ONTAP® 9

# Cluster Management Using OnCommand® System Manager

June 2021 | 215-12668_2021-06_en-us
doccomments@netapp.com

Updated for ONTAP 9.4

NetApp®

# Contents

**Understanding System Manager** ..... **13**

**Icons used in the application interface** ..... **14**

**System Manager window layouts** ..... **15**

**Window layout customization** ..... **16**

**System Manager enhancements** ..... **17**

**Setting up your cluster environment** ..... **19**
- Setting up the cluster by using OnCommand System Manager ..... 19
  - Setting up a cluster by using the template file ..... 21
  - Setting up the cluster manually ..... 23
- Accessing a cluster by using the OnCommand System Manager browser-based graphic interface ..... 28
- Configuring System Manager options ..... 29
- Viewing OnCommand System Manager log files ..... 29
  - How system logging works ..... 30
- Configuring a cluster by using System Manager ..... 30
  - Accessing a cluster by using the OnCommand System Manager browser-based graphic interface ..... 30
  - Setting up the cluster ..... 31
  - Setting up the network ..... 36
  - Setting up physical storage ..... 37
  - Setting up logical storage ..... 43

**Setting up SAML authentication** ..... **58**
- Enabling SAML authentication ..... 58
- Disabling SAML authentication ..... 59

**Setting up peering** ..... **61**
- Prerequisites for cluster peering ..... 61
- Creating intercluster LIFs ..... 62
- Creating cluster peer relationships ..... 62
- Creating SVM peