to UMC with admin credentials.
- 2 Click **Storage > Volumes**.
- 3 Browse and select the server that contains the NSS volume you want to split.
- 4 From the **Volumes** list, select the volumes that you want to split, and click **Split**.
- 5 Select the folder where the DFS junction will reside.  
All data below this folder is moved to the destination volume.
- 6 Select the destination volume to transfer the data.  
The trustee rights from the source volume are automatically applied to the destination volume.
- 7 Select **Start now** to begin the data transfer. The job may take a few seconds to initiate, changing the status to Scheduled. No action is required from your end during this process.

or

Alternatively, specify the date and time you want to schedule the transfer.

---

**NOTE:** Make sure that the volumes are active at the scheduled time.

---

- 8 Displays a summary of the Split job. Review it and click **Finish**.

The data transfer might take from a few minutes to several hours, depending on the volume of data being moved.

## What happens when jobs are paused?

Pause suspends a job until it is manually resumed or deleted. Only jobs that are currently in progress, scheduled, or suspended can be paused.

You can pause a Move or Split job to allow another job to run or to reduce the load on the system or network.

Select one or more active jobs, click **Pause**, then type a Comment to be displayed in the status report.

## What happens when jobs are resumed?

Resume continues the jobs from the point where it was paused. The data transfer continues depending on the job type.

---

**IMPORTANT:** You cannot resume a job that is completed, failed, or deleted.

---

Select one or more paused jobs, click **Resume** to complete the job.

## How to handle files skipped by Move or Split Jobs?

A Move or Split job displays a **Files skipped** status, when some of files are not moved because they were in use at the time DFS attempted to copy them to the destination volume.

### View skipped files:

- 1 Log in to UMC with admin credentials.
- 2 Click **Storage > Volumes > Jobs**.
- 3 Browse and select the server that contains the Move or Split jobs.

Jobs with a status of **Completed**, **Failed**, or **Canceled** remain in the status report for only seven days.

- 4 Select jobs with status **Files Skipped**, then click **View skipped files** to displays the files that were not transfered to the destination volume.

---

Column Name	Description
Name	Name of the file that was not transfered to the destination volume.
Type	The format of the file.
Path	Location of the file on the source volume.

---

### Retry skipped files

To copy the skipped files, click **Retry**.

You have to manually trigger the retry action to transfer the files. If the files are still in use, DFS returns to the **Files skipped** state.

### Finish

To complete the job and skip the files that were not transferred from the source volume, click **Finish**.

The source volume is deleted upon job completion, and the skipped files are no longer accessible. Before you click **Finish**, manually transfer any skipped files to the destination volume.

## How to cancel or delete jobs?

A job can be canceled up to a certain point during the Move or Split process. If data has been transfered beyond a specific stage, DFS returns an error message and prevents the job from being canceled. After the cancel command is initiated, UMC waits for the next convenient point to stop the job.

For example, if a large file is being transferred, UMC will wait for that file to finish transferring before canceling the job.

Data on the source remains unchanged. However, it is recommended to manually clean up any data that has been transfered to the destination volume.

Select one or more jobs, click **Cancel** to initiate the job cancellation process.

---

**NOTE:** If you cancel a job, you must initiate a new job, as partial data transfers are not supported.

---



# 8 Managing Pool Snapshots

This section describes the procedure for creating and managing pool snapshots on a server.

---

**NOTE:** The status of a snapshot is usually Offline, and its state is Active.

---

*Figure 8-1 Snapshot*



NAME	STATUS	STATE	PARTITION	PARTITION SIZE	% USED
POOL1_SNAPSHOT	Offline	Active	sdb1.2	1.00 GB	0%

- ◆ “What is a pool snapshot?” on page 69
- ◆ “What are the prerequisites for creating a pool snapshot?” on page 69
- ◆ “How to create a pool snapshot?” on page 69
- ◆ “How to list pool snapshots?” on page 70

## What is a pool snapshot?

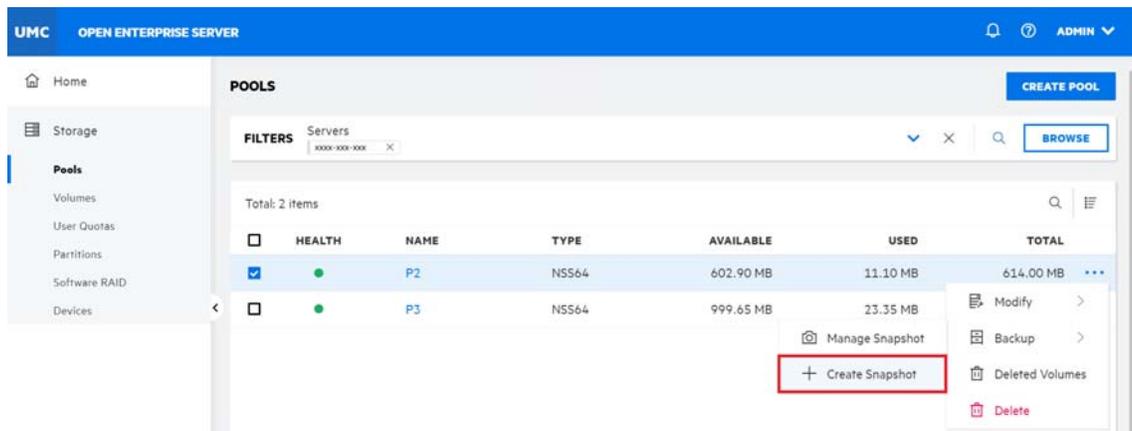
A pool snapshot is a metadata copy of a pool at a point-in-time. A pool snapshot improves the backup and restore services by saving time.

## What are the prerequisites for creating a pool snapshot?

- ◆ The pool you want to snapshot must already exist and be active.
- ◆ Free space must be available on a device to use it as the stored-on partition.
- ◆ Pool snapshots are not supported for shared NSS pools.

## How to create a pool snapshot?

- 1 In UMC, click  **Storage** > **Pools**.
- 2 Search or browse the servers to list the pools associated with them.
- 3 Select the pool, click More Options  icon, click **Backup**, and then select **Create Snapshot**.



**NOTE:** Creating a snapshot of a pool snapshot is not supported. If the selected pool is a pool snapshot, then the **Create Snapshot** option will not be available.

When you create a snapshot, both the original pool and the pool where the snapshot is stored must be active.

- 4 Specify the **Snapshot name**, select **Snapshot size**, select device from the list, and then click **CONFIRM**.

The 'CREATE SNAPSHOT' dialog box is shown. It contains the following fields: 'Snapshot name' with the value 'Pool\_6\_Snapshot', 'Snapshot size (GB)' with a dropdown set to '0.05', and 'Selected device' with 'sdb' selected. Below these fields is a table of available devices. The table has columns: DEVICE NAME, AVAILABLE, PHYSICAL SECTOR SIZE, and LOGICAL SECTOR SIZE. The 'sdb' device is highlighted in blue, indicating it is selected. The table shows 7 total available devices.

DEVICE NAME	AVAILABLE	PHYSICAL SECTOR SIZE	LOGICAL SECTOR SIZE
sda	0 B	512	512
sdb	21.45 GB	512	512
sdc	31.50 GB	512	512
sdd	0 B	512	512
sde	0 B	512	512
sdf	0 B	512	512
sdg	0 B	512	512

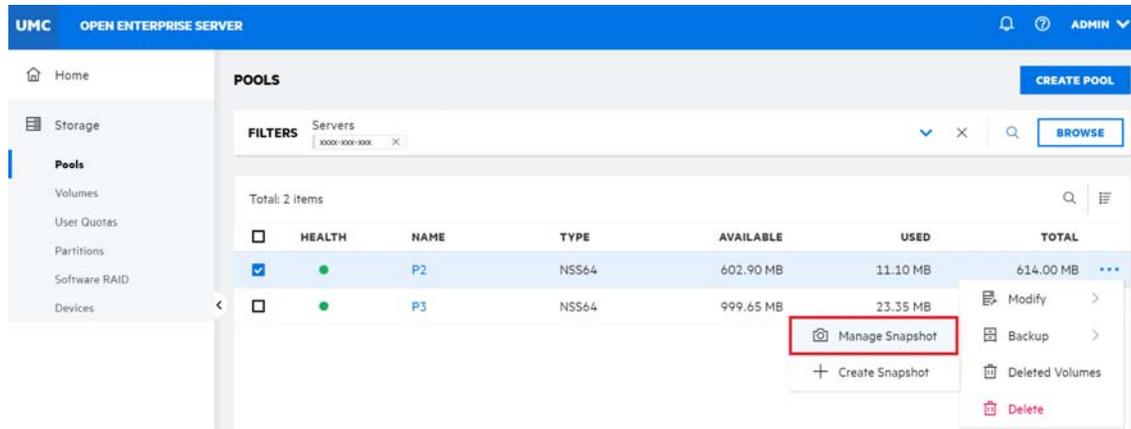
At the bottom of the dialog, there are 'CONFIRM' and 'CANCEL' buttons.

Minimum size required for a snapshot is 50MB. The newly created snapshot is available in the snapshot list with the status as offline. This snapshot is made online and accessed from the pool list for recovery.

## How to list pool snapshots?

- 1 In UMC, click **Storage > Pools**.
- 2 Search or browse the servers to list the pools associated with them.

3 Select the pool, click More Options **...** icon, click **Backup**, and then select **Manage Snapshot**.



4 Select the snapshot, click More Options **...** icon, and then select the required action.

**NOTE:** If you select multiple snapshots, the More Options **...** icon is available at the top right corner of the table.



- ◆ **Make Online:** This option makes the selected pool snapshot online so that you can access the data on it for data retrieval and backup. After the pool snapshot is made online, it is displayed in the pool list and its snapshot volumes are displayed in the volumes list.
- ◆ **Make Offline:** This option makes the selected pool snapshots and its associated volumes inaccessible through the pool list. It does not delete the data in the volumes.
- ◆ **Delete:** This option permanently deletes the selected pool snapshots from the server.



# 9 Managing NSS Volumes

This chapter describes the procedures for creating and managing NSS volumes on a server.

- ♦ “What is an NSS volume?” on page 73
- ♦ “What are the features that can be enabled for a new volume?” on page 73
- ♦ “What are the prerequisites for creating an encrypted volume with AES256?” on page 74
- ♦ “How to create a new NSS volume?” on page 75
- ♦ “How to list NSS volumes?” on page 77
- ♦ “How to view volume dashboard?” on page 78
- ♦ “How to deactivate and activate NSS volumes?” on page 78
- ♦ “How to mount or dismount a volume?” on page 80
- ♦ “How to rename a volume?” on page 81
- ♦ “How to delete a volume? Can I restore or permanently delete it?” on page 82
- ♦ “What is a volume object?” on page 83
- ♦ “How to update volume objects?” on page 83

## What is an NSS volume?

The logical volumes created on NSS storage pools are called NSS volumes. The **CREATE VOLUME** option on the **VOLUMES** page allows you to create an NSS volume in a pool. Depending on the physical space available, you can create any number of NSS volumes for each pool.

## What are the features that can be enabled for a new volume?

The below features can be enabled while creating a new volume.

### ♦ Salvage

The Salvage Files attribute enables deleted files to remain on the volume until the Purge Delay time expires or until space is needed on the volume for other data. Until the Purge Delay time expires, the Salvage feature tracks the deleted files and allows the deleted files to be salvaged and restored. If space is needed, the oldest deleted files are purged to clear space. Salvage is enabled by default. If the Salvage Files attribute is disabled, deleted files are purged immediately on deletion.

### ♦ User Quotas

The User Quotas (user space restrictions) attribute enables you to assign a maximum quota of space that a user’s data can consume across all directories in the volume.

#### ◆ **Directory Quotas**

The Directory Quotas attribute enables you to assign a maximum quota of space that a directory can consume.

#### ◆ **Active Directory**

This option lets you enable access to the AD users for the selected volume. For a volume (both NSS32 and NSS64) to be accessible to the AD users, it should be part of a pool that is AD media upgraded, and it should be AD-enabled.

#### ◆ **Compression**

The Compression attribute activates file compression in NSS volumes. Compression can be activated at creation time only and this choice persists for the life of the volume. Data in the volume is stored normally or in compressed form, depending on how frequently it is used. Compression parameters can be set at the server level to control compression behaviour.

#### ◆ **Encryption**

Encryption provides password-protected activation of encrypted NSS volumes. Encryption can be activated at creation time only, and this choice persists for the life of the volume.

#### ◆ **Event File List (EFL)**

NSS uses the Event File List (EFL) feature to track files that have changed on a volume during an interval called an epoch. It logs changes that are made to data and metadata for each active epoch on a specific NSS volume. You can use the API commands in scripts to start and stop an epoch, reset the event list for an epoch, and to affect how long epochs are retained.

---

**NOTE:** The feature, Event File List (EFL) is selected by default and you cannot deselect it.

---

## What are the prerequisites for creating an encrypted volume with AES256?

To create encrypted volumes with an AES-256 encryption algorithm, use the NSS64 pool type with pool media upgraded to AES. Use the `nsscon` commands in this section to upgrade the existing NSS media to support AES or to enable all future NSS pool creation to be automatically created with the AES Index support.

### **For the Existing NSS Pools**

```
nss /PoolMediaUpgrade=poolname /MediaType=AES
```

Upgrades the specified pool to support AES media.

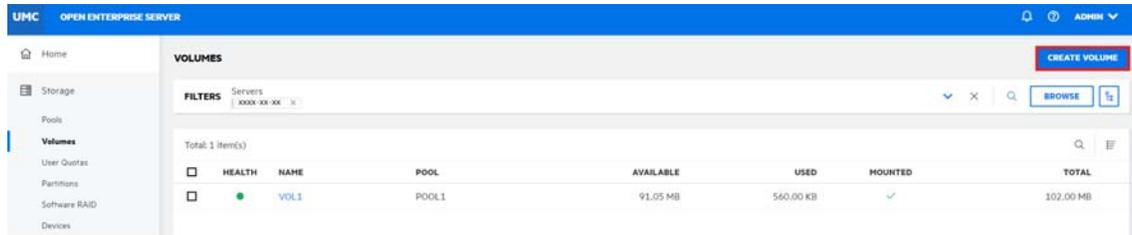
### **For the Newly Created NSS Pools**

The commands placed in the `nssstart.cfg` file persists across server reboots. If the NSS commands are added in the `nssstart.cfg` file, ensure those commands are not prefixed with `nss`.

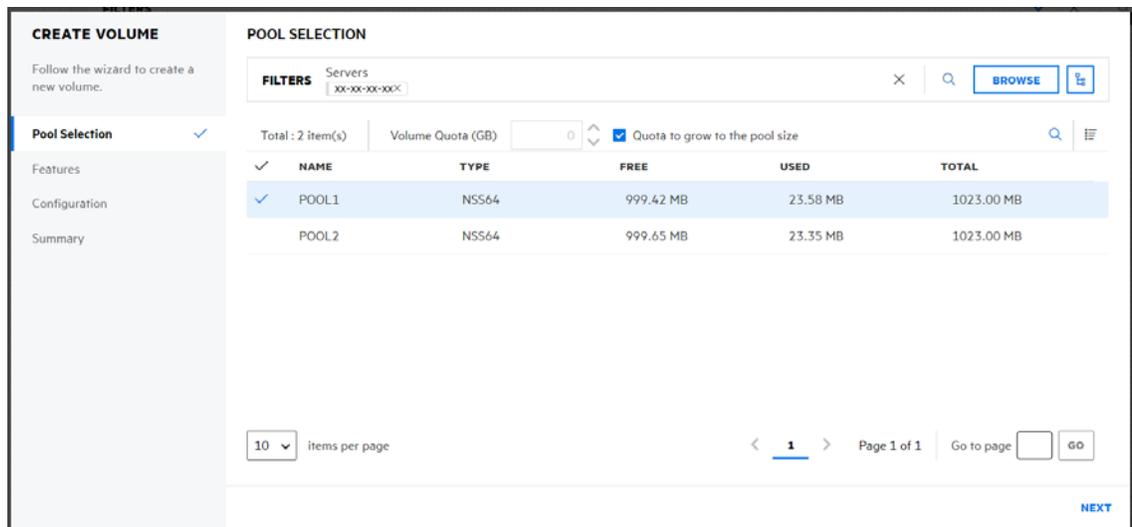
If these commands are issued from the command line, it persists only till a server reboot.

# How to create a new NSS volume?

- 1 In UMC, click  **Storage > Volumes**.
- 2 Click **CREATE VOLUME**.

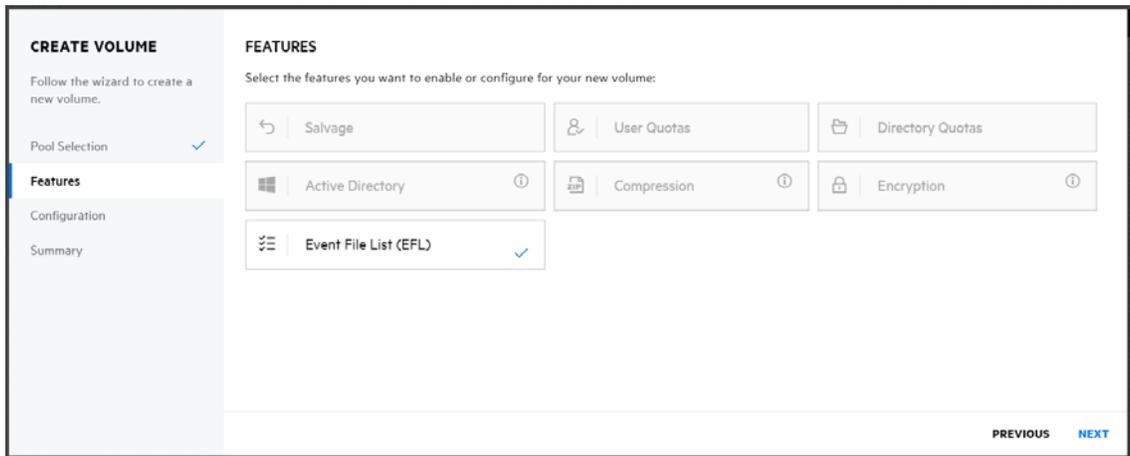


- 3 On the **POOL SELECTION** page, search or browse the server to select the pool where you want the new volume to reside, and click **NEXT**.



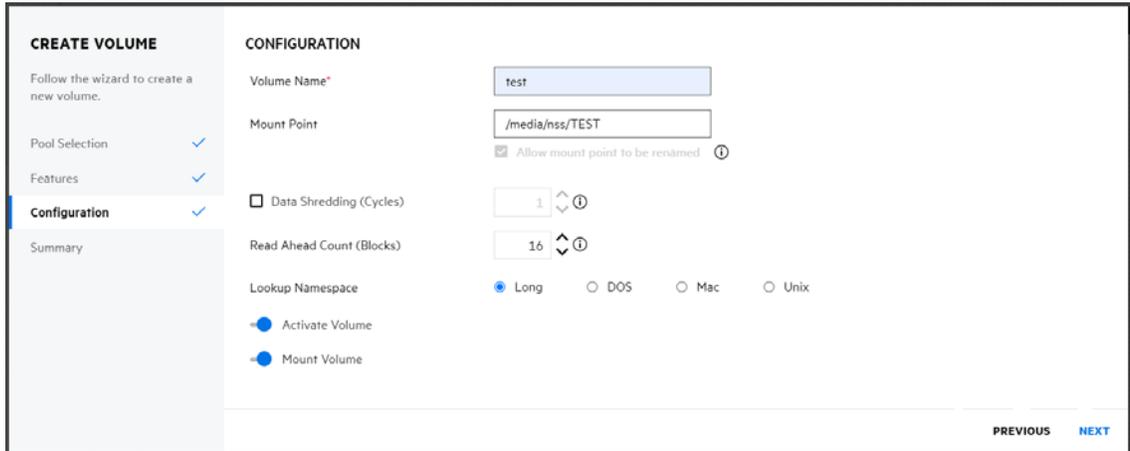
Specify the **Volume Quota** as required or select “**Quota to grow to the pool size**” checkbox to allow the volume to expand to the size of the pool.

- 4 On the **FEATURES** page, select the features you want to enable for the new volume, and click **NEXT**.

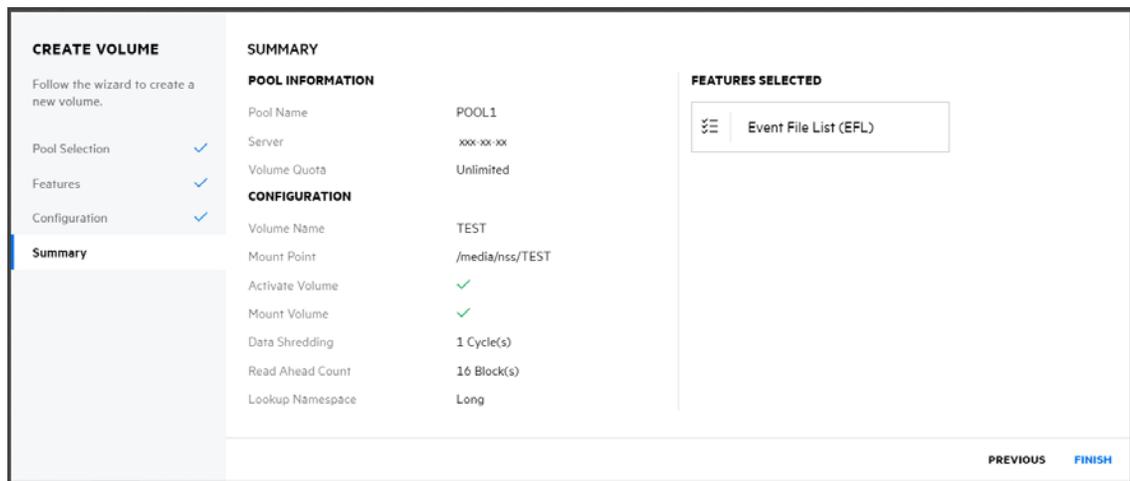


**NOTE:** The feature, Event File List (EFL) is selected by default and you cannot deselect it.

- 5 On the **CONFIGURATION** page, specify a name for the new volume, and click **NEXT**.  
Enable **Allow mount point to be renamed** to allow updates to the volume name or its path.



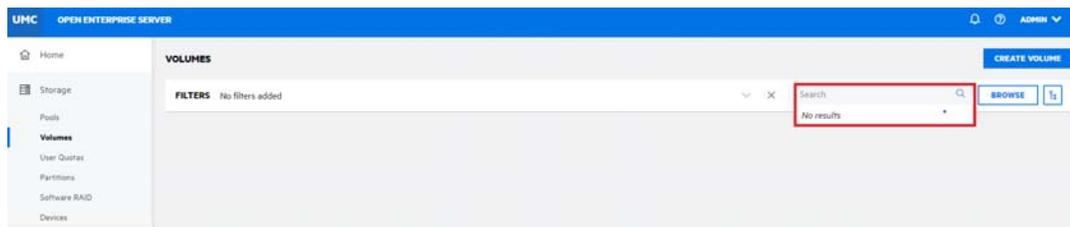
- 6 Review the details and click **FINISH**.



The new volume is available on the **VOLUMES** page.

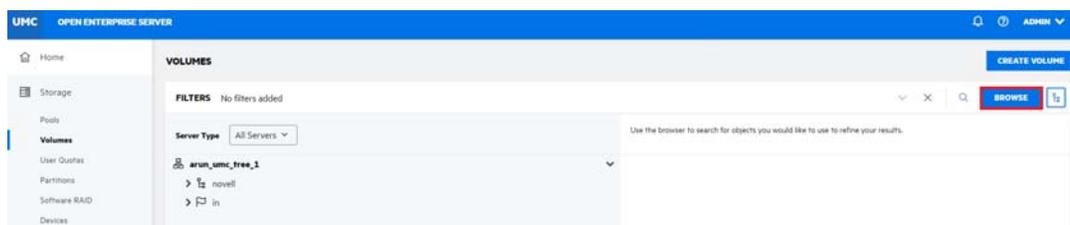
## How to list NSS volumes?

- 1 In UMC, click  **Storage > Volumes**.
- 2 Click the search Icon and specify the server name.



or

Click **Browse**, select **Server Type** to list the servers. Select the required servers from the list and click **APPLY**.

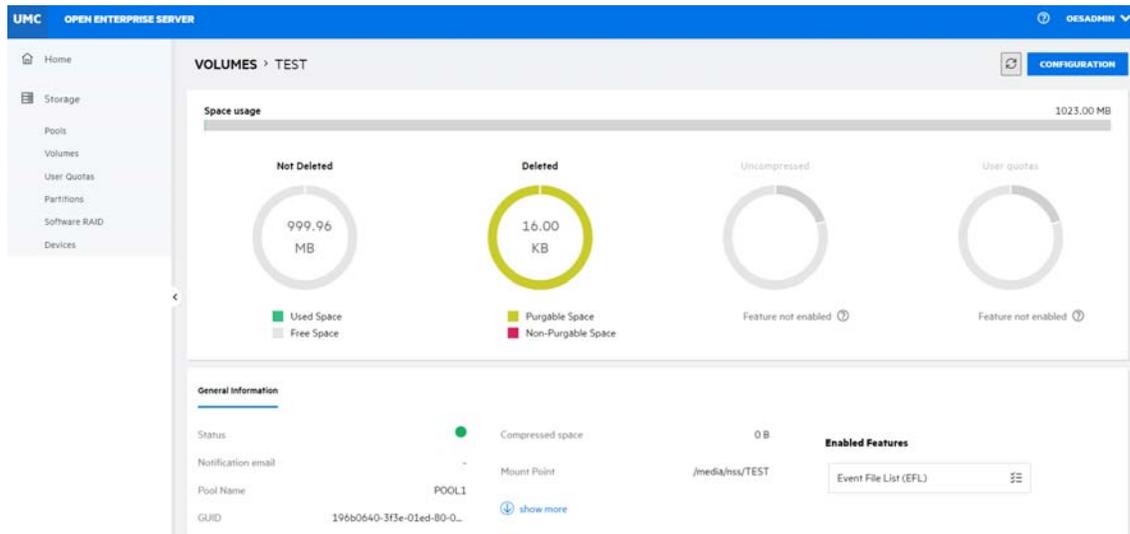


**NOTE:** When the **BROWSE** or tree view  icon is clicked, you cannot perform other actions outside the browse area. Click the same button again to close the browse or tree view area.

# How to view volume dashboard?

You can view the details of a volume like space usage, general information of the volume, and enabled features on the volume dashboard page.

- 1 In UMC, click  **Storage > Volumes**.
- 2 Search or browse the servers to list the volumes associated with them.
- 3 Select the volume, click More Options  icon, and then select **Dashboard**.



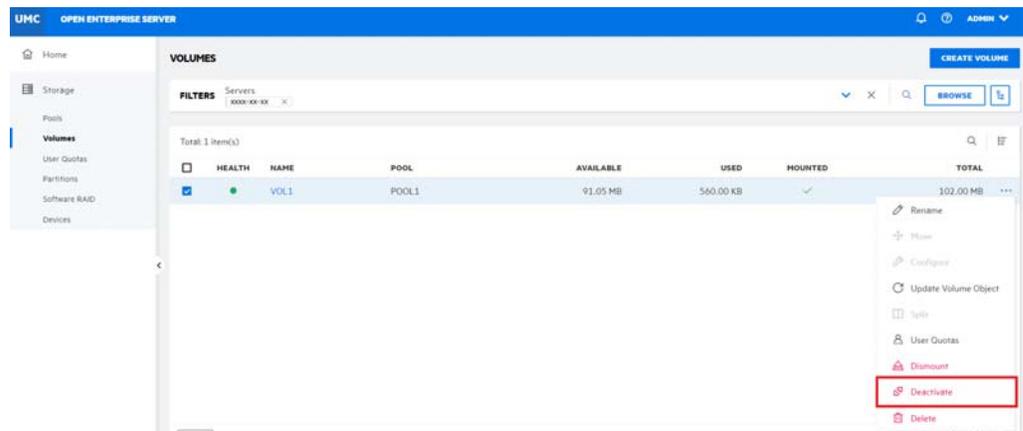
# How to deactivate and activate NSS volumes?

After configuring NSS volumes, you can activate and deactivate NSS volumes to make them available to users and applications. To view details of a volume, it must be active.

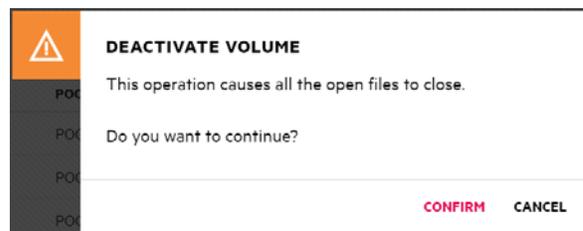
- 1 In UMC, click  **Storage > Volumes**.
- 2 Search or browse the servers to list the volumes associated with them.
- 3 **NOTE:** If you select multiple volumes, the More Options  icon is available at the top right corner of the table.

### 3a To Deactivate a Volume:

- 3a1 Select the volume, click More Options  icon, and then select **Deactivate**.



**3a2** Click **CONFIRM** to deactivate the selected volume.

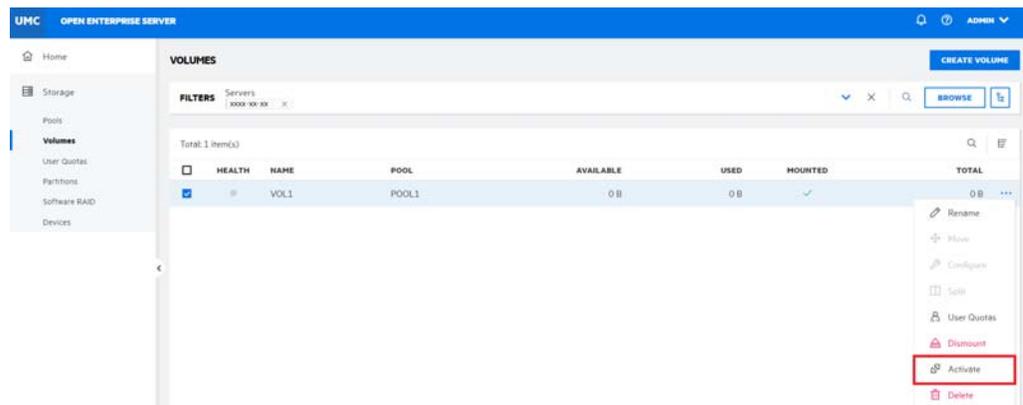


The deactivated volume details are not displayed on the **VOLUMES** page.

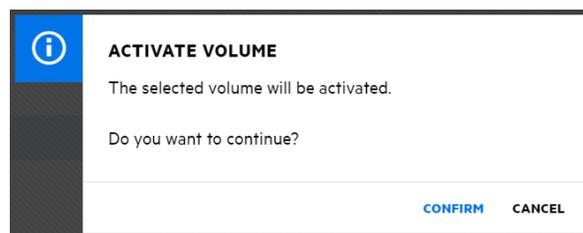
or

**3b To Activate a Volume:**

**3b1** Select the volume, click More Options **...** icon, and then select **Activate**.



**3b2** Click **CONFIRM** to activate the selected volume.



The activated volume details are displayed on the **VOLUMES** page.

After the page refreshes, each volume's state matches the state you specified. When a volume is already in the specified state, no change occurs.

## How to mount or dismount a volume?

After configuring the NSS volumes, you can mount and activate the volumes on the **VOLUMES** page to make them available to users and APIs. After mounting an NSS volume, it is only available to APIs until you activate it. Dismounting a volume makes it unavailable to the users and APIs.

---

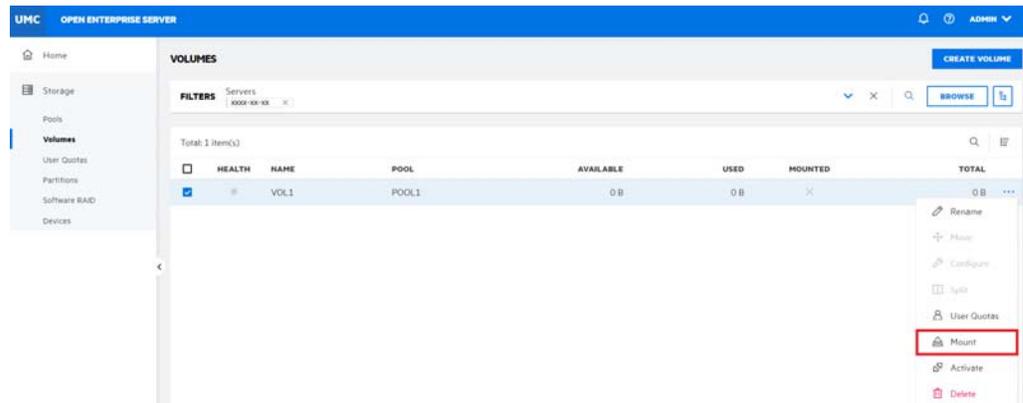
**NOTE:** If you **Mount** an encrypted volume, you are prompted for the related password.

---

- 1 In UMC, click  **Storage > Volumes**.
- 2 Search or browse the servers to list the volumes associated with them.
- 3 **NOTE:** If you select multiple volumes, the More Options  icon is available at the top right corner of the table.

### 3a To mount a Volume:

**3a1** Select the volume, click More Options  icon, and then select **Mount**.



**3a2** Click **CONFIRM** to mount the selected volume.



The page is refreshed and the **MOUNTED** state for the selected volume is changed to



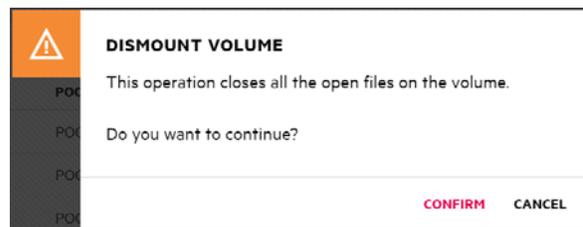
or

### 3b To Dismount a Volume:

3b1 Select the volume, click More Options  icon, and then select **Dismount**.



3b2 Click **CONFIRM** to dismount the selected volume.

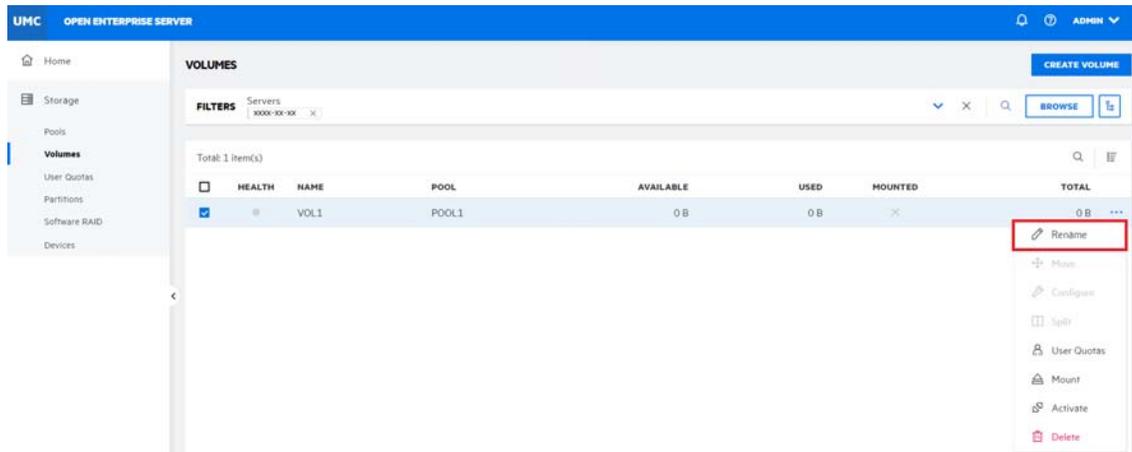


The page is refreshed and the **MOUNTED** state for the selected volume is changed to . The details of dismounted volumes are not displayed.

## How to rename a volume?

The **Rename** option on the **VOLUMES** page allows you to modify the name of the selected volume. For example, you want to change the name of a volume to reflect the department or organization name that uses it. Renaming a volume updates the corresponding eDirectory object.

- 1 In UMC, click  **Storage > Volumes**.
- 2 Search or browse the servers to list the volumes associated with them.
- 3 Select the volume to rename, click More Options  icon, and then select **Rename**.



4 Specify the new name for the volume, and then click **CONFIRM**.



The page is refreshed and the new volume name appears in the volumes list.

## How to delete a volume? Can I restore or permanently delete it?

Deleting a volume removes the data in the volume and frees the space to be used by other volumes in the same pool. When a volume is deleted, it is salvageable until Volume Purge Delay times out or you manually purge the deleted volumes.

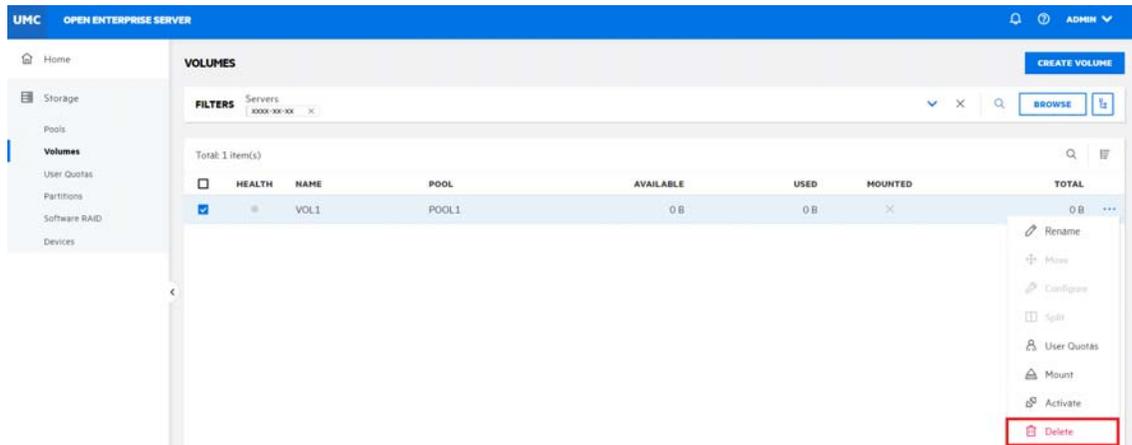
During the purge delay time, the deleted volume is salvageable, but the space belonging to the deleted volume is not available to other volumes. When the purging process begins, the volume is no longer salvageable.

- 1 In UMC, click  **Storage > Volumes**.
- 2 Search or browse the servers to list the volumes associated with them.
- 3 Select the volume, click More Options  icon, and then select **Delete**.

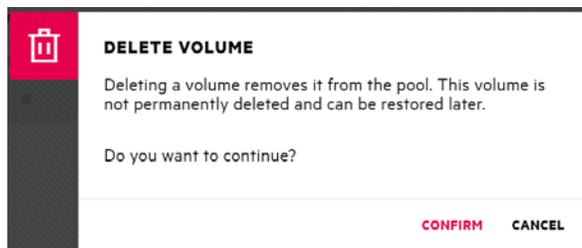
---

**NOTE:** If you select multiple volumes, the More Options  icon is available at the top right corner of the table.

---



4 Click **CONFIRM** to delete the selected volume.



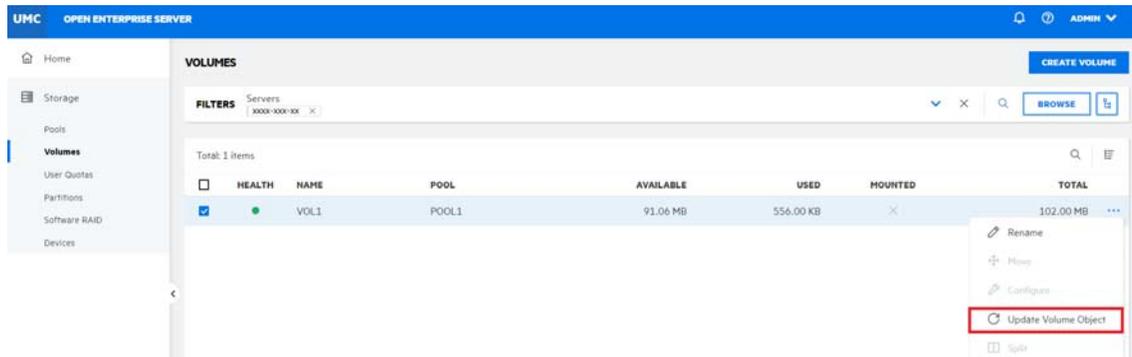
The deleted volumes are available in the **Deleted Volumes** list on the **Pools** page if the selected pool contains deleted volumes in it.

## What is a volume object?

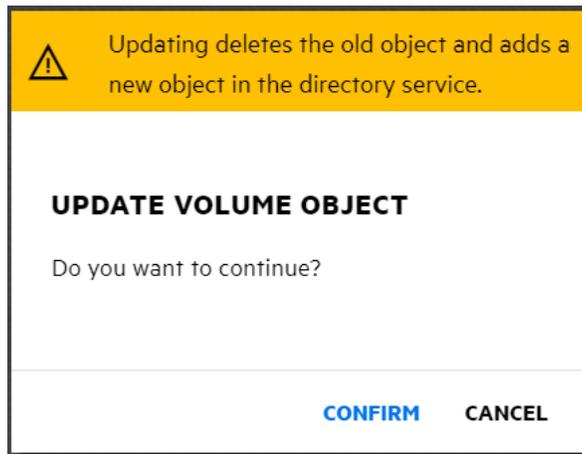
Volume objects represent a physical or logical volume on the network. Each NSS volume is represented by a volume object in eDirectory. The **Update Volume Object** option on the **Volumes** page allows you to add or replace a volume object for a volume at the same context level as the server.

## How to update volume objects?

- 1 In UMC, click  **Storage > Volumes**.
- 2 Search or browse the servers to list the volumes associated with them.
- 3 Select the volume, click More Options  icon, click **Modify**, and then select **Update Volume Object**.



4 Click **CONFIRM** to update the volume object of the selected volume.



If the volume object does not exist, NSS adds the volume object to the context level. If the volume object exists, NSS prompts to delete and replace the existing object or retain the existing object.

# 10 Managing User Quota

This chapter describes the procedure to view and manage user space restrictions for volumes on an OES server.

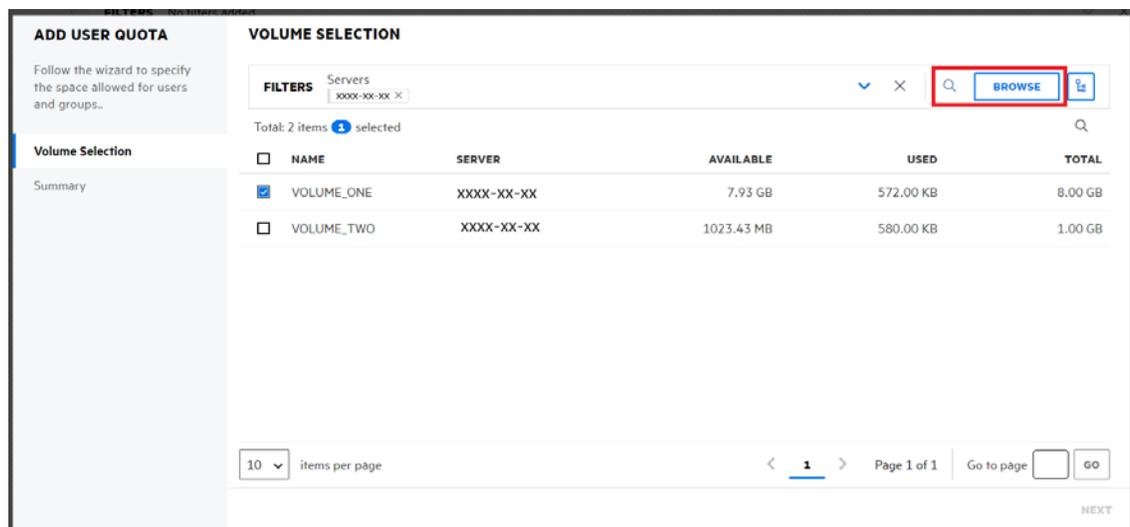
- ♦ “What are user quotas?” on page 85
- ♦ “How to add user quota?” on page 85
- ♦ “How to list user quotas?” on page 86
- ♦ “How to manage user quota?” on page 87
- ♦ “How to delete user quotas?” on page 88

## What are user quotas?

User Quotas are the space restrictions set for users of a volume on enabling the User Space Quotas attribute. A user quota specifies the maximum space a user’s data can consume in a volume. Space is allocated to users as required and user quota does not reserve the space for a user. You can overbook a volume and data can be set to grow to the size of the volume.

## How to add user quota?

- 1 In UMC, click  **Storage > User Quotas**, and click **ADD USER QUOTA**.
- 2 On the **VOLUME SELECTION** page, search or browse the servers, select the required volumes from the list, and then click **NEXT**.



**ADD USER QUOTA**

Follow the wizard to specify the space allowed for users and groups...

**VOLUME SELECTION**

FILTERS Servers [XXXX-XX-XX X] [BROWSE]

Total: 2 items selected

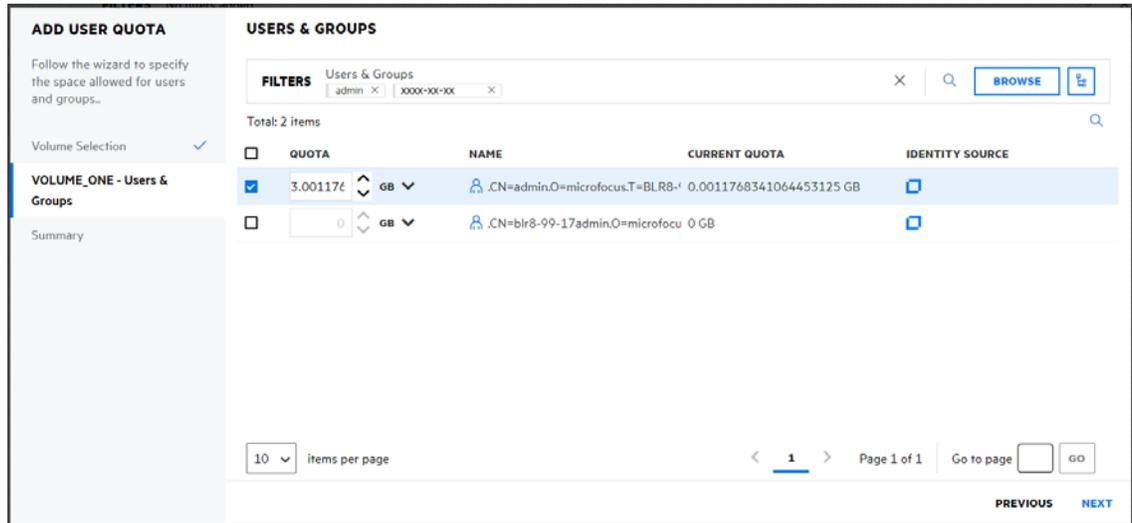
<input type="checkbox"/>	NAME	SERVER	AVAILABLE	USED	TOTAL
<input checked="" type="checkbox"/>	VOLUME_ONE	XXXX-XX-XX	7.93 GB	572.00 KB	8.00 GB
<input type="checkbox"/>	VOLUME_TWO	XXXX-XX-XX	1023.43 MB	580.00 KB	1.00 GB

10 items per page Page 1 of 1 Go to page [ ] GO

NEXT

- 3 On the **USERS & GROUPS** page, search or browse the users and groups to list the users.

- 4 Select the user(s), specify the storage space you want to assign to the selected user(s), and then click **NEXT**.

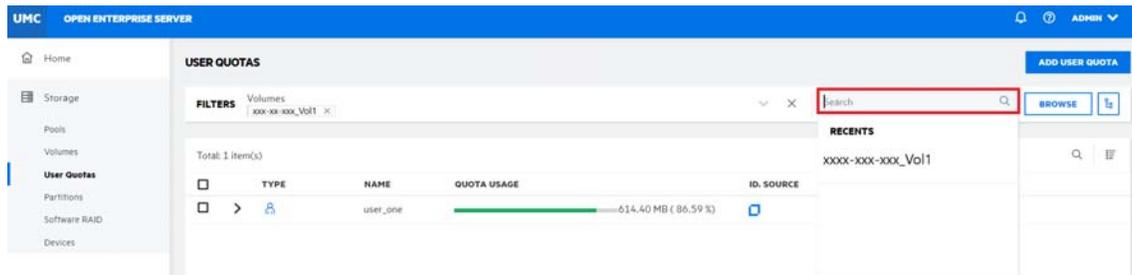


- 5 Review the details and click **FINISH**.

## How to list user quotas?

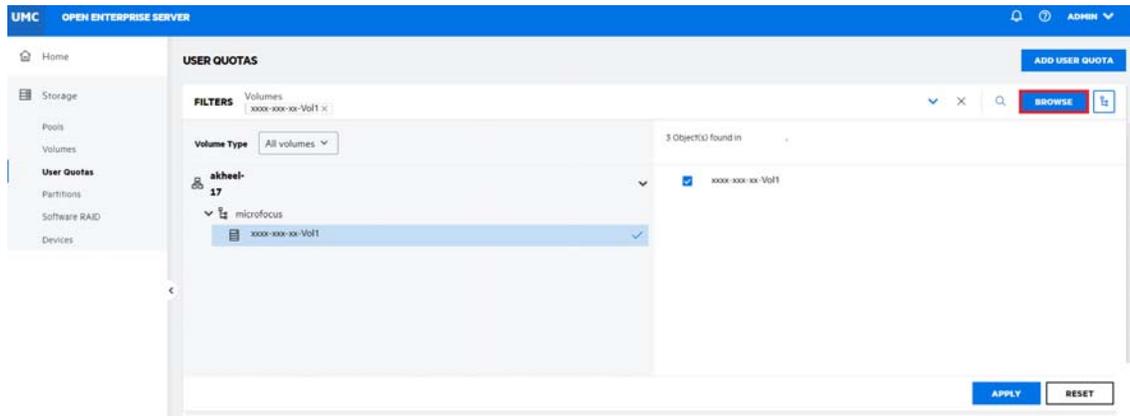
You can list the user quotas by selecting the volumes of a server.

- 1 In UMC, click  **Storage > User Quotas**.
- 2 Click the Search Icon and specify the volume name.



or

Click **Browse** and select **Server Type** to list the volumes. Select the required volumes and click **APPLY**.



**NOTE:** When the **BROWSE** or tree view icon is clicked, you cannot perform other actions outside the browse area. Click the same button again to close the browse or tree view area.

The list of the users with user quota assigned are displayed.

## How to manage user quota?

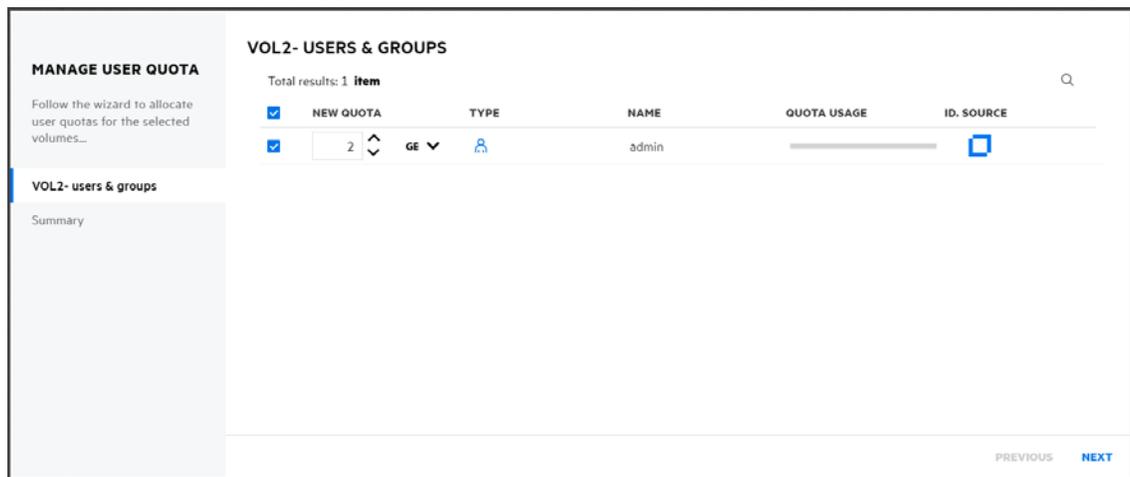
You can manage space restrictions of a user for a specific volume, regardless of whether the user has data on it or not.

- 1 In UMC, click **Storage > User Quotas**.
- 2 Search or browse the servers to list the volumes associated with them.
- 3 Select the required volumes from the list and click **APPLY**.
- 4 Select the user quota, click More Options **...** icon, and then select **Manage Quotas**.



**NOTE:** If you select multiple user quotas, the More Options **...** icon is available at the top right corner of the table.

- 5 On the **MANAGE USER QUOTA** page, specify the **NEW QUOTA** size, and click **NEXT**.



6 Review the details and click **FINISH**.

The new user quota is assigned to the user(s) for the selected volume(s).

## How to delete user quotas?

- 1 In UMC, click  **Storage > User Quotas**.
- 2 Search or browse the volumes to list the user quotas associated with them.
- 3 Select the user quota, click More Options  icon, and then select **Delete**.



**NOTE:** If you select multiple user quotas, the More Options  icon is available at the top right corner of the table.

- 4 Click **CONFIRM** to remove the user quota on the selected volume.

# 11 Managing NSS Partitions

This chapter describes the procedures for managing NSS partitions on a server.

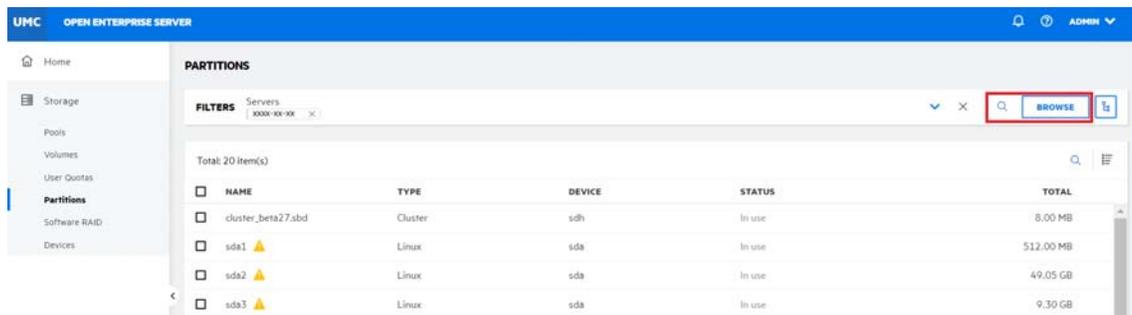
- ♦ “What is a partition?” on page 89
- ♦ “How to list NSS partitions?” on page 89
- ♦ “How to edit label of a partition?” on page 90
- ♦ “How to list volumes in a partition?” on page 90
- ♦ “What is NSS mirroring?” on page 91
- ♦ “How to mirror a partition?” on page 91
- ♦ “How to delete partitions?” on page 92

## What is a partition?

A partition is a logical division of a physical hard drive. NSS automatically creates the NSS partitions on the devices when you creates pools or RAID devices. You can view and label these NSS partitions from the **Partitions** page.

## How to list NSS partitions?

- 1 In UMC, click  **Storage > Partitions**.
- 2 Search or browse the servers to list the partitions associated with them.



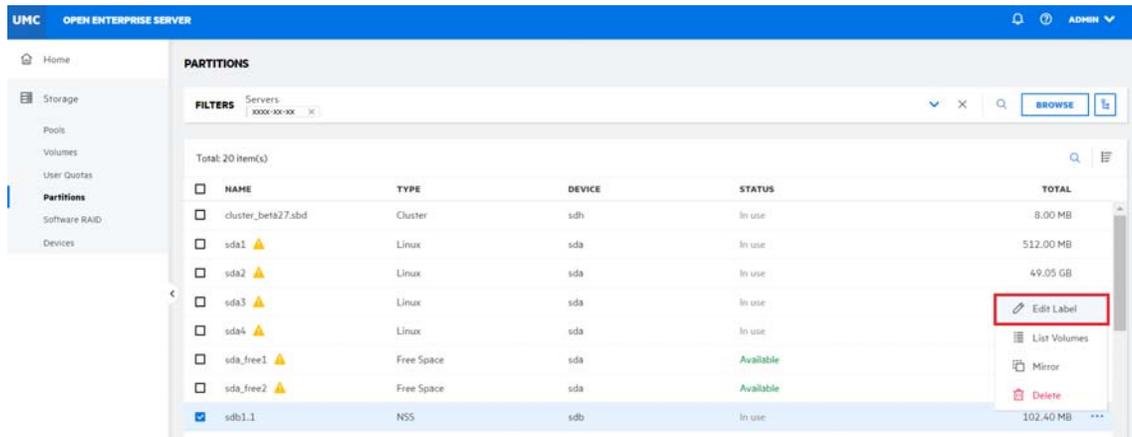
**NOTE:** When the **BROWSE** or tree view  icon is clicked, you cannot perform other actions outside the browse area. Click the same button again to close the browse or tree view area.

The actions that are performed on partitions are specific to partition type and vary depending on the type of partition selected.

# How to edit label of a partition?

A Label is the partition name assigned by the administrator and must be unique on a server. You can edit label of a partition by using the **Edit Label** option.

- 1 In UMC, click  **Storage > Partitions**.
- 2 Search or browse the servers to list the partitions associated with them.
- 3 Select the partition, click More Options  icon, and then select **Edit Label**.



- 4 Specify a new label to the partition and click **CONFIRM**.

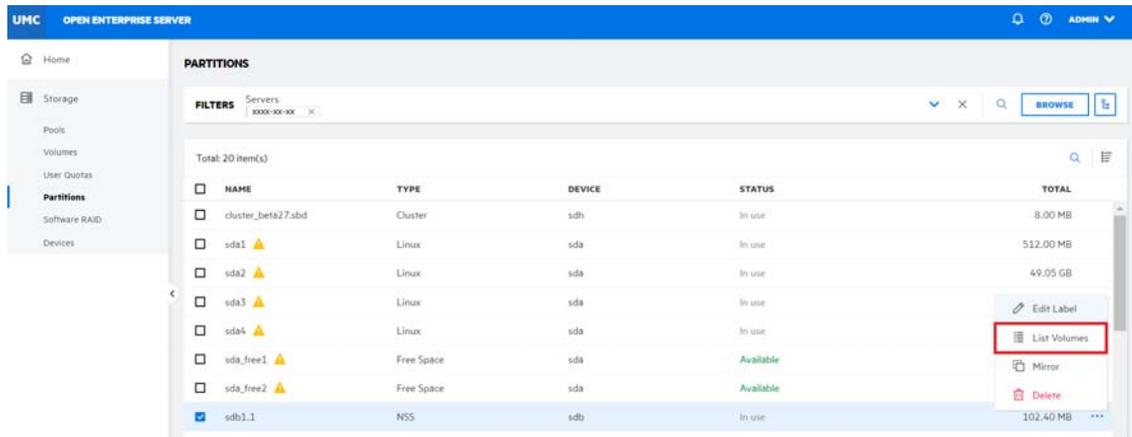


The 'EDIT LABEL' dialog box has a 'Label' field with the text 'Label Example' and 'CONFIRM' and 'CANCEL' buttons at the bottom.

The updated label is displayed in the partition list.

# How to list volumes in a partition?

- 1 In UMC, click  **Storage > Partitions**.
- 2 Search or browse the servers to list the partitions associated with them.
- 3 Select the partition, click More Options  icon, and then select **List Volumes**.



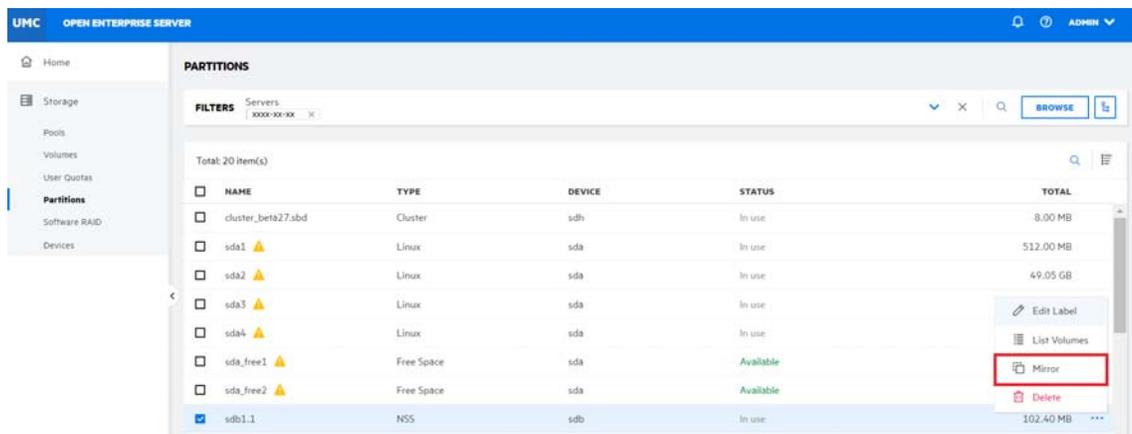
All the volumes available in the selected partition are listed.

## What is NSS mirroring?

NSS mirroring is a checkpoint-based synchronous mirroring solution. Data blocks are written synchronously to multiple storage devices. If the system crashes, data is still safe on the mirrored NSS volume on other servers.

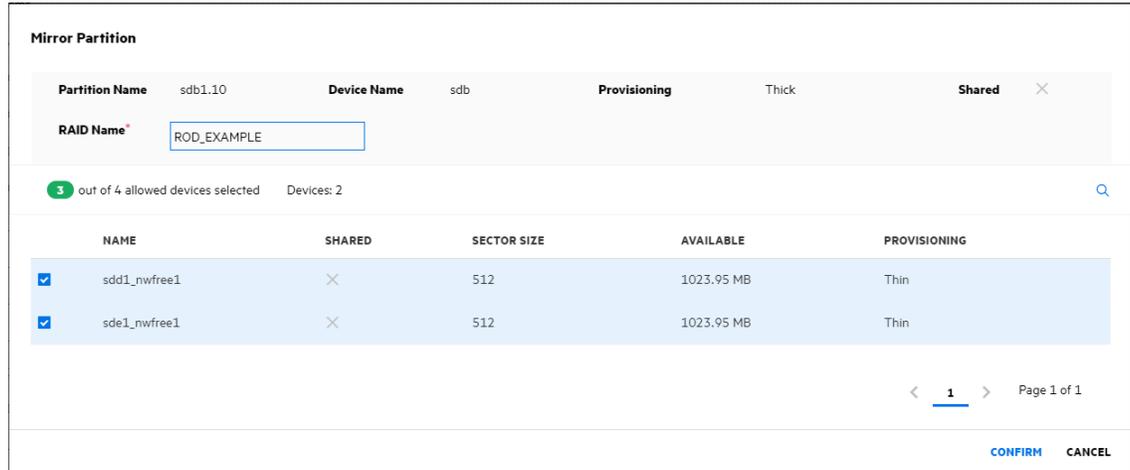
## How to mirror a partition?

- 1 In UMC, click  **Storage > Partitions**.
- 2 Search or browse the servers to list the partitions associated with them.
- 3 Select the partition, click More Options , and then select **Mirror**.



- 4 Specify the RAID name, select device(s) from the list, and then click **CONFIRM**.

**NOTE:** To create a mirror partition for a RAID device, the selected device(s) must have free space like the pool's size.

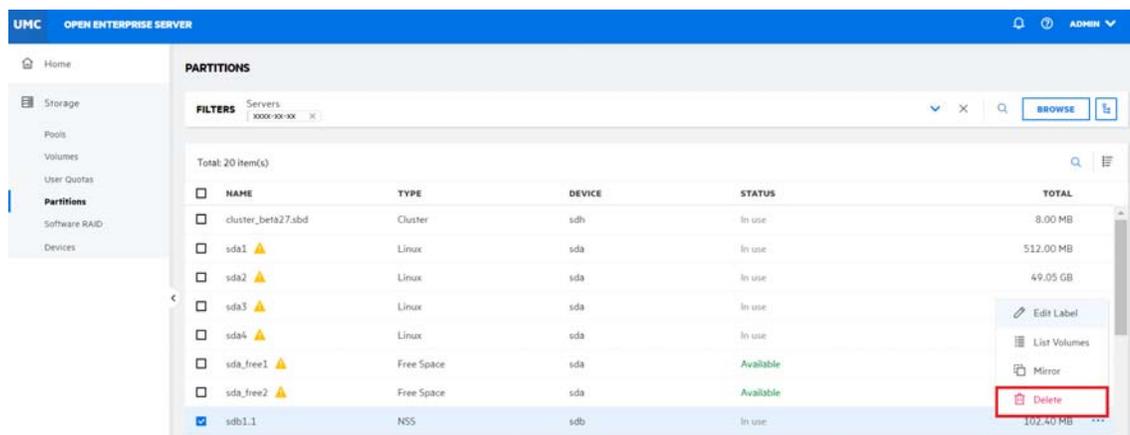


## How to delete partitions?

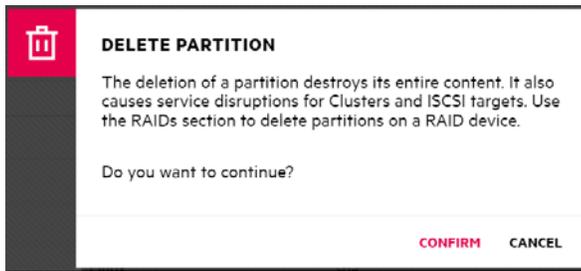
Deleting partitions deletes all the data in it. The delete option is only applicable for NSS partitions that are not part of RAID devices. For NSS software RAID devices, use the Software RAID's page to access and delete its partitions.

- 1 In UMC, click **Storage > Partitions**.
- 2 Search or browse the servers to list the partitions associated with them.
- 3 Select the partition, click More Options icon, and then select **Delete**.

**NOTE:** If you select multiple partitions, the More Options icon is available at the top right corner of the table.



- 4 Click **CONFIRM** to delete the selected partitions.



The selected NSS partitions are removed from the **PARTITIONS** list.



# 12 Managing NSS Software RAID Devices

This chapter describes the procedures for creating and managing NSS software RAID devices on a server.

- ♦ [“What is a software RAID?” on page 95](#)
- ♦ [“What RAIDs do NSS support?” on page 95](#)
- ♦ [“How to create a RAID device?” on page 96](#)
- ♦ [“How to list the RAID devices?” on page 98](#)
- ♦ [“How to view RAID device dashboard?” on page 98](#)
- ♦ [“How to rename a RAID device?” on page 99](#)
- ♦ [“How to increase the size of a RAID device?” on page 100](#)
- ♦ [“What happens when i delete a software RAID device?” on page 102](#)
- ♦ [“What happens when i delete a RAID1 device?” on page 102](#)
- ♦ [“How to delete a software RAID device?” on page 102](#)
- ♦ [“What is disk mirroring or remirroring?” on page 103](#)
- ♦ [“How to mirror or remirror a RAID 1 device?” on page 103](#)
- ♦ [“How to deactivate or activate a RAID device?” on page 104](#)

## What is a software RAID?

A software RAID is a configuration for storage devices that emulates a hardware RAID device. A software RAID combines partitioned space from multiple physical devices into a single virtual device that can be managed like any device. Each member device contributes an equal amount of space to the RAID. You can create partitions, pools, and volumes on a RAID device.

## What RAIDs do NSS support?

*Table 12-1 NSS supports three type of RAIDs.*

Type of RAID	Number of Partitions	Definition	Advantages	Disadvantages
RAID 0	2 to 14	Data striping	Improves storage performance	Does not provide data redundancy
RAID 1	2 to 4	Data mirroring	Provides data redundancy for failover and instant recovery	Does not improve performance; writes occur in parallel

Type of RAID	Number of Partitions	Definition	Advantages	Disadvantages
RAID 5	3 to 14	Data stripping with parity	Improves storage performance and enables limited data recovery.	Slightly degrades performance for writes to parity

## How to create a RAID device?

To set up a RAID device, you should allocate free space from any of your physical storage devices. NSS transparently presents the allocated free space as virtual partitions that represent NSS-managed physical partition areas on the participating drives.

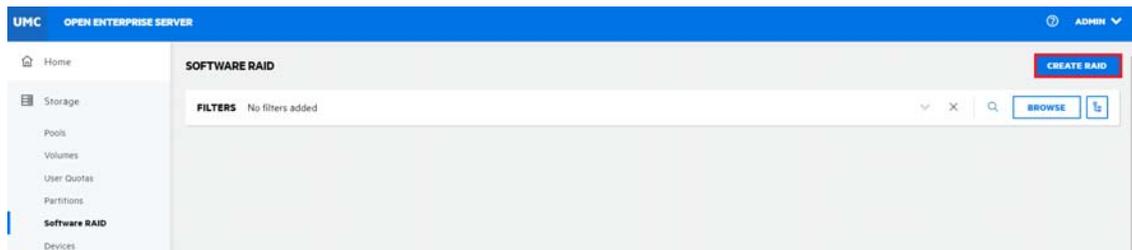
Partitions are the basic elements of a software RAID device. You can allocate partitions to the pools depending on the nature of the pools (shared or not shared for clustering) and the type of the RAID device.

### Guidelines for creating a software RAID device:

- ◆ Each partition in the RAID configuration should come from a different device. NSS lets you obtain RAID partitions from the same device, but this severely impedes the performance of your file system.
- ◆ Do not use space from a drive that contains your system partition (such as the root (/) or /boot partitions).
- ◆ You can use any combination of IDE or SCSI devices in a software RAID device. Make sure these devices have similar performance characteristics; otherwise, your performance might decrease.
- ◆ In a clustered solution using OES Cluster Services, for software RAIDs on shared disks:
  - ◆ You can have only one pool associated with that RAID device.
  - ◆ You must create an NSS pool and volume on that RAID device from the same server node before the pool can be migrated to other nodes in the cluster.

### To create a RAID device:

- 1 In UMC, click  **Storage > Software RAID**.
- 2 Click **CREATE RAID**.



- 3 On the **GENERAL INFORMATION** page, specify the RAID details, and click **NEXT**.

**CREATE RAID**

Follow the wizard to create a new RAID.

**General Information**

Devices

Summary

**GENERAL INFORMATION**

Name

Type RAID 0 ⓘ

Stripe Size (KB) 64 ⓘ

Partition Size (GB) 0 ⬆️⬆️

NEXT

- On the **DEVICES** selection page, select the server to list the devices associated with them. Select the devices to obtain space from, and click **NEXT**.

**NOTE:** If the specified partition size exceeds the amount of free space available on the physical devices, the RAID creation fails and returns an error message.

**CREATE RAID**

Follow the wizard to create a new RAID.

General Information ✓

**Devices**

Summary

**DEVICES**

**FILTERS** Servers xxxx-xx-xxxx X

Total: **4 item(s)** 3 of maximum 14 items Partition Size (GB) 1 ⬆️⬆️

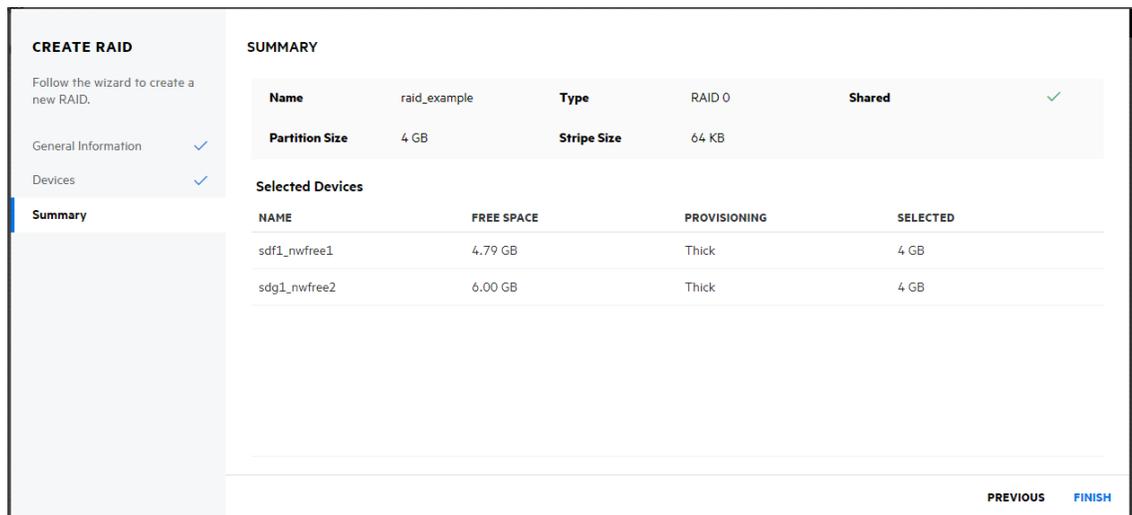
NAME	AVAILABLE	SHARED	SECTOR SIZE	PROVISIONING
<input checked="" type="checkbox"/> sdd1_nwfre...	2.00 GB	✓	512	Thick
<input checked="" type="checkbox"/> sde1_nwfre...	2.00 GB	✓	512	Thick
<input checked="" type="checkbox"/> sdf1_nwfre...	3.00 GB	✓	512	Thick
<input type="checkbox"/> sdg1_nwfre...	3.00 GB	✓	512	Thick

10 items per page

< **1** > Page 1 of 1 | Go to page  **GO**

PREVIOUS **NEXT**

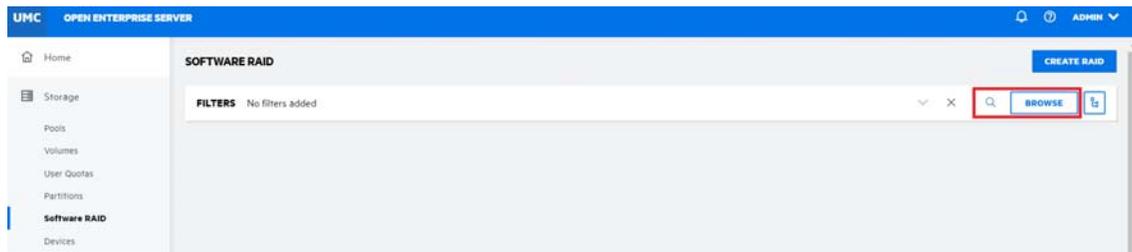
- Review the details and click **FINISH**.



The **Software RAID** page displays the newly created RAID device.

## How to list the RAID devices?

- 1 In UMC, click  **Storage > Software RAID**.
- 2 Search or browse the servers to list the RAID devices associated with them.



**NOTE:** When the  or tree view  icon is clicked, you cannot perform other actions outside the browse area. Click the same button again to close the browse or tree view area.

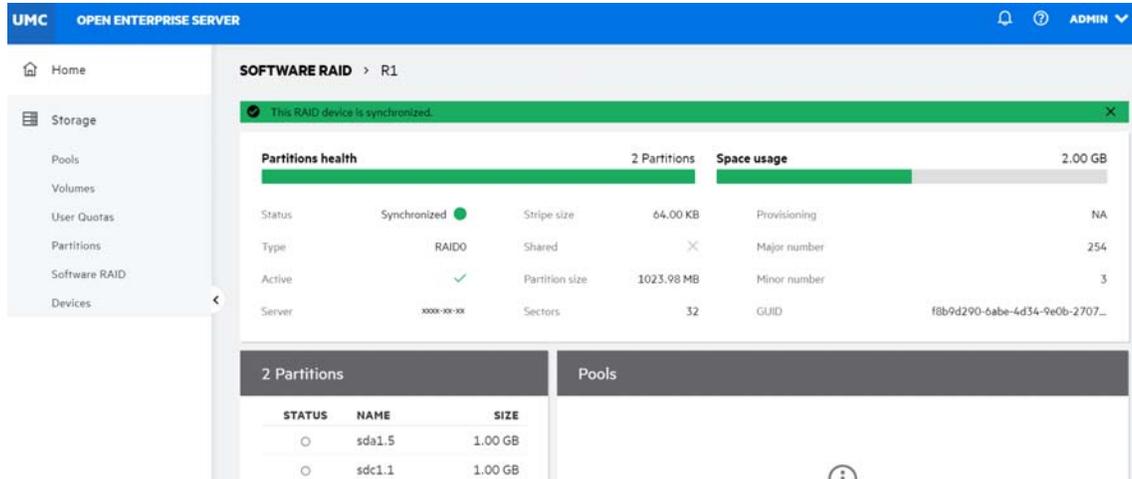
The list of the RAID devices available on the selected servers are displayed.

## How to view RAID device dashboard?

You can view the details of a RAID device like partition health, space usage, partitions, pools and general information on the **SOFTWARE RAID** dashboard page.

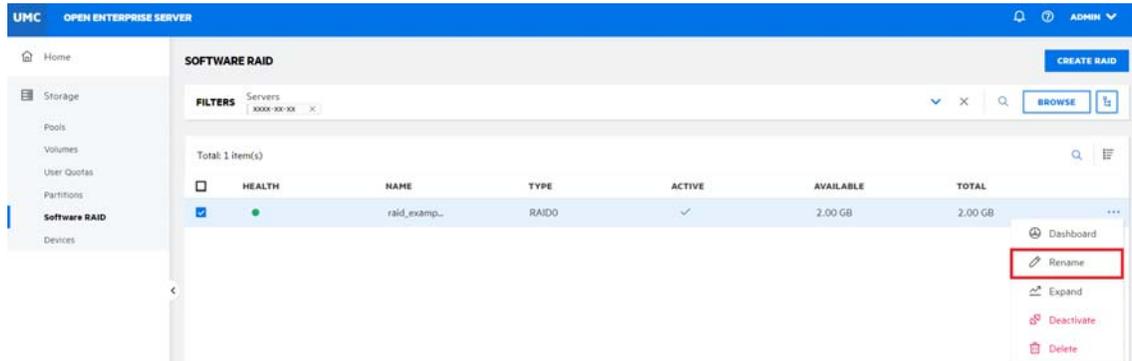
- 1 In UMC, click  **Storage > Software RAID**.
- 2 Search or browse the servers to list the RAID devices associated with them.

- 3 Select the RAID device, click More Options  icon, and then select **Dashboard**.



## How to rename a RAID device?

- 1 In UMC, click  **Storage > Software RAID**.
- 2 Search or browse the servers to list the RAID devices associated with them.
- 3 Select the RAID device, click More Options  icon, and then select **Rename**.



- 4 Specify a new name and click **CONFIRM**.

**RENAME RAID**

Existing Name    Example 3

New Name \*   

CONFIRM    CANCEL

The selected software RAID device is listed with its new name.

# How to increase the size of a RAID device?

You can increase the capacity of an existing software RAID 0, 1 or 5 device by adding partitions up to the maximum number for the type of RAID. You cannot modify the size of an individual partition after the device is created. The partition size is predetermined by the existing RAID.

Partitions can be added only if they match the shared state of current member devices. They must be all local or all shared and cannot be mixed.

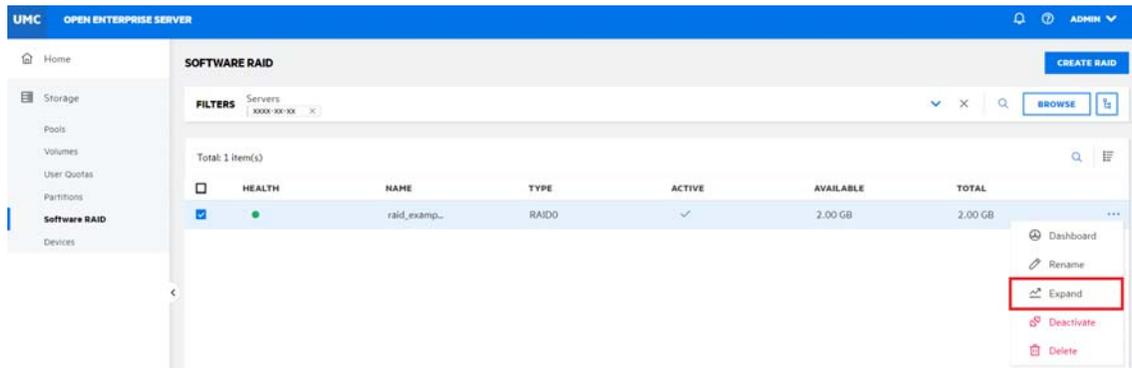
---

**IMPORTANT:** If the software RAID device is shared in a cluster, connect to the node where the RAID is currently active to manage the RAID and increase the size of the RAID.

---

## To add partitions on a RAID device:

- 1 In UMC, click  **Storage > Software RAID**.
- 2 Search or browse the servers to list the RAID devices associated with them.
- 3 Select the RAID device, click More Options  icon, and then select **Expand**.



If the software RAID device contains the maximum number of partitions, the **Expand** option is disabled.

- 4 On the **Devices** selection page, select the device(s), and click **NEXT**.

**Expand - RAID Device**

Follow the wizard to expand a RAID.

**Devices**

Summary

**Devices**

<b>Type</b>	RAID0	<b>Provisioning</b>	NA	<b>Shared</b>	Yes
<b>Partition Size</b>	102.38 MB	<b>Stripe Size</b>	64.00 KB	<b>Sector Size</b>	512 B

Total: 7 item(s) +2 of maximum 14 items

	NAME	IN USE	AVAILABLE	SHARED	SECTOR	PROVISIONING
<input type="checkbox"/>	sdb1_nwfree1	✓	29.90 GB	✓	512 B	Thin
<input type="checkbox"/>	sdcl1_nwfree1	✓	31.90 GB	✓	512 B	Thin
<input type="checkbox"/>	sdd1_nwfree1	✓	15.90 GB	✓	512 B	Thin
<input type="checkbox"/>	sde1_nwfree1	✓	15.90 GB	✓	512 B	Thin
<input checked="" type="checkbox"/>	sdf1_nwfree1	✗	3.00 GB	✓	512 B	Thick
<input checked="" type="checkbox"/>	sdg1_nwfree1	✗	5.00 GB	✓	512 B	Thick

10 items per page Page 1 of 1  GO

NEXT

The wizard lets you select the partitions with free space to meet the RAID's current partition size and are not members of the RAID.

5 Review the details and click **FINISH**.

**Expand - RAID Device**

Follow the wizard to expand a RAID.

Devices ✓

**Summary**

**Summary**

<b>Type</b>	RAID0	<b>Provisioning</b>	NA	<b>Shared</b>	Yes
<b>Partition Size</b>	102.38 MB	<b>Stripe Size</b>	64.00 KB	<b>Sector Size</b>	512 B

**Selected Devices**

NAME	FREE SPACE	PROVISIONING	SELECTED
sdf1_nwfree1	3.00 GB	Thick	102.38 MB
sdg1_nwfree1	5.00 GB	Thick	102.38 MB

PREVIOUS **FINISH**

The selected partitions are added to the RAID device, increasing its size.

# What happens when i delete a software RAID device?

Deleting a software RAID device removes the RAID relationship between the member partitions and the underlying storage structures. All data on the member partitions is deleted and cannot be restored. Before deleting the software RAID device, backup your data or move it to a different location if required.

# What happens when i delete a RAID1 device?

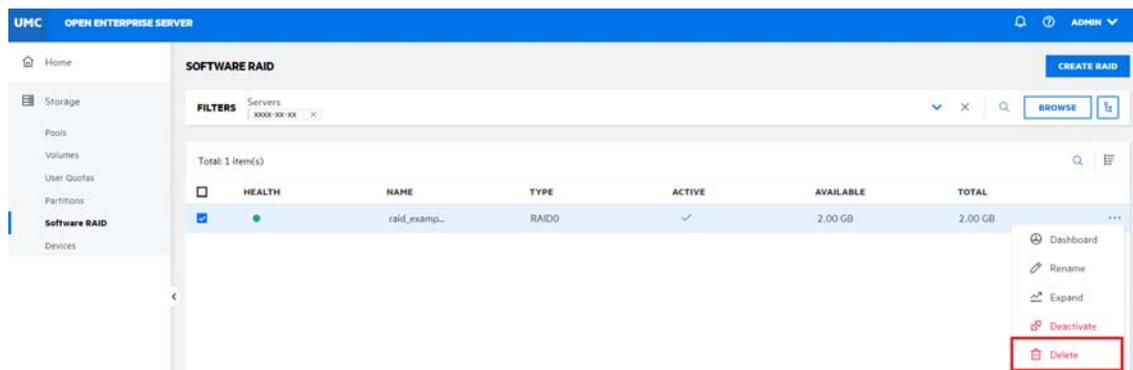
No data is lost and only the RAID1 device gets deleted under the following scenarios:

- When the RAID1 device has only one segment and if the device is consumed by a pool, deleting the RAID1 device deletes only the device. The segment is directly attached to the pool.
- When the RAID1 device has only one segment and if the device is an SBD mirror, deleting the RAID1 device deletes only the mirror. The mirror's segment becomes the SBD partition.

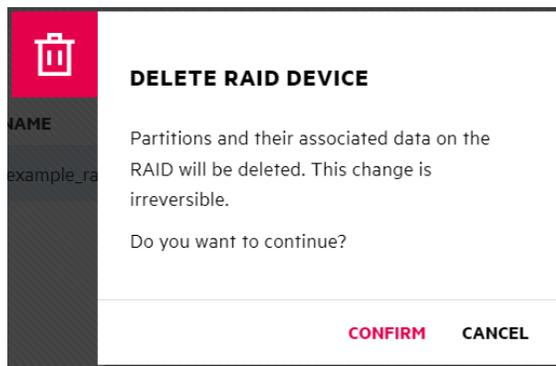
# How to delete a software RAID device?

- 1 In UMC, click  **Storage > Software RAID**.
- 2 Search or browse the servers to list the RAID device associated with them.
- 3 Select the RAID device, click More Options  icon, and then select **Delete**.

**NOTE:** If you select multiple RAID devices, the More Options  icon is available at the top right corner of the table.



- 4 Click **CONFIRM** to delete the selected RAID device.



The deleted software RAID device is not accessible from the **SOFTWARE RAID** page.

## What is disk mirroring or remirroring?

Disk mirroring or remirroring is the replication of data to two or more disks. Disk mirroring is a good choice for applications that require high performance and high availability. Disk mirroring or remirroring a RAID 1 device creates a copy of the data contained in that device.

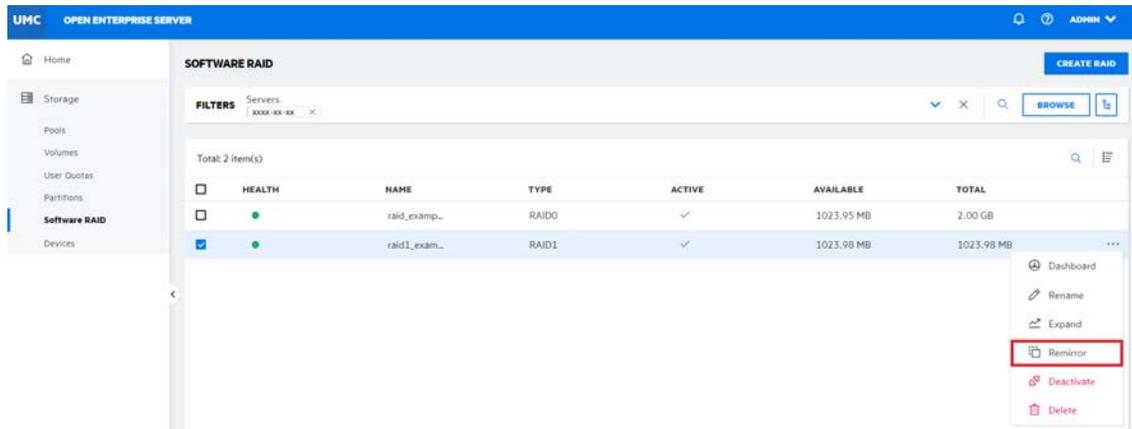
## How to mirror or remirror a RAID 1 device?

### Requirements for mirroring a software RAID 1 devices:

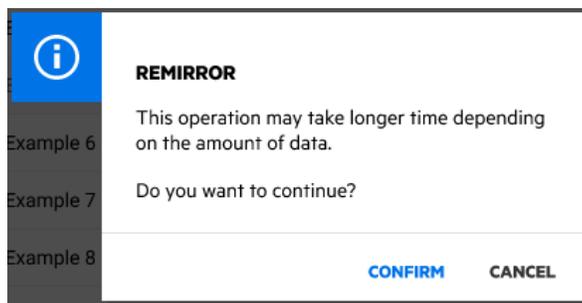
- ♦ Mirrored partitions must have the same partition type: NSS partitions to NSS partitions and traditional partitions to traditional partitions.
- ♦ Mirrored partitions should be set up on devices that have similar performance thresholds.
- ♦ You can mirror only partitions, each from its own OES partition. If a storage pool spans multiple devices, each of the individual partitions that make up that pool can be mirrored independently. The pool's partitions must be mirrored for the data in that pool to be fault tolerant.

### To Remirror a RAID 1 device:

- 1 In UMC, click  **Storage > Software RAID**.
- 2 Search or browse the servers to list the RAID device associated with them.
- 3 Select the RAID device, click More Options  icon, and then select **Remirror**.



4 Click **CONFIRM** to remirror the selected RAID device.



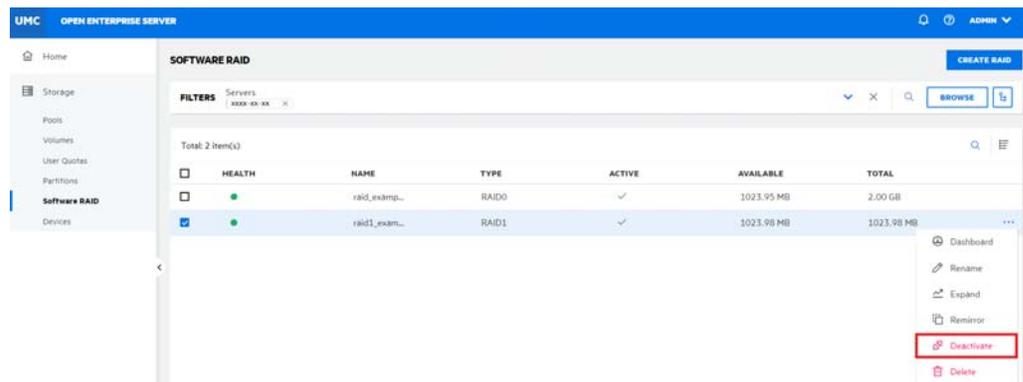
## How to deactivate or activate a RAID device?

You can activate and deactivate a RAID device to make them available to users. To view details of a RAID device, it must be active.

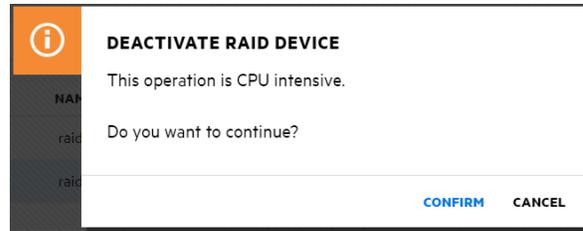
- 1 In UMC, click **Storage > Software RAID**.
  - 2 Search or browse the servers RAID to list the RAID devices associated with them.
- 
- 3 **NOTE:** Only one RAID device can be deactivated or activated at a time.

### 3a To Deactivate a RAID Device:

3a1 Select the RAID device, click More Options icon, and then select **Deactivate**.



**3a2** Click **CONFIRM** to deactivate the selected RAID device.

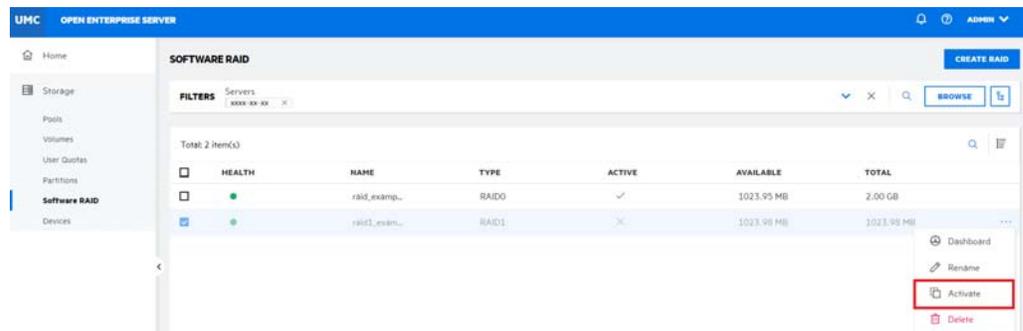


The deactivated RAID device details are not displayed on the **Software RAID** page.

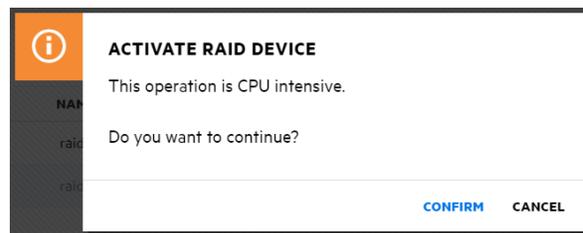
OR

**3b To Activate a RAID Device:**

**3b1** Select the RAID device, click More Options **⋮** icon, and then select **Activate**.



**3b2** Click **CONFIRM** to activate the selected RAID device.



The activated RAID device details are displayed on the **Software RAID** page.

After the page refreshes, each RAID device's state matches the state you specified. When a RAID device is already in the specified state, no change occurs.



# 13 Managing Devices

This chapter describes the procedures to manage devices connected to the servers.

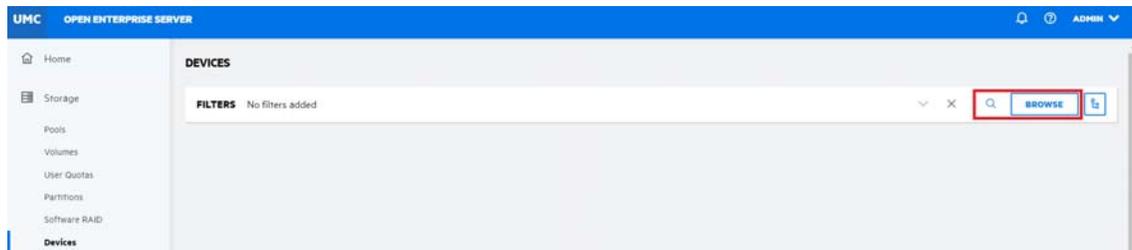
- ♦ “What is a device?” on page 107
- ♦ “How to list devices connected to the servers?” on page 107
- ♦ “What happens when a device is initialized?” on page 108
- ♦ “What happens when a device is shared?” on page 108
- ♦ “How can i initialize a device connected to a server?” on page 108
- ♦ “Why I need to reinitialize a device?” on page 109
- ♦ “How can i reinitialize a device?” on page 109
- ♦ “How to share or unshare an initialized device?” on page 110

## What is a device?

A device is a physical or virtual storage media available to a server. A device is directly attached to the server or connected via storage networking protocols.

## How to list devices connected to the servers?

- 1 In UMC, click  Storage > Devices.
- 2 Search or browse the servers to list the devices associated with them.



---

**NOTE:** When the  or tree view  icon is clicked, you cannot perform other actions outside the browse area. Click the same button again to close the browse or tree view area.

---

The list of available devices is displayed.

# What happens when a device is initialized?

Initializing a device deletes the partitions and its associated data. If the pool on this device has partitions on other devices, then the entire pool is deleted from those devices.

# What happens when a device is shared?

Sharing a device containing pools set all the pools on the device to shareable. If any of these pools span multiple devices, ensure that each device has the same share setting, otherwise the entire pool may become unusable.

Making a device shareable enables device sharing for those devices in high-availability clusters that you want to be part of a shared-disk storage solution. If the **Shareable for Clustering** option is enabled, the selected storage device can be shared by multiple computers in a cluster.

If a device is a member of a software RAID device, marking the device as shareable for clustering automatically sets all the other member devices of the RAID as shareable for clustering.

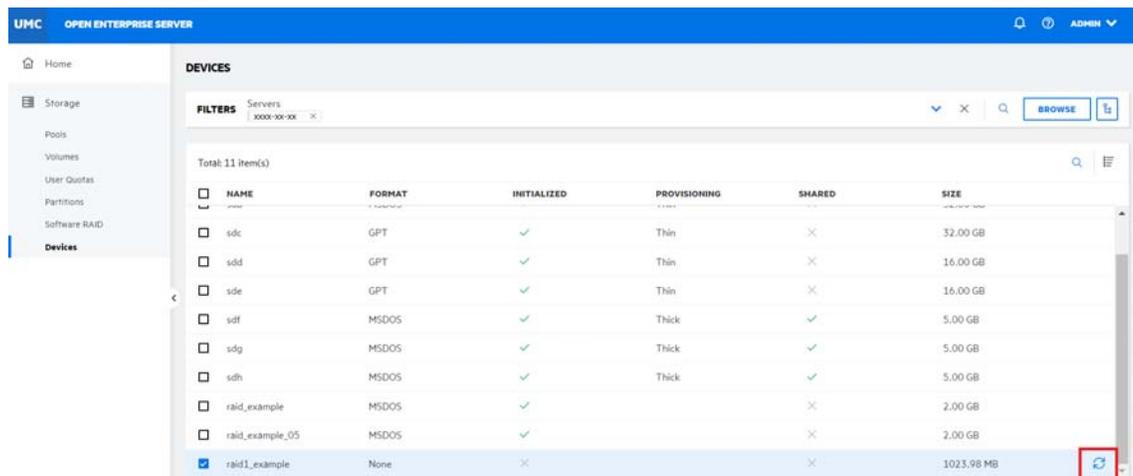
# How can i initialize a device connected to a server?

---

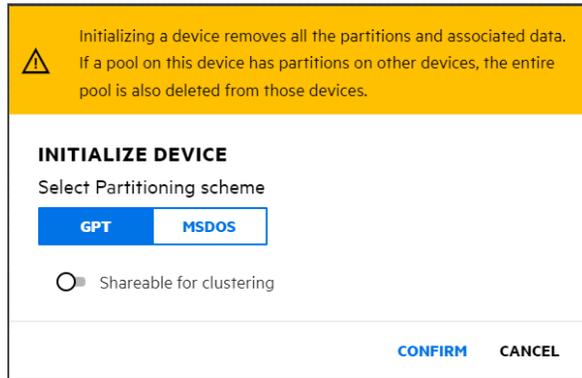
**WARNING:** Do not initialize the device that contains the operating system.

---

- 1 In UMC, click  **Storage > Devices**.
- 2 Search or browse the servers to list the devices associated with them.
- 3 Select the device, click More Options  icon, and then click **Initialize Icon**.



- 4 Select the partitioning scheme, click **Shareable for clustering** option for device sharing if required, and then click **CONFIRM**.
  - ◆ You can select the DOS partition table scheme that supports devices up to 2TB in size. It allows up to four partitions on a device.
  - ◆ You can select the GPT partition table scheme that supports device size up to 2E64 sectors (that is, up to 8388608petabytes (PB) based on the 512-byte sector size). It allows up to 128 partitions per disk. Each of its disks partitions is a logical device that is identified by a unique 128-bit (16-byte) GUID.



The status of the initialized device reflects in the device list.

## Why I need to reinitialize a device?

You can reinitialize an already initialized device if it is unusable. A device's reinitialization is an operation to clean up a device to start fresh in the event of a corruption or other similar event.

## How can i reinitialize a device?

---

**NOTE:** The **Reinitialize** option is available only if the device is already initialized.

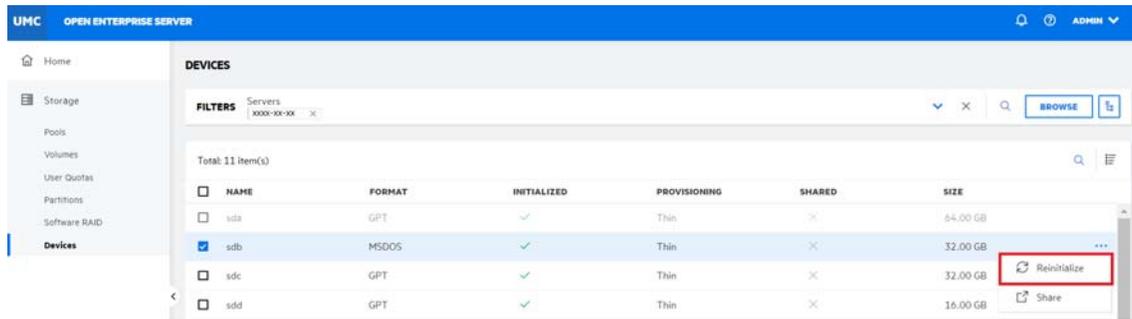
---

- 1 In UMC, click  **Storage > Devices**.
- 2 Search or browse the servers to list the devices associated with them.
- 3 Select an initialized device, click More Options  icon, and then select **Reinitialize**.

---

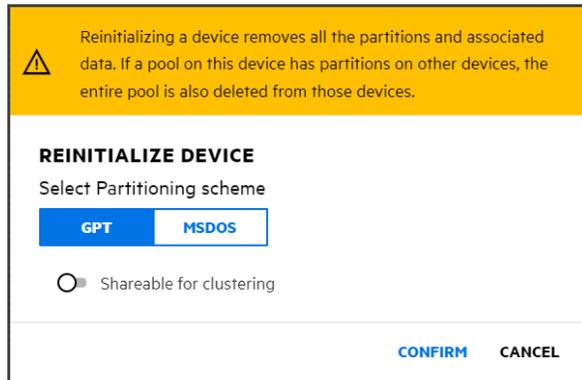
**NOTE:** If you select multiple devices, the More Options  icon is available at the top right corner of the table.

---



4 Select the partitioning scheme, click **Shareable of clustering** option if required, and then click **CONFIRM**.

- ◆ You can select the DOS partition table scheme that supports devices up to 2TB in size. It allows up to four partitions on a device.
- ◆ You can select the GPT partition table scheme that supports device sizes up to 2E64 sectors (that is, up to 8388608petabytes (PB) based on the 512-byte sector size). It allows up to 128 partitions per disk. Each of its disks partitions is a logical device that is identified by a unique 128-bit (16-byte) GUID.

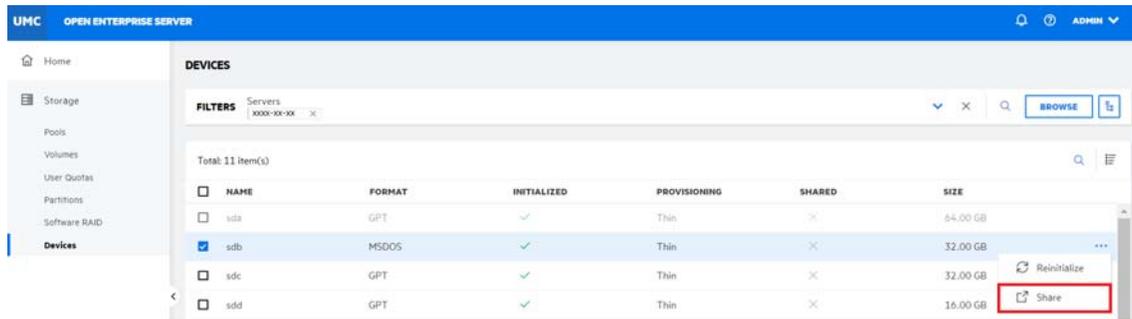


The status of the reinitialized device reflects in the device list.

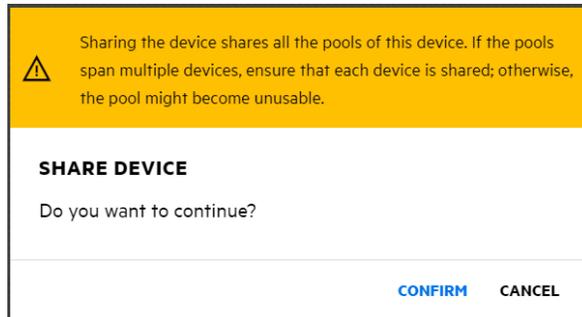
## How to share or unshare an initialized device?

- 1 In UMC, click **Storage > Devices**.
- 2 Search or browse the servers to list the devices associated with them.
- 3 Select an initialized device, click More Options icon, and then select **Share**.

**NOTE:** If you select multiple devices, the More Options icon is available at the top right corner of the table.



4 Click **CONFIRM** to share the selected device.



The status of the selected device reflects in the device list.

---

**NOTE:** The device is unshared following the same procedure. Unsharing a device fails if the device contains a pool (or any segment of a pool) that is cluster-enabled.

---



# V Files and Folders

- ◆ Chapter 14, “Managing Files and Folders,” on page 115
- ◆ Chapter 15, “Managing Rights,” on page 123



# 14 Managing Files and Folders

You can create, list, move, change owner, set directory quota, modify attributes, rename, move, delete, list deleted files, salvage, and purge deleted files here.

- ♦ [“How to view files and folders?” on page 115](#)
- ♦ [“How to create a new folder?” on page 115](#)
- ♦ [“How to modify properties of a volume, file, or folder?” on page 116](#)
- ♦ [“How to modify directory quota of a volume or folder?” on page 118](#)
- ♦ [“How to modify owner of a volume, file, or folder?” on page 118](#)
- ♦ [“How to modify attributes of a volume, file, or folder?” on page 119](#)
- ♦ [“How to view deleted files and folders?” on page 120](#)
- ♦ [“How to delete files and folders?” on page 120](#)
- ♦ [“How to salvage the deleted files and folders?” on page 120](#)
- ♦ [“How to purge the files and folders?” on page 121](#)
- ♦ [“How to rename a file or folder?” on page 121](#)
- ♦ [“How to move files and folders in a volume?” on page 121](#)
- ♦ [“How to resolve file move conflicts?” on page 122](#)

## How to view files and folders?

To view files and folders in a volume, perform the following steps:

1. Click **Files and Folders**  and use any one option to select the servers.
  - ♦ Click the **Search** icon, specify the required server name and select it from the drop-down list to view the available volumes.
  - ♦ Click **BROWSE**, select the required servers from the tree, and then click **APPLY**.
2. Click the Volume **Name** to view the files and folders in it.

## How to create a new folder?

Ensure to complete the following prerequisites before creating a folder in UMC.

- ♦ Users must have sufficient trustee rights to create folder at selected path.
- ♦ Target path or folder must be in the same tree as the logged-in user.

To create new folder in a volume, perform the following steps:

1. Click **Files and Folders**  and use any one option to select the servers.
  - ◆ Click the **Search** icon, specify the required server name and select it from the drop-down list to view the available volumes.
  - ◆ Click **BROWSE**, select the required servers from the tree, and then click **APPLY**.
2. Click the volume **Name** > **Add New Folder**  to create a new folder at folder level.

---

**NOTE:** You can perform the same action in a folder to create new sub folder.

---

3. Specify the new folder name and click **Confirm**.
4. (Optional) Select the newly created folder, click **More Options**  > **Properties** to view the folder details and trustees.
5. (Optional) Set the directory quota, owner, attributes, and trustees for the selected folder.

## How to modify properties of a volume, file, or folder?

To modify the properties of a volume, file, or folder, perform the following steps:

1. Click **Files and Folders**  and use any one option to select the servers.
  - ◆ Click the **Search** icon, specify the required server name and select it from the drop-down list to view the available volumes.
  - ◆ Click **BROWSE**, select the required servers from the tree, and then click **APPLY**.
2. Select the volume or click the volume **Name** to select the required file or folder, click **More Options**  > **Properties**.

The properties page consists of **Details** and **Trustees** tabs.

### Details tab

On the **Details** tab, you can modify Quota, Created By, and Attributes.

- ◆ **Quota:** Modify the existing quota in the **New Quota** field, select KB, MB, GB, or TB from the Units drop-down list, and then click **Confirm**.

### Modify Directory Quota

Directory VOL1

Used quota 0 Byte

Current quota 8.00 EB

New quota\* 85899 GB

Adjustment 0 Byte

Confirm Cancel

- ◆ **Modify Owner:** In this page, you can browse the server and select the required user or group to change the Owner.

### Modify Owner

Folder test file

Current User/Group .cn=admin.o=microfocus.t=blr8

Identity Source eDirectory Object Type All objects

9 Object(s) found in microfocus.

blr8- microfocus	OESCommonProxy_blr8- admin	admingroup
	blr8	novxregd
	novxtier	novksrvd
	www	wwwrun

Confirm Cancel

- ◆ **Attributes:** Turn on or turn off the toggle switch and click **Save** to modify the required attributes.

### Attributes

Archive	<input type="checkbox"/>
Hidden	<input type="checkbox"/>
Immediate Compression	X
Immediate Purge	<input type="checkbox"/>
Inhibit Delete	<input type="checkbox"/>
Inhibit Rename	<input type="checkbox"/>
Read Only	<input type="checkbox"/>

Save Cancel

## Trustees tab

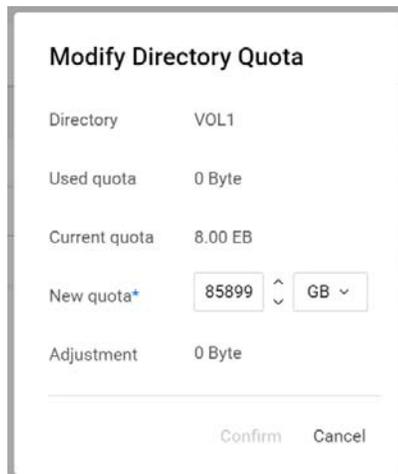
On the **Trustees** tab, you can view, add, and remove trustees with **Manage** option.

- ◆ **Refresh**: This option refresh the list of trustees for the selected volume, file, or folder.
- ◆ **Manage**: This action guides to right management page to manage trustee rights for the selected volume, file, or folder.

## How to modify directory quota of a volume or folder?

The directory quota for a volume or folder is not enabled by default. To modify the directory quota, perform the following steps:

1. Click **Files and Folders**  and use any one option to select the servers.
  - ◆ Click the **Search** icon, specify the required server name and select it from the drop-down list to view the available volumes.
  - ◆ Click **BROWSE**, select the required servers from the tree, and then click **APPLY**.
2. Select the volume or click the volume **Name** to select the required folder, click More Options  > **Properties**.
3. On the **Details** tab > **Quota**, click **Modify Quota** .
4. In the **Modify Directory Quota** box, update the new quota details, and then click **Confirm**.



Directory	VOL1
Used quota	0 Byte
Current quota	8.00 EB
New quota*	85899 <input type="text"/> GB <input type="text"/>
Adjustment	0 Byte

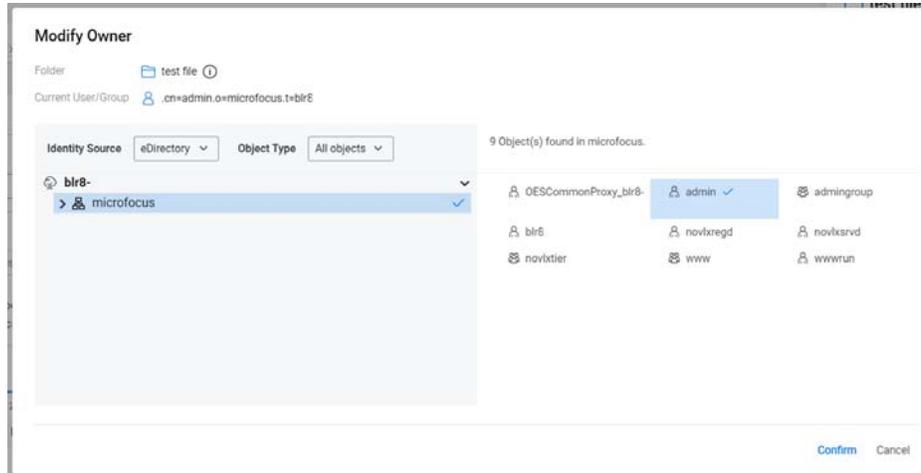
Confirm Cancel

## How to modify owner of a volume, file, or folder?

To modify owner of a volume, file, or folder, perform the following steps:

1. Click **Files and Folders**  and use any one option to select the servers.
  - ◆ Click the **Search** icon, specify the required server name and select it from the drop-down list to view the available volumes.
  - ◆ Click **BROWSE**, select the required servers from the tree, and then click **APPLY**.

2. Select the volume or click the volume **Name** to select the required file or folder, click **More Options**  > **Properties**.
3. On the **Details** tab > **Created By**, click **Modify user/group** .
4. On the **Modify Owner** page, browse the server to select the required user or group, and then click **Confirm**.



## How to modify attributes of a volume, file, or folder?

To modify attributes of a volume, file, or folder, perform the following steps:

1. Click **Files and Folders**  and use any one option to select the servers.
  - ◆ Click the **Search** icon, specify the required server name and select it from the drop-down list to view the available volumes.
  - ◆ Click **BROWSE**, select the required servers from the tree, and then click **APPLY**.
2. Select the volume or click the volume **Name** to select the required file or folder, click **More Options**  > **Properties**.
3. On the **Details** tab > **Attributes**, turn on or turn off the toggle switch, and then click **Save**.

### Attributes

- |                       |                                     |
|-----------------------|-------------------------------------|
| Archive               | <input type="checkbox"/>            |
| Hidden                | <input type="checkbox"/>            |
| Immediate Compression | <input checked="" type="checkbox"/> |
| Immediate Purge       | <input type="checkbox"/>            |
| Inhibit Delete        | <input type="checkbox"/>            |
| Inhibit Rename        | <input type="checkbox"/>            |
| Read Only             | <input type="checkbox"/>            |

Save

Cancel

## How to view deleted files and folders?

To view deleted files and folders, perform the following steps:

1. Click **Files and Folders**  and use any one option to select the servers.
  - ◆ Click the **Search** icon, specify the required server name and select it from the drop-down list to view the available volumes.
  - ◆ Click **BROWSE**, select the required servers from the tree, and then click **APPLY**.
2. Select the volume or click volume **Name** to select the folder in it, click **More Options** , and then select **Deleted Files & Folders**.

The list displays the deleted files and folders of the selected folder with the related information. You can salvage or purge these deleted files and folders if necessary.

- ◆ **Salvage:** You can restore the deleted files and folders using the **Salvage** option in **Deleted Files & Folders** location.
- ◆ **Purge:** You can permanently delete the deleted files and folders using the **Purge** option in **Deleted Files & Folders** location. Purged files and folders cannot be restored.

---

**NOTE:** On a junction, do not use the option **Deleted Files & Folders**, as it will not display the desired results.

---

## How to delete files and folders?

The deleted files and folders can be restored or permanently deleted from the **Deleted Files & Folders** location if required.

To delete files and folders in a volume, perform the following steps:

1. Click **Files and Folders**  and use any one option to select the servers.
  - ◆ Click the **Search** icon, specify the required server name and select it from the drop-down list to view the available volumes.
  - ◆ Click **BROWSE**, select the required servers from the tree, and then click **APPLY**.
2. Click volume name to select the required files and folders, click **More Options** , and then select **Delete**.
3. Click **Confirm** to delete the selected files and folders.

## How to salvage the deleted files and folders?

To salvage or restore the deleted files and folders, perform the following steps:

1. Click **Files and Folders**  and use any one option to select the servers.
  - ◆ Click the **Search** icon, specify the required server name and select it from the drop-down list to view the available volumes.
  - ◆ Click **BROWSE**, select the required servers from the tree, and then click **APPLY**.

2. Select the volume or click volume **Name** to select the folder in it, click **More Options** , and then select **Deleted Files & Folders**.
3. In the **Deleted Files & Folders** list, select the files and folders to restore, and then click **Salvage**.

The salvaged files and folders are restored to their respective locations.

## How to purge the files and folders?

To purge or permanently delete the files and folders, perform the following steps:

1. Click **Files and Folders**  and use any one option to select the servers.
  - ♦ Click the **Search** icon, specify the required server name and select it from the drop-down list to view the available volumes.
  - ♦ Click **BROWSE**, select the required servers from the tree, and then click **APPLY**.
2. Select the volume or click volume **Name** to select the folder in it, click **More Options** , and then select **Deleted Files & Folders**.
3. In the **Deleted Files & Folders**, select the required files and folders to permanently delete, and then click **Purge**.

The purge permanently deletes the selected files and folders from the volume and cannot be restored.

## How to rename a file or folder?

To rename a file or folder, perform the following steps:

1. Click **Files and Folders**  and use any one option to select the servers.
  - ♦ Click the **Search** icon, specify the required server name and select it from the drop-down list to view the available volumes.
  - ♦ Click **BROWSE**, select the required servers from the tree, and then click **APPLY**.
2. Click volume **Name** to select the file or folder in it, click **More Options** , and then select **Rename**.
3. Specify the New Name and click **Confirm**.

The list displays the selected file or folder with the new name.

## How to move files and folders in a volume?

To move files and folders in a volume, perform the following steps:

1. Click **Files and Folders**  and use any one option to select the servers.
  - ♦ Click the **Search** icon, specify the required server name and select it from the drop-down list to view the available volumes.
  - ♦ Click **BROWSE**, select the required servers from the tree, and then click **APPLY**.

2. Click volume **Name** to select the files and folders in it, click **More Options** , and then select **Move**.
3. In **Move Files** wizard, the File Information page lists the selected files and folders for move, click **Next**.
4. In the Target Location page, select the folder to move the selected files and folders, and click **Next**.

(Optional) You can click **Add New Folder** , specify the new folder name, and click  option to create new destination folder.

5. In the Summary page, verify the **Source** and **Destination** location, and then click **Finish**.

---

**NOTE:** **Resolve Conflicts** window is displayed, if the same file name exists in the destination location. To resolve this issue, see [“How to resolve file move conflicts?” on page 122](#).

---

The selected files and folders are moved to the new target location.

## How to resolve file move conflicts?

To resolve file move conflicts, perform the following steps:

1. Click **Files and Folders**  and use any one option to select the servers.
  - ♦ Click the **Search** icon, specify the required server name and select it from the drop-down list to view the available volumes.
  - ♦ Click **BROWSE**, select the required servers from the tree, and then click **APPLY**.
2. Click volume **Name** to select the files and folders in it, click **More Options** , and then select **Move**.
3. In **Move Files** wizard, the File Information page lists the selected files and folders for move, click **Next**.
4. In the Target Location page, select the folder to move the selected files and folders, and click **Next**.

(Optional) You can click **Add New Folder** , specify the new folder name, and click  option to create new destination folder.

5. In the Summary page, verify the **Source** and **Destination** location, and then click **Finish**.

Note: If the same file or folder names exists in the target location, **Resolve Conflicts** window is displayed.

6. In the **Resolve Conflicts** window, **Keep both** is selected as default, specify the **Prefix** or **Suffix** to rename all the conflicting files and folders.

You can also use **Overwrite** to replace the conflicting files and folders or **Skip** to ignore them.

7. Click **Continue** to finish the process.

# 15 Managing Rights

In Rights Management, you can add users or groups as trustees, modify the rights of existing trustees, replicate user or group rights, enable all rights for user or group, and remove all rights for user or group. For eDirectory users and groups, to view and modify the file system rights, you should either be an eDirectory administrator or a user who has administrative privileges.

- ♦ [“How to add trustee\(s\) for a volume, file, or folder?”](#) on page 123
- ♦ [“How to modify trustee rights for users and groups?”](#) on page 124
- ♦ [“How to view trustee rights of a volume, file, or folder?”](#) on page 124
- ♦ [“How to enable all rights for users and groups?”](#) on page 124
- ♦ [“How to disable all rights for users and groups?”](#) on page 125
- ♦ [“What are the various trustee rights?”](#) on page 125
- ♦ [“What are effective rights?”](#) on page 126
- ♦ [“How to view effective rights of users and groups?”](#) on page 126
- ♦ [“What are inherited rights?”](#) on page 127
- ♦ [“How to view inherited rights of a user or group?”](#) on page 127
- ♦ [“How to use inherited rights filter?”](#) on page 127
- ♦ [“How to copy or replicate rights of a user or group to other users and groups in the context tree?”](#) on page 128
- ♦ [“How to remove trustees for a selected path?”](#) on page 128

## How to add trustee(s) for a volume, file, or folder?

To add trustee(s) for a volume, file, or folder, perform the following steps:

1. Click **Files and Folders**  and use any one option to select the servers.
  - ♦ Click the **Search** icon, specify the required server name and select it from the drop-down list to view the available volumes.
  - ♦ Click **BROWSE**, select the required servers from the tree, and then click **APPLY**.
2. Click volume **Name** to select file or folder in it, click **More Options** , and then select **Manage Rights**.
3. On the **Manage Rights** page, click **Add Trustee**.
4. In the tree, select the servers to list the context users.
5. Select the users and groups, and then click **Confirm**.

The trustee rights for the newly added users and groups can be modified if necessary.

# How to modify trustee rights for users and groups?

To modify trustee rights for users and groups, perform the following steps:

1. Click **Files and Folders**  and use any one option to select the servers.
  - ◆ Click the **Search** icon, specify the required server name and select it from the drop-down list to view the available volumes.
  - ◆ Click **BROWSE**, select the required servers from the tree, and then click **APPLY**.
2. Select the volume or click volume **Name** to select the file or folder in it, click **More Options** , and then select **Manage Rights**.
3. On the **Manage Rights** page, select checkbox to modify the rights for the required users and groups, and then click **Apply Changes**.

S	R	W	C	E	M	F	A
<input type="checkbox"/>	<input checked="" type="checkbox"/>						

# How to view trustee rights of a volume, file, or folder?

To view and manage trustee rights of a volume, file, or folder, perform the following steps:

1. Click **Files and Folders**  and use any one option to select the servers.
  - ◆ Click the **Search** icon, specify the required server name and select it from the drop-down list to view the available volumes.
  - ◆ Click **BROWSE**, select the required servers from the tree, and then click **APPLY**.
2. Click volume **Name** to select the file or folder in it, click **More Options** , and then select **Manage Rights**.

The list displays the trustees and their rights on the selected volume, file, or folder. You can view, modify, add, remove, and replicate trustee rights here.

# How to enable all rights for users and groups?

To enable all rights for users and groups, perform the following steps:

1. Click **Files and Folders**  and use any one option to select the servers.
  - ◆ Click the **Search** icon, specify the required server name and select it from the drop-down list to view the available volumes.
  - ◆ Click **BROWSE**, select the required servers from the tree, and then click **APPLY**.
2. Select the volume or click volume **Name** to select the file or folder in it, click **More Options** , and then select **Manage Rights**.
3. On the **Manage Rights** page, select the users and groups.

---

**NOTE:** Use **Add Trustee** option to add users and groups if necessary.

---

4. Click **More Options** , select **Enable all rights**, and then click **Apply Changes**.  
All rights are enabled for the selected users and groups.

## How to disable all rights for users and groups?

To disable all rights for users and groups, perform the following steps:

1. Click **Files and Folders**  and use any one option to select the servers.
  - ◆ Click the **Search** icon, specify the required server name and select it from the drop-down list to view the available volumes.
  - ◆ Click **BROWSE**, select the required servers from the tree, and then click **APPLY**.
2. Select the volume or click volume **Name** to select the file or folder in it, click **More Options** , and then select **Manage Rights**.
3. On the **Manage Rights** page, select the users and groups.

---

**NOTE:** Use **Add Trustee** option to add users and groups if necessary.

---

4. Click **More Options** , select **Remove all rights**, and then click **Apply Changes**.  
All rights are disabled for the selected users and groups.

## What are the various trustee rights?

The table shows the list of available file-system trustee rights.

File-System Trustee Rights	Description
Supervisor (S)	Grants the trustee all rights to the directory or file and any subordinate items.  The Supervisor right cannot be blocked with an IRF (Inherited Rights Filter) and cannot be revoked. Users who have this right can also grant other users any rights to the directory or file and can change its Inherited Rights Filter.  Default = Off
Read (R)	Grants the trustee the ability to open and read files, and open, read, and execute applications.  Default = On
Write (W)	Grants the trustee the ability to open and modify (write to) an existing file.  Default = Off

File-System Trustee Rights	Description
Create (C)	Grants the trustee the ability to create directories and files and salvage deleted files.  Default = Off
Erase (E)	Grants the trustee the ability to delete directories and files.  Default = Off
Modify (M)	Grants the trustee the ability to rename directories and files, and change file attributes. It does not allow the user to modify the contents of the file.  Default = Off
File Scan (F)	Grants the trustee the ability to view directory and file names in the file system structure, including the directory structure from that file to the root directory.  Default = Off
Access Control (A)	Grants the trustee the ability to add and remove trustees for directories and files, to modify the rights assigned for trustees, and set the inherited rights filters.  This right does not allow the trustee to add or remove the Supervisor right for any user. Also, it does not allow to remove the trustee with Supervisor right.  Default = Off

## What are effective rights?

The effective rights granted to a trustee are a combination of explicit rights set at volume root or rights set on the file or folder and the inherited rights. Inherited rights are overridden by rights that are assigned explicitly for the trustee on a given path. If there are no trustees listed for effective rights, then the effective rights are the same as the inherited rights.

## How to view effective rights of users and groups?

To view the effective rights of users and groups, perform the following steps:

1. Click **Files and Folders**  and use any one option to select the servers.
  - ◆ Click the **Search** icon, specify the required server name and select it from the drop-down list to view the available volumes.
  - ◆ Click **BROWSE**, select the required servers from the tree, and then click **APPLY**.

2. Select the volume or click volume **Name** to select the file or folder in it, click **More Options** , and then select **Manage Rights**.
3. Select the **Effective & Inherited Rights** tab to list the users and groups, and their effective rights.

## What are inherited rights?

The inherited rights are trustee rights of subdirectories and files inherited from their parent directory. Typically, you set rights that you want to flow down to all users by assigning a Group object as the trustee of a directory located at the root of the volume. The trustee rights flow down through the file tree structure to its child subdirectories and files.

## How to view inherited rights of a user or group?

To view the inherited rights of a user or group, perform the following steps:

1. Click **Files and Folders**  and use any one option to select the servers.
  - ◆ Click the **Search** icon, specify the required server name and select it from the drop-down list to view the available volumes.
  - ◆ Click **BROWSE**, select the required servers from the tree, and then click **APPLY**.
2. Select the volume or click volume **Name** to select the file or folder in it, click **More Options** , and then select **Manage Rights**.
3. Select the **Effective & Inherited Rights** tab to list the users and groups, select the user or group, click **View Inherited Rights** .

The page displays the details of the inherited rights filters and effective rights of the user or group for the source file.

## How to use inherited rights filter?

Enabling inherited rights applies all the rights of the parent directory to the child directory. Disabling it restricts the rights to flow from the parent directory to the child directory.

To use the inherited rights filter, perform the following steps:

1. Click **Files and Folders**  and use any one option to select the server.
  - ◆ Click the **Search** icon, specify the required server name and select it from the drop-down list to view the available volumes.
  - ◆ Click **BROWSE**, select the required servers from the tree, and then click **APPLY**.
2. Select the volume or click volume **Name** to select the file or folder in it, click **More Options** , and then select **Manage Rights**.
3. Click **Inherited Rights Filter**, turn on or turn off the toggle switch to enable or disable inherited rights for all the selected users and groups for the file or folder.

# How to copy or replicate rights of a user or group to other users and groups in the context tree?

To copy or replicate rights of a user or group to other users and groups in the context tree, perform the following steps:

1. Click **Files and Folders**  and use any one option to select the servers.
  - ◆ Click the **Search** icon, specify the required server name and select it from the drop-down list to view the available volumes.
  - ◆ Click **BROWSE**, select the required servers from the tree, and then click **APPLY**.
2. Click volume **Name** to select the file or folder in it, click **More Options** , and then select **Manage Rights**.
3. In **Manage Rights** page, select the user or group, click **More Options** , and then select **Replicate**.
4. In **Replicate Rights** wizard, search or browse to list users and groups from the context tree.
5. Select the users and groups, and then click **APPLY**.

The selected users and groups displayed in the list and can be removed using **Remove**  option if necessary.

6. In **Replicate Rights** wizard, click **Confirm**.

The selected users and groups from the context tree must have the same rights of the user or group selected in **Manage Rights** page.

# How to remove trustees for a selected path?

To remove trustees for a selected path, perform the following steps:

1. Click **Files and Folders**  and use any one option to select the servers.
  - ◆ Click the **Search** icon, specify the required server name and select it from the drop-down list to view the available volumes.
  - ◆ Click **BROWSE**, select the required servers from the tree, and then click **APPLY**.
2. Select the volume or click volume **Name** to select the file or folder in it, click **More Options** , and then select **Manage Rights**.
3. In **Manage Rights** page, select the users and groups, click **More Options** , and then select **Remove**.
4. In **Remove Trustee** box, click **Remove**.

The access for the trustees are removed for the selected path.

# VI Storage Technology

This chapter describes the procedures for managing Distributed File Services (DFS) replica sites and junctions. For more information on DFS, see [Distributed File Services Administration Guide for Linux](#).

DFS is delivered as part of the Storage Services (novell-nss) userspace package. NSS must be installed and enabled on DFS replica servers to facilitate the DFS management context, as well as on any server where junctions are intended to be created.

---

**NOTE:** To access DFS, log in to UMC with your admin credentials, then click **Storage Technology > DFS**.

---

- ♦ [Chapter 16, “Managing Replica Sites,” on page 131](#)
- ♦ [Chapter 17, “Managing Junctions,” on page 137](#)



# 16 Managing Replica Sites

A replica site is the server that hosts an instance of the DFS Replica (VLDB) service and its associated replica file within a DFS management context. Each management context consists of one or two replicas, which can operate on any combination of supported DFS platforms. These servers can exist at the same level or below the management context in the eDirectory tree; however, they must not be part of a lower-level DFS management context.

---

**NOTE:** To access DFS, log in to UMC with your admin credentials, then click **Storage Technology > DFS > Replica Sites**.

---

- ◆ [“Naming Convention Changes” on page 131](#)
- ◆ [“How to list replica sites?” on page 131](#)
- ◆ [“Where to view details of a replica site?” on page 133](#)
- ◆ [“How to create a Management Context?” on page 133](#)
- ◆ [“How to add a Replica site?” on page 134](#)
- ◆ [“How to repair the DFS Replica service?” on page 134](#)
- ◆ [“How to configure the DFS Replica service?” on page 135](#)
- ◆ [“How to delete a Replica Site?” on page 135](#)
- ◆ [“What happens when a Replica Site is paused or stopped?” on page 135](#)

## Naming Convention Changes

**Management Context:** On creating a replica site, the selected Organization (O) or Organizational Unit (OU) becomes the management context. There is no separate workflow to create a management context.

**DFS Replica (VLDB):** The DFS Replica (VLDB) service provides the framework for locating volumes within the management context. Managing this service involves creating, day-to-day management, maintenance, and repair of the replica. In UMC, the term VLDB is replaced with DFS Replica service for better understanding.

## How to list replica sites?

- 1 Browse and select the DFS management context you want to manage. This displays the replica sites.

---

**NOTE:** When you browse the O or OU, it lists the existing management contexts.

---

- 2 The following information is displayed:

Column Name	Description
DFS status (Color Coding)	Status
Green	<b>Running:</b> The DFS Replica service is running.
Gray	<b>Stopped:</b> The DFS Replica service is stopped.
White	<b>Unknown:</b> UMC cannot determine the status of the replica site.
Replica status (Color Coding)	Status
Green	<b>Running:</b> The DFS Replica service is loaded and running.
Blue	<b>Repairing:</b> The DFS Replica service is being repaired.  Repair progress is not stored, so it is advisable not to interrupt the repair; otherwise, it would need to be restarted. The repair status is available in the <b>Details</b> section of each replica site.
Gray	<b>Stopped:</b> The DFS Replica service is stopped.  The service is manually stopped or after repair activating the service has failed and the status has changed to stopped.
White	<b>Unknown:</b> UMC cannot determine the status of the replica site.
Red	<b>Failed:</b> The DFS service has stopped, causing the DFS Replica service to is unload.  No volume operations can be performed on this replica site.
Server	The name of the replica site.
Management context	The name of a preexisting O or OU container that you chose from the eDirectory tree.

On selecting a replica site, the following actions can be performed:

- ◆ [Details](#)
- ◆ [Add](#)
- ◆ [Configure](#)
- ◆ [Pause](#)
- ◆ [Start and Resume](#)
- ◆ [Stop](#)
- ◆ [Repair DFS replica](#)
- ◆ [Delete](#)

## Where to view details of a replica site?

- 1 Browse and select the DFS management context you want to manage. This displays the replica sites.
- 2 Select a replica site, then select **Details**. The following information is displayed:

Parameter	Description
Status	State of the DFS Replica service.
Threads running	Displays the number of actual threads running for the service. Display the number of processing threads for the service.  The number of running threads can vary because of lack of memory on the server, or because the number of running threads is in the process of changing to meet the requested number.
Threads requested	Display the number of processing threads configured for the service. Range: 1 (default) to 16.
Running since	The date and time when the DFS Replica service was activated.
Management Context	The context of the selected replica site.
Path	The location of the replica database file. The default location is <code>/var/opt/novell/dfs</code> .
Last Repair	The date and time of the repair, level, and the status of the repair.

## How to create a Management Context?

A management context can support a maximum of two replica sites. On creating a replica site, the selected O or OU container becomes the management context.

- 1 Click **Create Replica Site**.
- 2 A wizard is displayed:
  - 2a **Management context:** Browse and select a container, then click **Next**.

---

**NOTE:** The selected container is designated as the management context for this replica.

---

- 2b **Servers:** Browse and select the server on which the DFS Replica service should be hosted. You can select a maximum of two servers.

- 2c **DFS replica location:** Select the default path (`/var/opt/novell/dfs`), or alternatively, select an NSS volume or a folder within the volume to store the DFS Replica (VLDB) database on the replica site. Click **Next**.

The name of the DFS Replica file itself cannot be specified or modified; it is always `vldb.dat`.

- 2d **Summary:** Review the summary of the created replica site and click **Finish**.

A new replica site is created in the selected management context.

## How to add a Replica site?

A maximum of two replica sites can be created for a DFS management context. These two replicas exchange databases (the entire database, not just the changes) whenever a change is made to the volumes. Upon receiving the database from the other replica, each replica merges it with its own, determining which entries have been added, deleted, or modified.

- 1 Browse and select the DFS management context you want to manage. This displays the existing replica sites.
- 2 Select a replica site, then select **Add**.
- 3 A wizard is displayed:
  - 3a **Servers:** Browse and select a server, then click **Next**.
  - 3b **DFS replica location:** Select the default path (`/var/opt/novell/dfs`) or a new folder to store the DFS Replica (VLDB) database on the replica site. Click **Next**.
  - 3c **Summary:** Review the summary of the replica site and click **Finish**.

A new replica site is added to the existing management context.

## How to repair the DFS Replica service?

The repair process rebuilds the DFS Replica database. Upon completion, the current active database is replaced with the repaired one. If there are two replica sites, the replica automatically synchronizes with the active repaired database. Until the repair database is activated, all DFS Replica requests (except those specifically referencing the repaired database) operate against the existing database. Thus, clients can access DFS junctions even during repair for volumes that still have correct entries in the DFS Replica database.

- 1 Browse and select the DFS management context you want to manage. This displays the replica sites.
- 2 Select a replica site, then select **Repair DFS Replica**.
- 3 Select one of the following repair levels, then click **OK**:
  - ♦ **Replace with the last saved copy:** Restores the last saved copy of the database using the automatically created backup file.
  - ♦ **Copy from another replica site:** Retrieves a copy of the database from another server that is currently running the DFS Replica service.  
This option is available only if there is more than one replica site.
  - ♦ **Rebuild from the eDirectory tree:** Rebuilds the database from scratch by recursively scanning the eDirectory tree down from the management context container and recording volume object information in the repaired database. This is a time-consuming activity and should be considered carefully.
- 4 Click **Confirm**. Monitor the status of the rebuild periodically until it completes. This duration can vary from a few minutes to several days depending on the selected repair level. To view the progress, select the replica site, and then **Details**.

During the repair process, the status displays as **Repairing**. If the option **Rebuild from the eDirectory tree** is selected, upon completion of the repair, DFS automatically reloads the DFS Replica service on the replica server and activates the database, changing the state to **Running**. If there is a second replica site, its database copy is automatically synchronized with the repaired database.

- 5 If any errors occur during the repair, refer to the following log file:

```
/var/opt/novell/log/dfs/vlrpr.log
```

## How to configure the DFS Replica service?

A few parameters of the DFS Replica service are configurable.

- 1 Browse and select the DFS management context you want to manage. This displays the replica sites.
- 2 Select a replica site, then select **Configure**.
- 3 A wizard is displayed:
  - 3a **Threads:** Edit the number of processing threads configured for the service. Range: 1 (default) to 16.
  - 3b **DFS Replica:** Select a path to store the DFS Replica (VLDB) database.
  - 3c **Run DFS replica service on server restart:** Enable this option if you want the service to start automatically when you restart the server.
  - 3d Click **Confirm** to save the changes for the replica site.

## How to delete a Replica Site?

Deleting a replica site deactivates and unloads the DFS Replica service, deletes the database file, and then updates the attribute for the DFS management context in eDirectory.

---

**IMPORTANT:** If the selected site is the last remaining replica site, deleting it also deletes its DFS management context.

---

- 1 Browse and select the DFS management context you want to manage. This displays the replica sites.
- 2 Select a replica site, then select **Delete**.
- 3 Click **Delete** to remove the selected replica site.

DFS synchronizes the changes with eDirectory, which might take up to 5 minutes.

## What happens when a Replica Site is paused or stopped?

**For example,**

In the management context “Operations”, there are two replica sites: 10.65.8.11 and 10.66.8.12. Below are the effects of pausing and resuming operations on these sites.

## Pause

10.65.8.11 is in a Paused state, while 10.66.8.12 is in a Running state.

- ◆ In UMC, the DFS **Replica status** (10.65.8.11) shows as Stopped.
- ◆ The DFS Replica service (10.65.8.11) is stopped but remains loaded. Volume operations performed on this site updates the DFS Replica database (10.65.8.11) and also synchronizes with 10.65.8.12.
- ◆ Users cannot access junctions available on the paused replica site (10.65.8.11).
- ◆ The available operations are: details, configure, resume, start, stop, and delete.

To pause a replica site:

- 1 Browse and select the DFS management context you want to manage. This displays the replica sites.
- 2 Select a replica site, then select **Pause**.
- 3 Click **Confirm** to pause the replica site. The **Replica status** changes to Stopped.

## Stop

10.65.8.11 is in a Running state, while 10.66.8.12 is in a Stopped state.

- ◆ In UMC, the **DFS status** is Stopped and DFS **Replica status** is Failed.
- ◆ Since the DFS service is stopped, it unloads the DFS Replica service. Volume operations performed on this site are not updated in the DFS Replica database (10.66.8.12) but are updated on its replica site (10.65.8.11) as replication continues.
- ◆ Users cannot access junctions available on the stopped replica site (10.66.8.12).
- ◆ The available operation is delete.

To stop a replica site:

- 1 Browse and select the DFS management context you want to manage. This displays the replica sites.
- 2 Select a replica site, then select **Stop**.
- 3 Click **Confirm** to stop the replica site. The **DFS status** has changed to Stopped and the DFS **Replica status** to Failed.

# 17 Managing Junctions

A DFS junction serves as a logical placeholder for data stored on a different NSS volume. Each junction points to a single target location.

To administrators, the junction appears in the file structure as a folder. However, users typically see the junction as a sub-folder and is unaware of its existence. If the target path is unavailable or if the DFS Replica service for the target's management context is not running, users cannot access the target data. Clients that are not DFS-aware see a junction as a file that they have no rights to access.

---

**NOTE:** To access DFS, log in to UMC with your admin credentials, then click **Storage Technology > DFS > Junctions**.

---

- ♦ [“What are the guidelines for creating or managing junctions?” on page 137](#)
- ♦ [“How to create a junction?” on page 137](#)
- ♦ [“Where to view junctions?” on page 138](#)
- ♦ [“How to configure junctions?” on page 140](#)
- ♦ [“How to delete junctions?” on page 140](#)
- ♦ [“How to synchronize rights between the Source and Target locations?” on page 140](#)

## What are the guidelines for creating or managing junctions?

- ♦ Junctions can exist between the source and target volumes within the same or different DFS management contexts.
- ♦ When creating a junction, a new folder can be created. This functionality is exclusive to UMC.
- ♦ Only eDirectory users can be added as trustees to a junction.
- ♦ Both the junction and target locations inherit trustees and their rights relative to their actual locations, following the OES Trustee Model. Using the **Sync** functionality, you can synchronize the explicit rights of a junction between the source and target locations. For more information on trustee rights, see [“What are the various trustee rights?” on page 125](#).

## How to create a junction?

To create junction, follow these steps:

- 1 Click **Create junction**.
- 2 A wizard is displayed:
  - 2a Browse and select the DFS management context in which to create the junction.

---

**NOTE:** Junction are created only within the selected management context.

---

## 2b Source path:

**2b1 Name:** Enter the name of the junction.

**2b2** Browse and select the NSS volume or folder where you want to create the junction, then click **Continue**.

To navigate the volume, click the object.

**2c Target path:** Browse and select the NSS volume or folder where you want the junction to point, then click **Continue**.

The target NSS volume or folder is where the data resides.

**2d Source trustee rights:** Set eDirectory trustees and their rights for the source. Browse and select one or more users to set as trustees, then click **Apply**.

**2d1 Assigned rights:** Select trustee and assign required rights. By default, the trustee is listed with a minimum of `Read` and `File Scan` rights. Modify the trustee rights if necessary.

---

**NOTE:** All trustee operations supported in Files & Folders can be performed in this page (Source and Target).

---

**2d2 Effective rights:** Rights are not available because the junction has not been created.

**2e Target trustee rights:** Set eDirectory trustees and their rights for the target. Browse and select the users set on source along with any additional users. Then, set the trustees rights and click **Apply**.

**2e1 Assigned rights:** Select trustee and assign required rights. By default, the trustee is listed with a minimum of `Read` and `File Scan` rights.

---

**IMPORTANT:** For file visibility, users need at least `Read` and `File Scan` rights on the target location.

---

**2e2 Effective rights:** Effective rights on the junction target include explicitly defined rights on the junction itself and rights inherited from the junction's parent directory. These rights are not editable.

**2f Summary:** Review the summary of the newly created junction and click **Finish**.

On the junction listing page, select the server or volume to view the newly created junction.

## Where to view junctions?

A junction is a virtual folder that points to the root of a target NSS volume or to any of its directories. You can view the list of junctions in two locations:

- ♦ [“DFS > Junctions” on page 139](#)
- ♦ [“Files & Folders” on page 139](#)

## DFS > Junctions

- 1 Browse and select the servers or volumes to list the junctions.
- 2 (Conditional) When connecting to a server for the first time, you must scan all the volumes, to cache the junction information. Click **Scan now** or **Run scan** to list the junctions.

After creating new junctions, click **Refresh** to update the cache and display the newly added junctions in the junctions list.

- 3 The following information is displayed:

Column Name	Description
Status (Colour Coding)	The states of the junction are Available or Broken.
Green	<b>Available:</b> The data at the target location is accessible through the junction.
Red	<b>Broken:</b> The target location that the junction points to is unavailable.
Name	The name specified by the administrator.
Management context	The management context of the selected server or volume.
Source path	A folder path on the volume or root of the volume where the junction resides.
Target path	A folder path on the volume or root of the volume where the data resides.
OES target	The target server is an OES server.
Last modified	The timestamp indicating when the junction was last modified.

When selecting a junction, the following actions can be performed:

- ◆ Details - The same information is available on the junction listing page. Additional information is the creation date of the junction.
- ◆ Rename
- ◆ Configure
- ◆ Sync rights - source to target
- ◆ Sync rights - target to source
- ◆ Delete

## Files & Folders

- 1 Browse and select the servers to list the volumes.
- 2 Click the volume to view the junctions. Junctions are listed as a file in the volume or its folders.

## How to configure junctions?

When configuring a junction, the source path and junction name cannot be modified.

- 1 Browse and select the servers or volumes to list the junctions.
- 2 Select a junction, then select **Configure**.
  - 2a The target path, source trustee rights and target trustee rights can be modified.
  - 2b **Summary:** Review the changes made and click **Finish**.

On the junction listing page, select the server or volume to view the modified junction.

## How to delete junctions?

Deleting a junction removes the junction file and its associated trustees, trustee rights, and inherited rights set on the junction. The data and trustee rights at the target location are not affected.

- 1 Browse and select the servers or volumes to list the junctions.
- 2 Select a junction, then select **Delete**.

---

**NOTE:** To avoid security or visibility issues, ensure to verify trustee settings at the target location before or after deletion.

---

- 3 Click **Delete** to remove the selected junctions.

## How to synchronize rights between the Source and Target locations?

To synchronize all the assigned rights of a trustee, follow these steps:

- 1 Browse and select the servers or volumes to list the junctions.
- 2 Select a junction and choose either the **Sync rights - source to target** option or the **Sync rights - target to source** option. This action copies the trustee rights from the source to the target or vice versa.
- 3 To validate the rights, click **Files & Folders**.
- 4 Browse and select the servers to list the volumes.
- 5 Select the volume, then select **Manage Rights**. This displays the users with their modified rights.

# VII Service configuration

- ◆ Chapter 18, “Managing NCP,” on page 143
- ◆ Chapter 19, “Managing SMDR,” on page 149
- ◆ Chapter 20, “Managing TSAFS,” on page 151



# 18 Managing NCP

This chapter describes the configuration options used by NCP server. For more information on NCP server settings, see [NCP Server for Linux Administration Guide](#).

- ♦ [“How to manage the local code page?”](#) on page 143
- ♦ [“How to manage NCP server caching settings?”](#) on page 143
- ♦ [“How to manage encryption and MFA on an NCP server?”](#) on page 144
- ♦ [“How to manage NCP server setting locks?”](#) on page 144
- ♦ [“How to manage NCP server communications settings?”](#) on page 145
- ♦ [“How to manage NCP server volumes settings?”](#) on page 145
- ♦ [“How to manage NCP server logging settings?”](#) on page 146
- ♦ [“How to manage NCP server performance tuning settings?”](#) on page 147
- ♦ [“How to manage NCP server user ID updates?”](#) on page 147

## How to manage the local code page?

The NCP server supports most of the code pages used for files and subdirectory names. NCP server by default uses the code page corresponding to the code page used by the Linux server operating system that is specified at install time.

If you want to select a different local code page, follow the steps:

- 1 In **Configuration > NCP**, search or browse to select the server.
- 2 In **Server Environment**, select a new local code page from the **Local code page** drop-down and click **Save**.

## How to manage NCP server caching settings?

- 1 In **Configuration > NCP**, search or browse to select the server.
- 2 In **NCP > Server Environment** drop-down, select **NCP > Caching**.
  - ♦ **Maximum cached files per sub-directories** are the number of files that can be cached for a sub-directory.
  - ♦ **Maximum cached files per volume** are the number of files that can be cached for a volume.
  - ♦ **Maximum cached sub-directories per volume** are the number of sub-directories that can be cached for a volume.
  - ♦ **Maximum lazy close files** are the number of file handles that can be lazy closed.
- 3 Specify the required values and click **Save**.

# How to manage encryption and MFA on an NCP server?

The NCP server security allows you to manage encryption and MFA on an NCP server.

1 In **Configuration** > **NCP**, search or browse to select the server.

2 In **NCP** > **Server Environment** drop-down, select **NCP** > **Security**.

3 **Encryption:**

- ◆ **Enable, Disable, or Enforce** to manage encryption capabilities between the NCP server and NCP clients.
- ◆ Specify the Grace Period and Cipher strength.

---

**NOTE:** If Encryption is enforced, then cipher strength is set to low and grace period is disabled.

---

**MFA:**

- ◆ **Enforce** or **Disable** MFA to manage the connections to access NCP server.
- ◆ Specify the **Grace Period**.

---

**NOTE:** If MFA is enforced, then Grace Period is disabled.

---

**Auditing:**

When auditing is enabled, the security configuration setting changes performed on an NCP server are logged.

4 Select the required security options and then click **Save**.

# How to manage NCP server setting locks?

1 In **Configuration** > **NCP**, search or browse to select the server.

2 In **NCP** > **Server Environment** drop-down, select **NCP** > **Locks**.

3 You can manage the following options:

- ◆ **Cross-protocol locks**

Cross-protocol locks prevent the same file from being concurrently accessed for modifications from both CIFS and NCP client.

- ◆ **Oplock support level**

NCP opportunistic locking allows the client to cache file data for better performance. You can select anyone of the options from the drop-down.

- ◆ Disable
- ◆ Exclusive locks
- ◆ Shared and exclusive locks

- ◆ **Lock range mask**

Allow applications to acquire a lock above the address (0x7fffffffffffffff) region.

- ◆ **Byte range lock spin time**

Specify the range between 0 and 5000 (milliseconds) to avoid lock collisions when the LockTimeOut is sent as 0 in a Byte Range Request from the client.

- ◆ **Log lock statistics**

When an NCP volume lock is held more than the configured time, the NCP server displays a message in the ncpserv.log file with the relevant details.

4 Select and specify the required NCP server locks options and click **Save**.

## How to manage NCP server communications settings?

1 In **Configuration > NCP**, search or browse to select the server.

2 In **NCP > Server Environment** drop-down, select **NCP > Communications**.

3 You can manage the following options:

- ◆ **First watchdog packet**

Enable and specify the time when the NCP server should send a ping packet if no client activity is detected.

- ◆ **Broadcast message**

Enable or disable broadcast messages from the NCP server.

- ◆ **TCP/NCP keep alive interval**

Specify the time when the NCP server must send a TCP packet if no client activity is detected.

- ◆ **NCP keep alive interval**

Enable and specify the time when the NCP server should send a TCP packet if no client activity is detected.

4 Select and specify the required NCP server communications options and click **Save**.

## How to manage NCP server volumes settings?

1 In **Configuration > NCP**, search or browse to select the server.

2 In **NCP > Server Environment** drop-down, select **NCP > Volumes**.

3 You can manage the following:

- ◆ **Commit file**

Ensures that all the data written to a file by NCP client is written to the disk.

- ◆ **Execute attribute support**

Enables to use NCP "execute only" attribute with the user mode execute bit on a file or sub-directory.

- ◆ **Keep NSS file deleter IDs**

Retains the deleter ID when a file is deleted in NSS volumes.

- ◆ **Sendfile support**

The NCP server sends the file read data to the clients directly to the Linux Kernel Ring 0 environment. This option is not supported for encrypted connections.

- ◆ **Sync trustees to NSS at volume mount**  
Resynchronizes trustees for an NSS volume when it is mounted for NCP.
  - ◆ **Warn users – volume is full**  
Warn users when no space is available on the volume.
  - ◆ **Warn users – volume path is unavailable**  
Warn users when the volume path is no longer present.
  - ◆ **Warn users – volume space is low**
  - ◆ **Low volume warning threshold**  
Specify the low watermark threshold for volume (in blocks) to warn users when space is low. An NSS block is 4 KB.
  - ◆ **Low volume warning reset threshold**  
Specify the high watermark threshold for volume (in blocks). An NSS block is 4 KB. Sets the high watermark threshold (in MB), which is the level where the low watermark threshold is reset, and users no longer receive the low-space message.
  - ◆ **Trustee build wait time**  
Specify the time that the NCP server waits to build the trustee cache during volume mount.
- 4 Select and specify the required NCP server volumes options and click **Save**.

## How to manage NCP server logging settings?

- 1 In **Configuration > NCP**, search or browse to select the server.
- 2 In **NCP > Server Environment** drop-down, select **NCP > Logging**.
- 3 You can manage the following:
  - ◆ **NCPServ log level**  
Select the log level. The logs are available in the `/var/opt/novell/log/ncpserv.log` file.
  - ◆ **NCP2NSS log level**  
Select the log level. The logs are available in the `/var/opt/novell/log/ncp2nss.log` file.
  - ◆ **NCPCON log level**  
Select the log level. The logs are available in the `/var/opt/novell/log/ncpcon.log` file.
  - ◆ **Log cache statistics**  
Enables logging the cache statics of the NCP server to the `/var/opt/novell/log/ncpserv.log` file.
  - ◆ **Log ID broker statistics**  
Enables logging the ID broker errors to the `/var/opt/novell/log/ncpserv.log` file.
  - ◆ **Log memory statistics**  
Enables logging the memory statistics to the `/var/opt/novell/log/ncpserv.log` file.

- ◆ **Log eDirectory object history**

Enables NCP to send a notification to NSS when an object is deleted or renamed in the eDirectory and logs the event in the `/opt/novell/ncpserv/sbin/objecthistory.txt` file.

4 Specify and select the required NCP server logging setting and click **Save**.

## How to manage NCP server performance tuning settings?

1 In **Configuration > NCP**, search or browse to select the server.

2 In **NCP > Server Environment** drop-down, select **NCP > Performance tuning**.

3 You can manage the following:

- ◆ **Connection memory buffer pool size**

Specify the buffer pool size to be used for certain NCP verb replies. Changing this option requires ndsd service restart. For more information, see [Augmented Size of NCP Verbs 87\\_20 and 89\\_20 Replies](#) section in [NCP Server for Linux Administration Guide](#).

- ◆ **Concurrent async requests**

Specify the maximum number of asynchronous threads that can be created to process eDirectory or NCP requests.

- ◆ **Additional SSG threads**

Specify the number of additional SSG threads that can be used to process the incoming NCP file service request. These threads are used when the fixed 25 NCP threads are busy.

- ◆ **CPU affinity**

CPU affinity is applied to SSG threads in the NCP server to improve the encryption performance. 50% of the active CPUs are used for CPU affinity with the same number of SSG Threads.

4 Specify and select the required NCP server performance tuning settings and click **Save**.

## How to manage NCP server user ID updates?

1 In **Configuration > NCP**, search or browse to select the server.

2 In **NCP > Server Environment** drop-down, select **NCP > User ID updates**.

UID update mode allows you to set the frequency of the maintenance thread to update the UIDs.

3 Select the required **UID update mode** and click **Save**.



# 19 Managing SMDR

This chapter describes the configuration options used by Storage Management Data Requester (SMDR).

SMS provides remote backup and restore services using SMDR. When configuring SMDR, the changes are saved in the `/etc/opt/novell/sms/smdrd.conf` file on OES servers. SMDR reads this configuration file to verify if any values have been modified.

For more information on SMS, see [Storage Management Services Administration Guide for Linux](#).

## How to configure SMDR?

- 1 Log in to UMC with admin credentials.
- 2 Click **Configuration** > **SMDR**.
- 3 Browse and select the server for which you want to modify the SMDR configuration settings.
- 4 **Encryption (TLS):** To enhance the security of remote backup connections created by SMDR, you can modify the TLS version used. By default, SMDR uses TLS version 1.3 for encryption. However, it can be configured to use TLS version 1.2 if needed. When TLS version 1.3 is enabled, connections using TLS 1.2 are still accepted.
- 5 **IP address:** In a multi-homed environment, you can configure the IP address on which SMDR listens. If a server is assigned with multiple IP addresses, specify the desired IP address for SMDR to use.

---

**NOTE:** This option is unavailable if multiple servers are selected in [Step 3](#). By default, SMDR uses the first bound IP address of the server.

---

- 6 **Discovery Mechanism:** SMDR uses SLP (Service Location Protocol) and HOSTS/DNS for discovery and name resolution. Based on your selection, the priority of the discovery mechanisms is updated in the `/etc/opt/novell/sms/smdrd.conf` file.
- 7 **Autoload TSANDS:** By default, this option is disabled. When enabled, this setting is automatically loaded and applied upon restarting the OES server or SMS service.
- 8 **Autoload TSAFS:** By default, this option is enabled. This setting is automatically loaded and applied upon restarting the OES server or SMS service restarts.
- 9 **Enable GroupWise backup:** By default, this option is disabled. When enabled, TSAFS supports the backup of GroupWise database files.
- 10 On modifying the above parameters, you must restart SMDR service.

```
systemctl restart novellsmdrd.service
```

This command restarts the `smdrd` daemon.



# 20 Managing TSAFS

This chapter describes the configuration options used by Target Service Agent for File System (TSAFS).

TSAFS provides configurable parameters to help optimize its performance. Changes made to TSAFS configuration are saved in the `/etc/opt/novell/sms/tsafs.conf` file on OES servers. When TSA is loaded, it reads this configuration file to verify if any values have been modified.

For more information on SMS, see [Storage Management Services Administration Guide for Linux](#).

## How to configure TSAFS?

- 1 Log in to UMC with admin credentials.
- 2 Click **Configuration** > **TSAFS**.
- 3 Browse and select the servers for which you want to modify the TSAFS configuration settings.
- 4 **Read buffer size:** This parameter controls the number and size of read requests to the file system.  
  
By default, the **Read Buffer Size** is set to 65536 bytes, with a configurable range from 16384 bytes to 262144 bytes. It is recommended to set this value as an integral multiple of the file system block size.
- 5 **Read threads per job:** This parameter controls the number of simultaneous read requests to the file system, determining the rate at which the read-ahead cache is built.  
  
By default, the **Read threads per job** is set to 4, with a configurable range from 1 to 32.
- 6 **Read thread allocation:** This parameter controls the maximum number of read threads that may be allocated to process a single data set.  
  
By default, the **Read Thread Allocation** is set to 100 (%), with a configurable range from 10 (%) to 100 (%). It is recommended to set this value to 100 (%) if the backup application requests data sets serially.
- 7 **Read ahead throttle:** This parameter limits the number of simultaneous data sets that are cached. In certain runtime scenarios, it helps override the **Read Thread Allocation** to ensure the completion of processing large data sets.  
  
By default, the **Read Ahead Throttle** is set to 2, with a configurable range from 1 to 32.
- 8 **Cache memory threshold:** This parameter controls the maximum amount of server memory that the TSA uses to store cached data sets.  
  
By default, the **Cache Memory Threshold** is set to 25 (%), with a configurable range from 1 (%) to 25 (%).
- 9 **Enable caching:** By default, this option is enabled. It specifies whether the TSA should perform predictive caching during backups. Caching can improve backup performance for certain workloads by prefetching files into memory.

- 10 Enable clustering:** By default, this option is enabled. If the backup server does not support clusters, the option is disabled. It specifies whether the TSA should be cluster-aware and recognize cluster pools as resources for backup or restore.
- 11** On modifying the above parameters, you must reload TSAFS service.

```
smsconfig -u tsafs  
smsconfig -l tsafs
```

This command loads TSAFS with updated configuration settings.

# VIII

## File Access Protocols

This chapter describes the procedures for managing NCP and CIFS shares, connections, and their global configurations on a server. For more information, see [OES 23.4: NCP Server for Linux Administration Guide](#) and [OES CIFS for Linux Administration Guide](#).

---

**NOTE:** The servers must be on OES 24.1 to list the NCP servers.

---

- ♦ [Chapter 21, “Managing NCP Shares,” on page 155](#)
- ♦ [Chapter 22, “Managing NCP Connections,” on page 165](#)
- ♦ [Chapter 23, “Managing CIFS Shares,” on page 169](#)
- ♦ [Chapter 24, “Managing CIFS Connections,” on page 179](#)
- ♦ [Chapter 25, “Managing Invalid Users,” on page 183](#)
- ♦ [Chapter 26, “Managing User Context,” on page 187](#)



# 21 Managing NCP Shares

- ♦ “What is an NCP share and how to manage it?” on page 155
- ♦ “How to list NCP shares?” on page 156
- ♦ “How to verify trustees for an NCP share? (OES 23.4)” on page 156
- ♦ “How to verify the rights of an NCP share?” on page 157
- ♦ “How to resync trustees an NCP share? (OES 23.4)” on page 157
- ♦ “How to resync the rights of an NCP share?” on page 157
- ♦ “How to enable or disable encryption on an NCP share?” on page 158
- ♦ “How to enable or disable MFA on an NCP share?” on page 158
- ♦ “What are accessed files and how to view them? (OES 23.4)” on page 158
- ♦ “What are open files and how to view them?” on page 159
- ♦ “What are the prerequisites for adding a secondary volume?” on page 159
- ♦ “How to add secondary volume?” on page 159
- ♦ “How to view secondary volume?” on page 160
- ♦ “How to remove secondary volume?” on page 160
- ♦ “How to manage security for sub-folders on an NCP share? (OES 23.4)” on page 161
- ♦ “How to manage sub-folder security on an NCP share?” on page 162
- ♦ “How to enable or disable write permission for an NCP share?” on page 162
- ♦ “How to activate or deactivate an NCP share?” on page 162

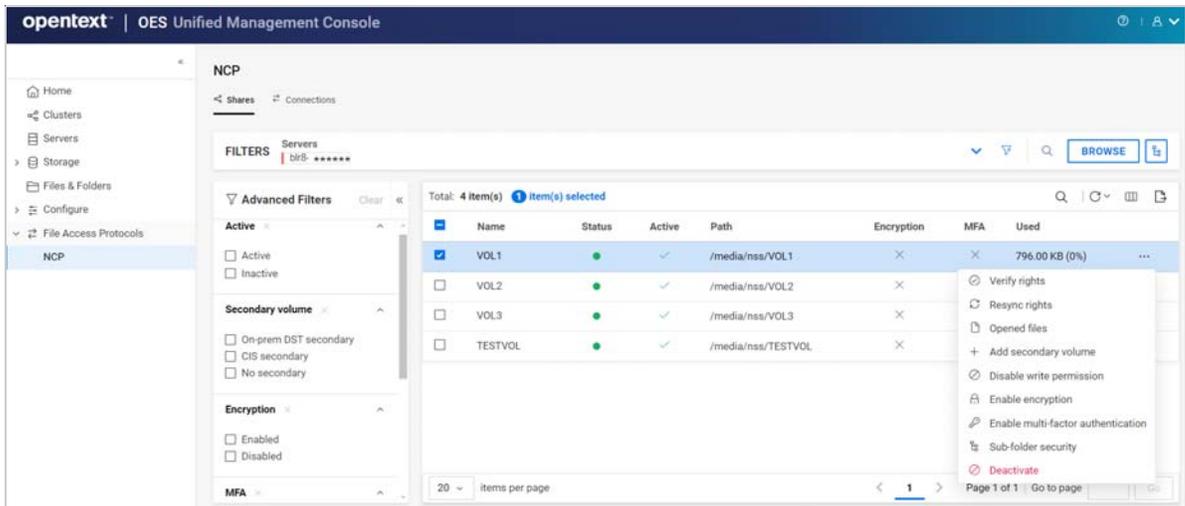
## What is an NCP share and how to manage it?

NCP volumes are NCP shares on Linux POSIX file systems such as Ext3, XFS, and Reiser. Storage Services (NSS) volumes are a special type of NCP volume.

The directory and file access are controlled with the OES trustee model for file system trustees and trustee rights. Users access NCP volume data using the Client for Open Enterprise Server software on their Windows or Linux workstations.

These are few actions that can be performed on a share:

- ♦ Verify and resync rights
- ♦ View and manage open files
- ♦ Manage encryption and MFA
- ♦ Activate or deactivate



## How to list NCP shares?

- 1 In UMC, click **File Access Protocols** > **NCP**.
- 2 Click the Search icon and specify the server name.

or

Click **Browse** and select server type to list their associated servers. Select the required servers from the list, and then click **APPLY**.

---

**NOTE:** When the **BROWSE** or tree view **📁** icon is clicked, you cannot perform other actions outside the browse area. Click the same button again to close the browse or tree view area.

---

This displays the list of NCP shares available on the server.

## How to verify trustees for an NCP share? (OES 23.4)

The verify trustees option shows the difference in trustees rights information between the NSS and NCP server for the specified NCP share. This action can be performed on multiple shares at a time.

- 1 In UMC, click **File Access Protocols** > **NCP**.
- 2 Click the Search icon and specify the server name.

or

Click **Browse** and select **Server Type** to list their associated servers. Select the required servers from the list, and then click **APPLY**.

- 3 Select the NCP share, click More Options **⋮** icon, and then select **Verify trustees**.

---

**NOTE:** Beginning with OES 24.1, **Verify trustees** is changed to **Verify rights**.

---

## How to verify the rights of an NCP share?

The verify rights option shows the difference in trustee's rights information between the NSS and NCP server of the specified NCP share. This action can be performed on multiple shares at a time.

- 1 In UMC, click **File Access Protocols > NCP**.
- 2 Click the Search icon and specify the server name.  
or  
Click **Browse** and select **Server Type** to list their associated servers. Select the required servers from the list, and then click **APPLY**.
- 3 Select the NCP share, click More Options  icon, and then select **Verify rights**.

## How to resync trustees an NCP share? (OES 23.4)

The resync trustees option synchronizes the trustees rights from NSS to NCP server for the selected share. This action can be performed on multiple shares at a time.

- 1 In UMC, click **File Access Protocols > NCP**.
- 2 Click the Search icon and specify the server name.  
or  
Click **Browse** and select **Server Type** to list their associated servers. Select the required servers from the list, and then click **APPLY**.
- 3 Select the NCP share, click More Options  icon, and then select **Resync trustees**.

---

**NOTE:** Beginning with OES 24.1, **Resync trustees** is changed to **Resync rights**.

---

- 4 In Resync box, click **Confirm**.

## How to resync the rights of an NCP share?

The resync rights option synchronizes the trustee's rights from NSS to NCP server of the selected NCP share. This action can be performed on multiple shares at a time.

- 1 In UMC, click **File Access Protocols > NCP**.
- 2 Click the Search icon and specify the server name.  
or  
Click **Browse** and select **Server Type** to list their associated servers. Select the required servers from the list, and then click **APPLY**.
- 3 Select the NCP share, click More Options  icon, and then select **Resync rights**.
- 4 In Resync box, click **Confirm**.

## How to enable or disable encryption on an NCP share?

- 1 In UMC, click **File Access Protocols > NCP**.
- 2 Click the Search icon and specify the server name.  
or  
Click **Browse** and select **Server Type** to list their associated servers. Select the required servers from the list, and then click **APPLY**.
- 3 Select the NCP share, click More Options **⋮** icon, and then select **Enable encryption**.
- 4 In Enable encryption box, click **Confirm**.

This enables encryption on the selected share and only encrypted connections can access this share. This can be performed on multiple volumes at a time.

You can follow the same procedure to disable encryption if it is already enabled. When encryption is disabled, all connections are allowed to access this share.

## How to enable or disable MFA on an NCP share?

- 1 In UMC, click **File Access Protocols > NCP**.
- 2 Click the Search icon and specify the server name.  
or  
Click **Browse** and select **Server Type** to list their associated servers. Select the required servers from the list, and then click **APPLY**.
- 3 Select the NCP share, click More Options **⋮** icon, and then select **Enable multi-factor authentication**.
- 4 In Enable multi-factor authentication box, click **Confirm**.

This enables multi-factor authentication on the selected share. This can be performed on multiple volumes at a time.

You can follow the same procedure to disable multi-factor authentication if it is already enabled.

## What are accessed files and how to view them? (OES 23.4)

Accessed file lists the NCP share files that are in open state by an NCP connection. These files can be closed manually.

- 1 In UMC, click **File Access Protocols > NCP**.
- 2 Click the Search icon and specify the server name.  
or  
Click **Browse** and select **Server Type** to list their associated servers. Select the required servers from the list, and then click **APPLY**.
- 3 Select the NCP share, click More Options **⋮** icon, and then select **Accessed files**.  
This displays the list of open files. This operation can be performed on multiple shares at a time.

---

**NOTE:** Beginning with OES 24.1, **Accessed files** is changed to **Open files**.

---

- 4 Select the file from the list and then click .

This performs the logical closure of the selected file on the NCP server. This can be performed on multiple files at a time.

## What are open files and how to view them?

Open files are those files that are kept in open state by an NCP connection. These files can be closed manually.

- 1 In UMC, click **File Access Protocols > NCP**.

- 2 Click the Search icon and specify the server name.

or

Click **Browse** and select **Server Type** to list their associated servers. Select the required servers from the list, and then click **APPLY**.

- 3 Select the NCP share, click More Options  icon, and then select **Open files**.

This displays the list of open files. This operation can be performed on multiple shares at a time.

- 4 Select the file from the list and then click .

This performs the logical closure of the selected file on the NCP server. This can be performed on multiple files at a time.

## What are the prerequisites for adding a secondary volume?

- ♦ Ensure that the primary volume is active before adding a secondary volume.
- ♦ The primary volume must not have any secondary volume mounted on it.
- ♦ One primary volume can have only one secondary volume.
- ♦ Primary and secondary volume operations are supported only for NSS volumes.

## How to add secondary volume?

- 1 In UMC, click **File Access Protocols > NCP**.

- 2 Click the Search icon and specify the server name.

or

Click **Browse** and select **Server Type** to list their associated servers. Select the required servers from the list, and then click **APPLY**.

- 3 Select the volume, click More Options  icon, and then select **Add secondary volume**.

You can add one secondary volume to a primary volume. On selecting multiple volumes, this option is disabled.

- 4 In **Add secondary volume**, select the secondary volume, and then click **Confirm**.

This adds the selected secondary volume to the primary volume on the server.

## How to view secondary volume?

By using advance filter, you can view DST or CIS secondary volumes. You can select the Secondary path column to view secondary volume path details.

## How to remove secondary volume?

- 1 In UMC, click **File Access Protocols > NCP**.

- 2 Click the Search icon and specify the server name.

or

Click **Browse** and select **Server Type** to list their associated servers. Select the required servers from the list, and then click **APPLY**.

- 3 You can remove one secondary volume to a primary volume at a time. On selecting multiple volumes, this option is disabled.

Select the share, click More Options  icon, and then select **Remove secondary volume**.

Multiple secondary volume removal is not supported.

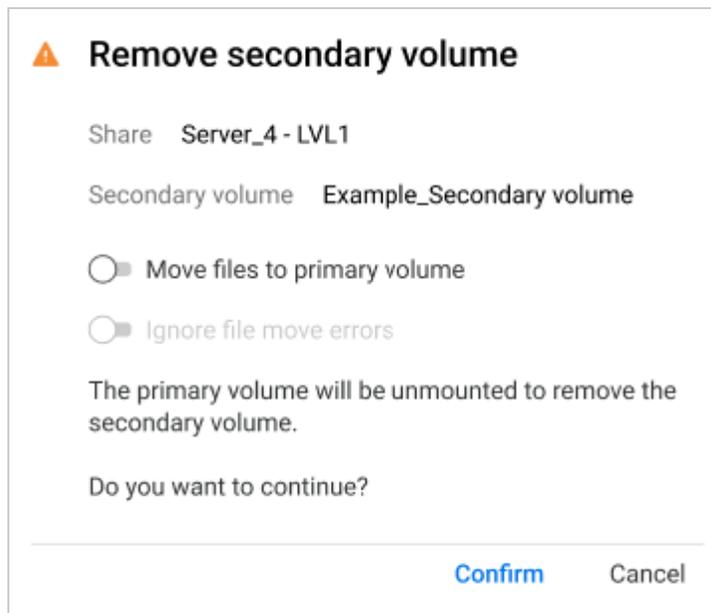
- 4 Select the required options, and then click **Confirm**.

- ◆ **Move files to primary volume**

Move all files from the secondary volume to the primary volume before removing the secondary volume.

- ◆ **Ignore file move errors**

Ignore all error messages while moving the files from secondary volume to primary volume in order to complete the process.



This removes the secondary volume from the primary volume of the server.

## How to manage security for sub-folders on an NCP share? (OES 23.4)

Encryption and multi-factor authentication are security options to manage sub-folder security in a volume.

1 In UMC, click **File Access Protocols > NCP**.

2 Click the Search icon and specify the server name.

or

Click **Browse** and select **Server Type** to list their associated servers. Select the required servers from the list, and then click **APPLY**.

3 Select the NCP share, click More Options **⋮** icon, and then select **Manage sub-folders**.

---

**NOTE:** Beginning with OES 24.1, **Manage sub-folders** is changed to **Sub-folder security**.

---

4 Select the folder, click More Options **⋮** icon, and select Encryption or Multi-factor authentication options to modify the security.

This can be performed on multiple sub-folders at a time.

## How to manage sub-folder security on an NCP share?

Encryption and multi-factor authentication are security options to manage sub-folder security in a volume.

- 1 In UMC, click **File Access Protocols > NCP**.
- 2 Click the Search icon and specify the server name.  
or  
Click **Browse** and select **Server Type** to list their associated servers. Select the required servers from the list, and then click **APPLY**.
- 3 Select the NCP share, click More Options **⋮** icon, and then select **Sub-folder security**.
- 4 In **Sub-folder security**, select the folder, click More Options **⋮** icon, and select Encryption or Multi-factor authentication options to modify the security.

This can be performed on multiple sub-folders at a time.

## How to enable or disable write permission for an NCP share?

- 1 In UMC, click **File Access Protocols > NCP**.
- 2 Click the Search icon and specify the server name.  
or  
Click **Browse** and select **Server Type** to list their associated servers. Select the required servers from the list, and then click **APPLY**.
- 3 Select the NCP share, click More Options **⋮** icon, and then select **Enable write permission**.
- 4 In Enable write box, click **Confirm**.

This enables the write permission for the selected NCP share.

Follow the same steps to disable the write permission on this share. These actions can be performed on multiple shares.

## How to activate or deactivate an NCP share?

Activate an NCP share to make it available for users and applications. To view the details of a share, it must be active. The details of deactivated shares are not available.

- 1 In UMC, click **File Access Protocols > NCP**.
- 2 Click the Search icon and specify the server name.  
or  
Click **Browse** and select **Server Type** to list their associated servers. Select the required servers from the list, and then click **APPLY**.

This displays the list of volumes available on the selected servers.

---

**3 NOTE:** If you select multiple shares, the More Options  icon is displayed at the top right corner of the table.

---

**3a To deactivate an NCP share:**

**3a1** Select the share, click More Options  icon, and then select **Deactivate**.

**3a2** In **Deactivate**, click **Confirm**.

This closes all the open connections to the selected NCP share. The files are not deleted, but the share must be active to access them.

or

**3b To activate an NCP share:**

**3b1** Select the share, click More Options  icon, and then select **Activate**.

**3b2** In **Activate**, click **Confirm**.

This activates the selected share and all the files are made available to the associated connections.



# 22 Managing NCP Connections

NCP connection is available with OES 24.1 or later version.

- ♦ “How to view NCP connections?” on page 165
- ♦ “What actions can be performed on NCP connections?” on page 165
- ♦ “How to send a broadcast message to all NCP connections?” on page 166
- ♦ “How to clear unauthenticated NCP connection?” on page 166
- ♦ “How to view open files, NCP shares, and details of an NCP connection?” on page 166
- ♦ “How to send message to an NCP connection?” on page 167
- ♦ “How to clear an NCP connection?” on page 167

## How to view NCP connections?

To view the list of NCP connections, perform the following steps:

1. In UMC, click **File Access Protocols > NCP**.
2. Click the Search icon and specify the server name.

or

Click **Browse**, select Server Type to list their associated servers. Select the required servers from the list, and then click **APPLY**.

---

**NOTE:** When the **BROWSE** or tree view  icon is clicked, you cannot perform other actions outside the browse area. Click the same button again to close the browse or tree view area.

---

3. Click **NCP > Connections**.

This displays the list of available NCP connections on the selected servers.

## What actions can be performed on NCP connections?

You can perform the following actions on the NCP connections

- ♦ **Broadcast message to all connections on selected servers**
- ♦ **Clear unauthenticated connections on selected servers**
- ♦ **Clear all connections on selected servers**

## How to send a broadcast message to all NCP connections?

You can send message to all the NCP connections using the **Broadcast message to all connections on selected servers** option in **Actions** drop-down.

1. In UMC, click **File Access Protocols > NCP**.
2. Click the Search icon and specify the server name.  
or  
Click **Browse**, select Server Type to list their associated servers. Select the required servers from the list, and then click **APPLY**.
3. Click **NCP > Connections**.
4. Click the **Actions** drop-down and select **Broadcast message to all connections on selected servers**.
5. Specify the message and click **Send**.

The Character limit of broadcast message is 256.

This delivers the specified broadcast message to all the NCP connections for the selected servers.

## How to clear unauthenticated NCP connection?

You can clear all the unauthenticated NCP connections from the list using the **Clear unauthenticated connections on selected servers** option in **Actions** drop-down.

1. In UMC, click **File Access Protocols > NCP**.
2. Click the Search icon and specify the server name.  
or  
Click **Browse**, select Server Type to list their associated servers. Select the required servers from the list, and then click **APPLY**.
3. Click **NCP > Connections**.
4. Click the **Actions** drop-down and select **Clear unauthenticated connections on selected servers**.
5. In **Clear all unauthenticated connections**, click **Confirm**.

You can follow the same steps to clear all connections. Click the **Actions** drop-down and select **Clear all connections on selected servers**.

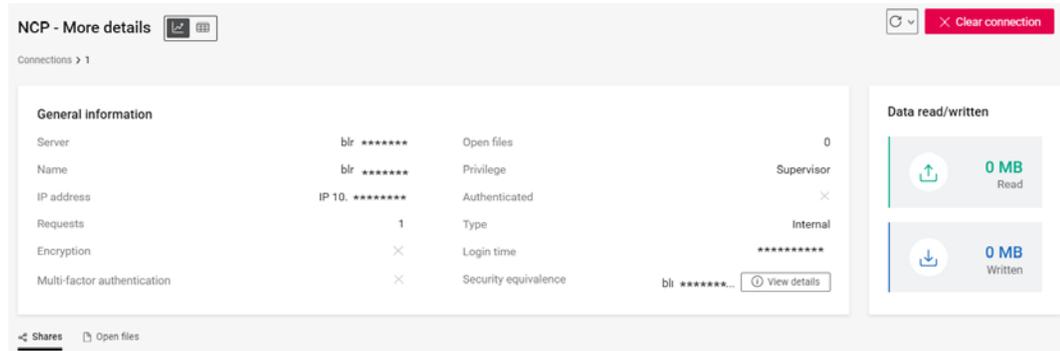
## How to view open files, NCP shares, and details of an NCP connection?

You can view the details of an NCP connection using the more details option. This includes general information, data read or write, list of related shares, and open file information.

1. In UMC, click **File Access Protocols > NCP**.
2. Click the Search icon and specify the server name.  
or

Click **Browse**, select Server Type to list their associated servers. Select the required servers from the list, and then click **APPLY**.

3. Select the NCP connection, click More Options  icon, and then select **More details**.



You can view the NCP shares associated with the connection by selecting the **Shares**.

The files that are left in open state through an NCP connection can be viewed by selecting the **Open files**. Open files option is available from OES 24.1.1 or later.

You can use the  dashboard view or  table view icons to display the NCP connections.

## How to send message to an NCP connection?

1. In UMC, click **File Access Protocols > NCP**.
2. Click the Search icon and specify the server name.

or

Click **Browse**, select Server Type to list their associated servers. Select the required servers from the list, and then click **APPLY**.

3. Click **NCP > Connections**.
4. Select the NCP connection, click More Options  icon, and then select **Send message**.
5. Specify the message and click **Send**.

The Character limit of message is 256.

This delivers the specified message to the selected NCP connection and can also be performed on multiple connections at a time.

## How to clear an NCP connection?

1. In UMC, click **File Access Protocols > NCP**.
2. Click the Search icon and specify the server name.

or

Click **Browse**, select Server Type to list their associated servers. Select the required servers from the list, and then click **APPLY**.

3. Click **NCP > Connections**.

4. Select the NCP connection, click More Options  icon, and then select **Clear connection**.
5. In **Clear connection**, click **Confirm**.

This clears the NCP connection on the selected servers and this action can also be performed on multiple connections at a time.

# 23 Managing CIFS Shares

Managing CIFS shares are available from OES 24.3 or later versions.

- ♦ “How to create a new CIFS share?” on page 169
- ♦ “How to list CIFS shares?” on page 170
- ♦ “How to remove a CIFS share?” on page 170
- ♦ “What is encryption on a CIFS share?” on page 170
- ♦ “How to manage encryption on a CIFS share?” on page 170
- ♦ “What is folder redirection on a CIFS share?” on page 172
- ♦ “What is Mac backup on a CIFS share?” on page 172
- ♦ “What is the character limit for CIFS share name and comment box?” on page 172
- ♦ “How to filter the CIFS shares?” on page 172
- ♦ “How to manage folder redirection on a CIFS share?” on page 173
- ♦ “How to manage Mac backup on a CIFS share?” on page 173
- ♦ “What are the various rights and how to manage it on CIFS shares?” on page 174
- ♦ “How to add trustees for a CIFS share?” on page 174
- ♦ “What is the CIFS share limitation that a server can host?” on page 175
- ♦ “How to modify an existing CIFS share?” on page 175
- ♦ “What are open files in a CIFS share?” on page 176
- ♦ “How to view the open files in a CIFS share?” on page 176
- ♦ “How to close open files of CIFS shares?” on page 176
- ♦ “What are the various access modes for open files?” on page 177

## How to create a new CIFS share?

1. In UMC, click **File Access Protocols > CIFS**.
2. Click **Create Share**.
3. In **Create Share** wizard > **Path**, search or browse the servers to select the volume, and click **Next**.

---

**NOTE:** You can select only one volume for creating the CIFS share.

---

4. On the **Configuration** page, specify the share name, add comment (optional), and click **Next**.  
You can manage encryption, folder redirection, and Mac backup using the respective toggle switches.
5. On the **Summary** page, verify the general information and configuration settings, and then click **Finish**.  
You can view the newly created CIFS share in the share list.

## How to list CIFS shares?

- 1 In UMC, click **File Access Protocols > CIFS**.
- 2 Click the search icon and specify the server name.

or

Click **Browse** and select server type to list their associated servers. Select the required servers from the list, and then click **APPLY**.

---

**NOTE:** When the **BROWSE** or tree view  icon is clicked, other actions outside the browse area are disabled. Click the same button again to close the browse or tree view area.

---

This displays the list of CIFS shares available on the selected servers.

## How to remove a CIFS share?

Removing a CIFS share does not delete the data in it. The association between the CIFS share and the path is revoked and cannot be restored.

---

**NOTE:** This **remove** option is available only for custom data shares.

---

1. In UMC, click **File Access Protocols > CIFS**.
2. Search or browse the servers to list the shares.
3. Select the share, click More options  icon, and then click **Remove**.

This removes the selected CIFS share from the list. You can remove multiple share at a time. You must create a new share and select the share path to access the data in it. For more information on creating a share, see [“How to create a new CIFS share?” on page 169](#).

## What is encryption on a CIFS share?

If encryption is enabled on a share, only encrypted client connections can access the share. You can enable or disable encryption on a CIFS share while creating a new share or by selecting an individual share. For more information on managing encryption, see [“How to manage encryption on a CIFS share?” on page 170](#).

Encryption can be enabled or disabled at the share level whereas if encryption is applied at the global level then no need to apply it at the share level. You can enable encryption on individual shares if encryption is disabled at the global level.

## How to manage encryption on a CIFS share?

You can enable or disable encryption on a share while creating a share or on an existing share.

## Enable encryption while creating a new share

---

**NOTE:** The system shares do not support encryption, folder redirection, and Mac backup.

---

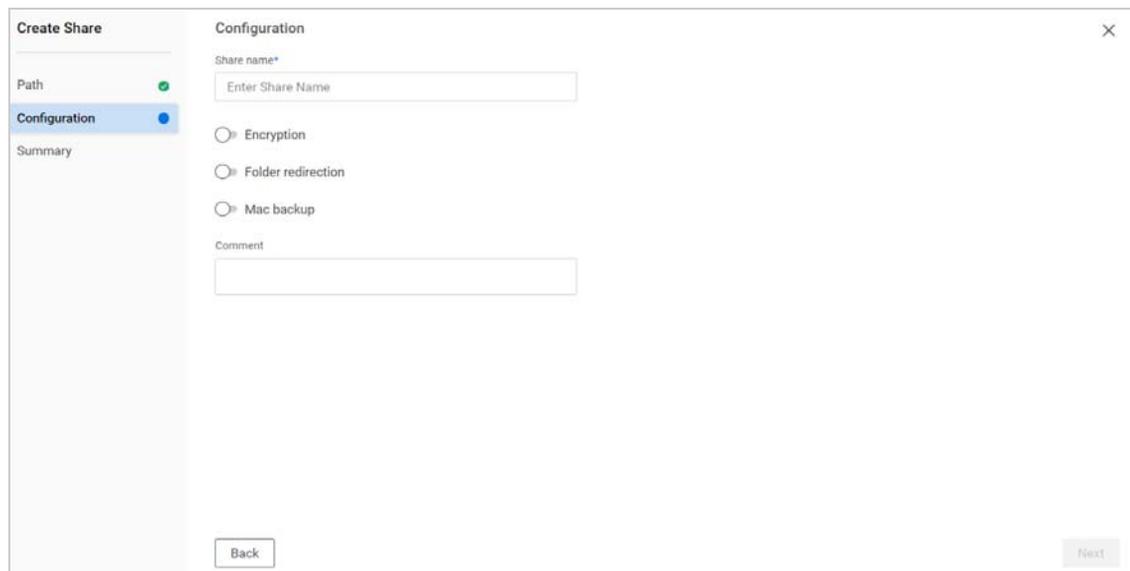
1. In UMC, click **File Access Protocols > CIFS**.
2. In the **Create Share** wizard > **Path**, search or browse the servers to select the volume, and click **Next**.

---

**NOTE:** You can select only one volume for creating the CIFS share.

---

3. On the **Configuration** page, specify the share name, add comment (optional), and click **Next**.



The screenshot shows the 'Create Share' wizard in the Configuration step. On the left, a sidebar lists 'Path' (with a green checkmark), 'Configuration' (with a blue dot), and 'Summary'. The main area is titled 'Configuration' and contains a 'Share name\*' field with the placeholder 'Enter Share Name'. Below this are three radio button options: 'Encryption', 'Folder redirection', and 'Mac backup', all of which are currently unselected. A 'Comment' text box is located below the radio buttons. At the bottom of the wizard, there are 'Back' and 'Next' buttons.

Encryption is disabled by default. You can manage encryption, folder redirection, and Mac backup using the respective toggle switches.

4. On the **Summary** page, verify the details and click **Finish**.

## Enable encryption on an existing share

1. In UMC, click **File Access Protocols > CIFS**.
2. Search or browse the servers to list the shares.
3. Select the share, click More options **⋮** icon, and then click **Enable Encryption**.

This enables encryption on the selected CIFS share. Follow the same procedure to disable encryption if it is already enabled. Encryption can be managed on multiple shares at a time.

## What is folder redirection on a CIFS share?

Folder redirection allows users and administrators to redirect the path of a folder to another location. The new location can be on a local computer or on a network file share. Users can manage the files as it is in the local directory. The files in the folder can be accessed from any computer on the network.

For more information on managing folder redirection, see [“How to manage folder redirection on a CIFS share?” on page 173.](#)

---

**NOTE:** The support for this feature is available only for AD users.

---

## What is Mac backup on a CIFS share?

The Mac backup allows users or administrators to manage the shares to back up their data on the Mac clients. The users or administrators must have read, write, create, erase, modify, and file scan permissions to perform this action. For more information on rights, see [“What are the various rights and how to manage it on CIFS shares?” on page 174.](#)

## What is the character limit for CIFS share name and comment box?

A CIFS share name can be up to 80 characters long and can contain any single byte characters, but should not begin or end with an underscore `_` or contain multiple underscores `_`.

(Optional) You can provide a description in the comment box for the CIFS share. The maximum allowed length is 47 characters.

## How to filter the CIFS shares?

You can use Advanced Filters to filter CIFS shares based on the following criteria:

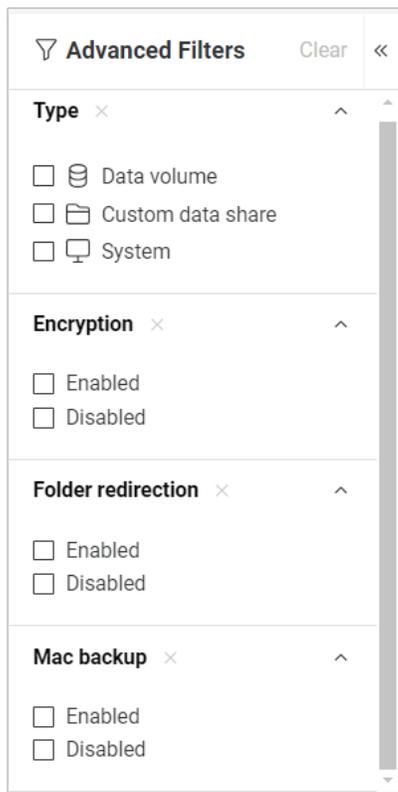
- ◆ **Type** - CIFS shares can be filtered based on the type such...
  - ◆ **Data volume** shares are shares created for normal NSS volumes.
  - ◆ **Custom data** shares are created for directories under NSS volumes.
  - ◆ **System** shares are created for some specific functionalities such as IPC\$, \_ADMIN, and so on...

---

**NOTE:** The system shares do not support encryption, folder redirection, and Mac backup.

---

- ◆ **Encryption** - CIFS shares can be filtered based on the encryption enabled or disabled state. If encryption is disabled, any connection can access the share.
- ◆ **Folder redirection** - CIFS shares can be filtered based on the folder redirection enabled or disabled state.
- ◆ **Mac backup** - CIFS shares can be filtered based on the Mac backup enabled or disabled state.



## How to manage folder redirection on a CIFS share?

---

**NOTE:** The system shares do not support encryption, folder redirection, and Mac backup.

---

Folder redirection allows you to redirect a path of one folder to another location and this path can be accessed from any computer on the network.

1. In UMC, click **File Access Protocols > CIFS**.
2. Search or browse the servers to list the shares.
3. Select the share, click More options **⋮** icon, and then select **Enable folder redirection**.
4. Click **Confirm**.

This enables the folder redirection on the selected share. You can follow the same procedure to disable it. This action can be performed on multiple shares at a time.

You can also manage folder redirection using the **Edit** option while modifying the share.

## How to manage Mac backup on a CIFS share?

---

**NOTE:** The system shares do not support encryption, folder redirection, and Mac backup.

---

1. In UMC, click **File Access Protocols > CIFS**.
2. Search or browse the servers to list the shares.

3. Select the share, click More options  icon, and then select **Enable Mac backup**.
4. Click **Confirm**.

This enables the Mac backup on the selected share. You can follow the same procedure to disable it. This action can be performed on multiple shares at a time.

You can also manage Mac backup using the **Edit** option while modifying the share.

## What are the various rights and how to manage it on CIFS shares?

You can manage trustee rights on CIFS share using the **Manage rights** option.

1. In UMC, click **File Access Protocols > CIFS**.
2. Search or browse the servers to list the shares.
3. Select the share, click More options  icon, and then select **Manage rights**.
4. On the **Manage rights** page, use the checkbox to manage the required rights.

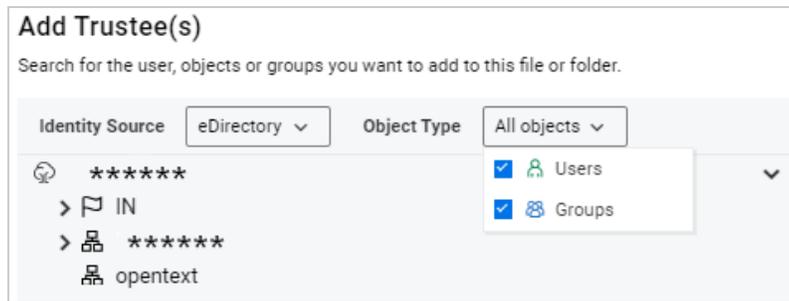
S	R	W	C	E	M	F	A
<input type="checkbox"/>	<input checked="" type="checkbox"/>						

Rights	Description
S - Supervisor	Users have all rights on the file or directory and can manage the Access Control right.
R - Read	Users can open and read files in the directory.
W - Write	Users can open and write to files in the directory.
C - Create	Users can create files and subdirectories, and can also salvage or restore them.
E - Erase	Users can erase files and directories, and can also purge or permanently delete them.
M - Modify	Users can modify the meta-data of the file or directory.
F - File Scan	Users can display and search on file and directory names in the file system structure.
A - Access Control	Users can add and remove trustees, and change trustee rights to files and directories.

## How to add trustees for a CIFS share?

1. In UMC, click **File Access Protocols > CIFS**.
2. Search or browse the servers to list the shares.

3. Select the share, click More options  icon, and then select **Manage rights**.
4. On the **Manage rights** page, click **Add Trustee**.
5. In the **Add Trustee(s)** wizard, navigate through the server tree and select the required trustees or users.



You can modify the object type using the **All objects** drop-down.

6. Click **Confirm**.

This adds the selected trustee(s) to the volume. For more information on inherited rights and effective rights, see [“What are inherited rights?” on page 127](#) and [“What are effective rights?” on page 126](#) in Chapter 15, [“Managing Rights,” on page 123](#).

## What is the CIFS share limitation that a server can host?

A server can host up to **65535** CIFS shares.

## How to modify an existing CIFS share?

You can modify the CIFS share path and configuration setting of an existing share using the **Edit** option.

---

**NOTE:** The **Edit** option is only supported on custom shares.

---

1. In UMC, click **File Access Protocols > CIFS**.
2. Search or Browse the servers to list the shares.
3. Select the share, click More options  icon, and then select **Edit**.
4. In **Edit Share** wizard > **Path**, navigate through the server tree to select the new share path, and click **Next**.

---

**NOTE:** Only one path can be selected for a share.

---

5. On the **Configuration** page, specify the share name and comment (optional), and click **Next**.  
You can manage encryption, folder redirection, and Mac backup using the toggle switches.
6. On the **Summary** page, verify the details, and then click **Finish**.

This updates the selected CIFS share path and configuration settings.

# What are open files in a CIFS share?

Open files are those files that are left in open state by a CIFS connection at share level. These files can be closed manually.

## How to view the open files in a CIFS share?

1. In UMC, click **File Access Protocols > CIFS**.
2. Search or browse the servers to list the shares.
3. Select the share, click More options **⋮** icon, and then select **Open files**.

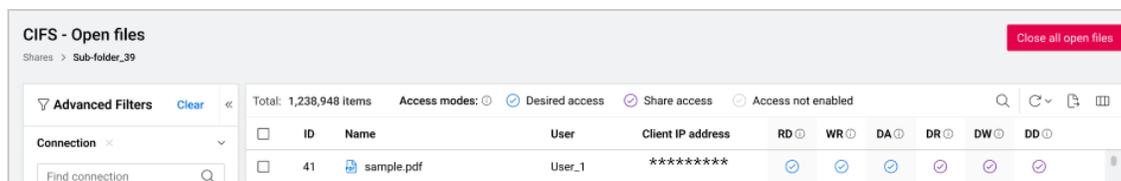
This displays the list of open files. This action is supported on a single share selection. You can view the details of the open files, related shares, users, and assigned rights.

## How to close open files of CIFS shares?

You can close all open files with the **Close all open files** option or you can close individual file or files in a CIFS share using the **✕** cross. This option allows you to manage open files in CIFS shares of multiple servers at a time.

### Close all open files

1. In UMC, click **File Access Protocols > CIFS**.
2. Search or browse the servers to list the shares.
3. Select the share, click More options **⋮** icon, and then select **Open files**.
4. To close all open files at once, click the **Close all open files** button.



This action closes all the files available in the open files list.

### Close individual open file

1. In UMC, click **File Access Protocols > CIFS**.
2. Search or browse the servers to list the shares.
3. Select the share, click More options **⋮** icon, and then select **Open files**.

4. On the **Open files** page, select and click the  cross icon to close an individual file.



You can close multiple open files at a time.

## What are the various access modes for open files?

Details of the CIFS connection includes the access modes in which the CIFS server opened the file for the user.



Desired Access	Description	Shared Access	Description
RD	Right to read data from the file.	DR	Right to read data from the file is denied.
WR	Right to write data into the file.	DW	Right to write data into the file is denied.
DA	Right to delete the file.	DD	Right to delete or rename the file is denied.



# 24 Managing CIFS Connections

This chapter consists of FAQs for viewing the CIFS connections, open files, associated shares, and security equivalence of the connection.

- ♦ “How to list and view the information related to CIFS connections?” on page 179
- ♦ “How to view the open files of a CIFS connection?” on page 180
- ♦ “How to view the shares associated with a CIFS connection?” on page 180
- ♦ “How to view the security equivalence of a CIFS connection?” on page 181

## How to list and view the information related to CIFS connections?

To list and view the information related to a CIFS connections, perform the following:

1. In UMC, click **File Access Protocols > CIFS > Connections**.
2. Click the search icon and specify the server name.

or

Click **Browse**, select **Server Type** to list their associated servers. Select the required servers from the list, and then click **APPLY**.

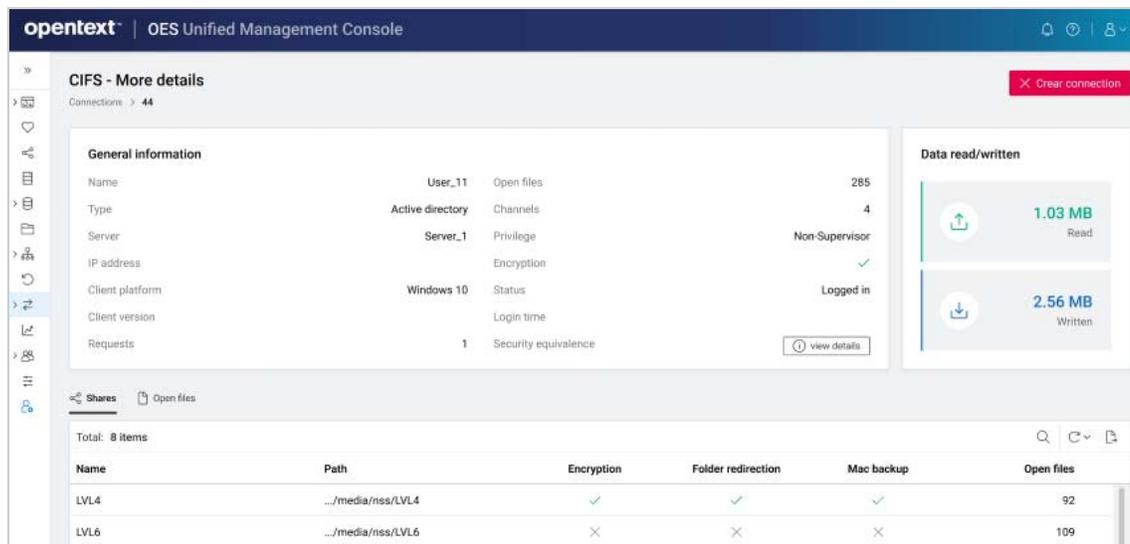
---

**NOTE:** When the  or tree view  icon is clicked, you cannot perform other actions outside the browse area. Click the same button again to close the browse or tree view area.

---

This displays the list of available CIFS connections on the selected servers.

3. To view the information related to a CIFS connection, select a connection and click **More details** icon.



The **CIFS - More details** page displays the general information, data managed, shares, and open files of the CIFS connection.

## How to view the open files of a CIFS connection?

To view the open files of a CIFS connection, perform the following:

1. In UMC, click **File Access Protocols > CIFS > Connections**.
2. Search or browse the servers to list the connections.
3. Select a connection and click **More details** icon.
4. On the **CIFS - More details** page, click **Open files** tab.



The page displays the list of open files accessed by the CIFS connection. For more information on access modes, see [“What are the various access modes for open files?”](#) on page 177.

## How to view the shares associated with a CIFS connection?

To view the shares associated with a CIFS connection, perform the following:

1. In UMC, click **File Access Protocols > CIFS > Connections**.
2. Search or browse the servers to list the connections.
3. Select a connection and click **More details** icon.
4. On the **CIFS - More details** page, click **Shares**.

Name	Path	Encryption	Folder redirection	Mac backup	Open files
VOL1	/media/nss/VOL1	x	x	x	0

The page displays the list of shares accessed by the CIFS connection.

## How to view the security equivalence of a CIFS connection?

To view the security equivalence of a CIFS connection, perform the following:

1. In UMC, click **File Access Protocols > CIFS > Connections**.
2. Search or browse the servers to list the connections.
3. Select a connection and click **More details** icon.
4. On the **CIFS - More details** page, click **View details** adjacent to security equivalence field.

The page displays **Security equivalence for:** window with the users and FQDN details for the CIFS connection.



# 25 Managing Invalid Users

This chapter consists of FAQs for viewing, adding, removing, and updating invalid users and blocked users.

- ♦ “How to list invalid users and blocked users?” on page 183
- ♦ “Who is an invalid user?” on page 183
- ♦ “Who is a blocked user?” on page 184
- ♦ “How to add a user to the blocked user list?” on page 184
- ♦ “How to unblock an invalid user?” on page 184
- ♦ “How to unblock a blocked user?” on page 184
- ♦ “How to change an invalid user to a blocked user?” on page 184

## How to list invalid users and blocked users?

The **Invalid users** tab supports only a single server selection. If multiple servers are selected during other CIFS operations, and you select the **Invalid users** tab, an empty page is displayed.



1. In UMC, click **File Access Protocols > CIFS > Invalid users**.
2. Click the search icon and specify the server name.  
or  
Click **Browse**, select the required server from the list, and then click **APPLY**.

---

**NOTE:** When the **BROWSE** or tree view icon is clicked, you cannot perform other actions outside the browse area. Click the same button again to close the browse or tree view area.

---

This displays the list of invalid users and blocked users on the selected server.

## Who is an invalid user?

An invalid user can be a user who does not exist in eDirectory or an admin added him to the invalid users list. The authentication request from this user is ignored based on the configured timeout period. The timeout period of an invalid is between 0 and 525600 minutes.

## Who is a blocked user?

A blocked user is a user whose authentication request is ignored permanently. Unblock the blocked user from the list to start considering authentication requests.

## How to add a user to the blocked user list?

1. In UMC, click **File Access Protocols > CIFS > Invalid users**.
2. Search or browse the server to list the invalid users.
3. Click **Block user**.
4. Specify the user name and click **Confirm**.

This adds the user to the blocked user list.

## How to unblock an invalid user?

Unblocking an invalid user allows the authentication request to be processed for the user.

1. In UMC, click **File Access Protocols > CIFS > Invalid users**.
2. Search or browse the server to list the invalid users.
3. Select an invalid user, click More options  icon, and then select **Unblock**.



This unblocks the selected invalid user from the list. You can unblock multiple users at a time.

## How to unblock a blocked user?

Unblocking a blocked user allows the authentication request to be processed for the user.

1. In UMC, click **File Access Protocols > CIFS > Invalid users**.
2. Search or browse the server to list the blocked users.
3. Select a blocked user and click **Unblock**.



This unblocks the selected blocked user from the list. You can unblock multiple users at a time.

## How to change an invalid user to a blocked user?

1. In UMC, click **File Access Protocols > CIFS > Invalid users**.
2. Search or browse the server to list the invalid users.

3. Select an invalid user and then click **Block user**.



This updates the selected invalid user as a blocked user.



# 26 Managing User Context

This chapter consists of FAQs for viewing, adding, and removing user contexts.

- ♦ “How to list the user contexts?” on page 187
- ♦ “How to add a user context?” on page 187
- ♦ “How to remove a user context?” on page 187

## How to list the user contexts?

User context is an eDirectory container where CIFS search for users during login.

1. In UMC, click **File Access Protocols > CIFS > User contexts**.
2. Click the search icon and specify the server name.

or

Click **Browse**, select Server Type to list their associated servers. Select the required servers from the list, and then click **APPLY**.

---

**NOTE:** When the  or tree view  icon is clicked, you cannot perform other actions outside the browse area. Click the same button again to close the browse or tree view area.

---

This displays the list of available user contexts on the selected servers.

## How to add a user context?

Adding a user context allows the users in the container to access the share.

1. In UMC, click **File Access Protocols > CIFS > User contexts**.
2. Click **Add user context**.
3. In the **Add user context** wizard > **Servers**, select the servers from the directory tree and click **Next**.
4. On the **Containers** page, browse the server to select the containers and click **Next**.
5. On the **Summary** page, verify the containers and servers, and then click **Finish**.

This adds the user context to the list. You can add multiple contexts at a time.

## How to remove a user context?

Removing a user context restricts the users in the container from accessing the share.

1. In UMC, click **File Access Protocols > CIFS > User contexts**.
2. Search or browse the servers to list the user contexts.

3. Select a user context and click Remove  icon.

This removes the user context from the list. You can perform this action on multiple selection.

# IX Reports

- ◆ [Chapter 27, “Cluster Reports,” on page 191](#)



# 27 Cluster Reports

A cluster report helps to diagnose problems with the cluster nodes and resources.

- ♦ [“How to generate a cluster report?” on page 191](#)
- ♦ [“How to view reports?” on page 191](#)
- ♦ [“Report Failures” on page 192](#)

## How to generate a cluster report?

- 1 Log in to UMC with your admin credentials.
- 2 Click **Clusters**.
- 3 Select a cluster, then select **Run report**. Alternatively, you can access this option from the dashboard by clicking **Actions > Run report**.

The report includes information of the selected cluster, such as its current cluster configuration, cluster nodes, cluster resources, each cluster resource's policies and load, unload, and monitor scripts, and resource mutual exclusion groups.

## How to view reports?

- 1 Log in to UMC with your admin credentials.
- 2 Click **Reports**.
- 3 Browse and select the cluster objects you want to view. The reports are listed for those objects on which you have generated reports by using **Run report** option in **Clusters**.
- 4 The following information is displayed:

Column Name	Description
Status (Color Coding)	Status
Green	<b>Available:</b> The reports have been successfully generated.
Blue	<b>In progress:</b> The report generation has been triggered but not yet completed.
Red	<b>Failed:</b> The report generation failed. Try generating the report after some time.

- 5 Select the cluster object and click **Open report**. The report displays cluster status, resource mutual exclusion groups, cluster options, and cluster resources.

## Report Failures

When running a report on a cluster, if the report generation is not initiated, it could be due to the Redis service not being active.

To verify the status of the service, run:

```
systemctl status redis@umc.service
```

If the service is inactive, restart it:

```
systemctl restart redis@umc.service
```

After restarting the service, run the report again and verify the status of the report in the Reports page.

# X Troubleshooting

- ◆ [“Known Issues” on page 195](#)
- ◆ [“UMC Health Script” on page 196](#)
- ◆ [“Missing Node Modules” on page 198](#)
- ◆ [“Unable to Connect to the Database” on page 198](#)
- ◆ [“Warning: Entered Hostname is Incorrect” on page 199](#)
- ◆ [“Volume Issues” on page 199](#)
- ◆ [“Login Failures” on page 199](#)
- ◆ [“Action to Perform in case of Cache-Related Issues” on page 200](#)
- ◆ [“Cluster Issues” on page 200](#)



# 28 Troubleshooting

This section presents the information on some of the troubleshooting issues in UMC.

- ♦ [“Known Issues” on page 195](#)
- ♦ [“UMC Health Script” on page 196](#)
- ♦ [“Missing Node Modules” on page 198](#)
- ♦ [“Unable to Connect to the Database” on page 198](#)
- ♦ [“Warning: Entered Hostname is Incorrect” on page 199](#)
- ♦ [“Volume Issues” on page 199](#)
- ♦ [“Login Failures” on page 199](#)
- ♦ [“Action to Perform in case of Cache-Related Issues” on page 200](#)
- ♦ [“Cluster Issues” on page 200](#)

## Known Issues

- ♦ NCP shares or connections are not listed when any of the selected virtual servers are unreachable, offline, or comatose.  

To resolve this issue, ensure that the virtual servers associated with NCP shares or connections are reachable and online before attempting to list them.
- ♦ A top-level context cannot be added despite the presence of user objects at a deep level. As a workaround, add any other sub-container such as DC, O, or OU to include user context.
- ♦ On the **Configure > Servers > Server Settings** page, settings are modified for only one or first server, even if two servers are displayed in the **FILTERS** section.
- ♦ If the DNS record is not updated with the hostname to the IP address of the UMC server, then the UMC server is not listed on the OES Welcome page. To resolve this issue, add the IP address and hostname to the DNS record.
- ♦ If the security status of an NCP sub-folder is updated, the page fails to display the sub-folder list. You must manually refresh the sub-folder security list in UMC by using the refresh icon to fetch the updated list.
- ♦ If the CIFS service is not available on the cluster node, UMC is unable to fetch the connections in the cluster dashboard. The node is grayed out, and no actions can be performed on the node through UMC.
- ♦ After taking a resource online or offline, you must manually refresh the resource table to view the updated status.
- ♦ Salvaging file fails if a file with the same name exists in the associated folder.
- ♦ Creating a pool on a shared device is not allowed if the setup is not cluster configured.
- ♦ If the pool object already exists, joining the pool to an AD domain will fail. Clean up the object in the active directory and try again.

- ◆ When the `BROWSE` or tree view `📁` icon is clicked, you cannot perform other actions outside the browse area. Click the same button again to close the browse or tree view area.
- ◆ If you cannot browse through UMC after logging into it, ensure that the Compare, Read, and Write permissions on **All Attributes Rights** and the Browse permission on **Entry Rights** are enabled at the tree level for the logged-in users.
- ◆ If the UMC screen is not properly displayed or scaled on the web browser, ensure to set the display resolution to 1920 x 1080 or 1920 x 927 and the zoom level to 100%.

## UMC Health Script

The `umcServiceHealth` Script verifies the health of the UMC server and all services running on the server.

### Syntax

```
umcServiceHealth [options]
```

### Options

Options	Description
<code>-h --help</code>	Displays the help screen.
<code>-s --service-check</code>	Verifies the health of the dependent services. The services are: <ul style="list-style-type: none"> <li>◆ <code>apache2.service</code></li> <li>◆ <code>postgresql.service</code></li> <li>◆ <code>ndsd.service</code></li> <li>◆ <code>microfocus-umc-server.service</code></li> <li>◆ <code>microfocus-umc-backend.service</code></li> <li>◆ <code>docker.service</code></li> <li>◆ <code>docker-edirapi.service</code></li> <li>◆ <code>redis@umc.service</code></li> </ul>
<code>-e --edirapi-check</code>	Verifies the health of edirapi and Identity Console container ( <code>identityconsole-oes</code> ).
<code>-c --cert-check</code>	Verifies the health of the server certificate, displaying details such as: <ul style="list-style-type: none"> <li>◆ Certificate expiry date</li> <li>◆ Public Key SAN details</li> <li>◆ Private Key status</li> </ul>
<code>-u --edirObj-check</code>	Verifies the health of the <code>umcConfig</code> object under the security context of the eDirectory.
<code>-d --db-check</code>	Verifies the health of the PostgreSQL database (internal or remote). <p><b>IMPORTANT:</b> <code>umcServiceHealth.sh -dautofix</code> - mention that <code>.sh</code> is not required. It will be cleaned up soon.</p>

---

Options	Description
-n --nodeModule-check	Verifies the availability of the <code>node_modules</code> folder.
-r --redis-check	Verifies the health of <code>Redis</code> .
-a --all-check	Verifies the health of the UMC server and run the other checks.

---

## Autofix

The `autofix` script automatically fixes detected issues without your intervention. If an issue is encountered upon executing the health script, to resolve it, run the same script with the `autofix` option enabled. This option can be used along with:

- ◆ `--service-check`
- ◆ `--db-check`
- ◆ `--nodeModule-check`
- ◆ `--all-check`

The `autofix` script does not resolve issues related to critical components like `eDirectory` and server certificate. Hence, these options (`--edirapi-check`, `--cert-check`, and `--edirObj-check`) are not supported as they require your validation and intervention for proper resolution.

## Examples

- ◆ To verify the health of the dependent services, run the script:

```
umcServiceHealth -s
```

Displays the status of dependent services on this server. The `Apache` service is down, it displays its state and command to restart the service. Alternatively, you can rerun this command with `autofix` to resolve the issue.

Figure 28-1 umcServiceHealth Script

```
***** ~ # umcServiceHealth -s
#####
[UMC Server Health Check]
Script executed on: [2024-09-11:13:17:25:IST]
#####

[Service Status Check]
postgresql.service is active
mysqld.service is active
microfocus-umc-server.service is active
microfocus-umc-backend.service is active
docker.service is active
docker-tc.service is active
redis@umc.service is active
apache2.service is inactive
ERROR: apache2 is in inactive state. Make sure service is up and running by executing systemctl restart apache2.service

#####
#####

[Apache Module Check]
apache headers module is enabled
apache proxy_http module is enabled

#####
#####

[Summary]
Service Status Check: 1 issue found

#####
#####

END on : [2024-09-11:13:17:25:IST]
Script Execution Time: 0.22 seconds
#####
```

To fix the Apache issue automatically, run the script with `autofix` option.

```
umcServiceHealth -sautofix
```

The Apache service is successfully restarted.

- ◆ To verify health of Redis, run the script:

```
umcServiceHealth -r
```

This displays the permissions Redis has on certificate files and health of the parameters in the `/etc/redis/umc.conf` file.

To resolve the Redis issues listed by the script, re-install the `microfocus-oes-umc-server` RPM. This regenerates the `/etc/redis/umc.conf` file to fix the issues.

## Missing Node Modules

This issue arises because the node module is corrupted or the node folder is missing.

To resolve this issue, run the health script (`umcServiceHealth`) with `autofix` option.

```
umcServiceHealth -sautofix
```

## Unable to Connect to the Database

In the UMC status file, an error is logged stating, “Unable to connect to the database.”

This issue can occur if in the UMC login screen, the `treename` field is empty because it is unable to get the details from the database.

Run the health script (`umcServiceHealth`) to verify the status and resolve the issue.

## Warning: Entered Hostname is Incorrect

During UMC configuration, when specifying the database details, a warning is displayed indicating that hostname is incorrect. This issue arises because of an incorrect DNS record, which prevents the database from being reached. The `y2log` file logs a message stating, "Could not translate host name to address."

To resolve this issue, ensure that the hostname provided is resolvable by the DNS.

## Volume Issues

### Failing to List Pools or Volumes

Make sure that the backend service is working properly. Use the command `systemctl status microfocus-umc-backend.service`.

### Unable to Perform Storage Operations as an Admin Equivalent User

Try performing `/ForceSecurityEquivalenceUpdate` from the NSS console.

### Creating Volume with AES 256 Encryption Failed

Before creating a volume, perform `/PoolMediaUpgrade=pool_name /MediaType=AES` from NSS console.

## Login Failures

If you cannot login to UMC, run the health script (`umcServiceHealth`) to verify the status of the services and resolve the issues.

Alternatively you can perform this tasks manually by verifying the services `edirapi` container, `microfocus-umc-server`, and `postgresql`.

Run the following commands to verify the status:

- ♦ `systemctl status docker-edirapi.service`
- ♦ `systemctl status microfocus-umc-server.service`
- ♦ `systemctl status postgresql.service`

Run the following commands to restart services:

```
systemctl restart docker.service
```

```
systemctl restart docker-edirapi.service
```

## Action to Perform in case of Cache-Related Issues

Make sure to clear the browser cookies or perform UMC operations from a private window.

## Cluster Issues

### Renaming Cluster Pool or Volume Failed

Renaming a cluster pool or volume may show inconsistent behavior. If you are unable to list pools or volumes after renaming, open UMC from another window in incognito mode.

### Status of Healthy Cluster is Down or Unknown

If the status of a healthy cluster is `Down` or `Unknown`, then increase the timeout value `CLUSTER_LISTING_FAILURE_TIMEOUT = 2000` in the `/opt/novell/umc/apps/umc-server/prod.env` file. The default value is 2000 ms and due to network latency, it might not be able to retrieve the correct status of the cluster. Also, if this parameter is missing in the `prod.env` file, ensure to add it so cluster listing timeout occurs after the specified time.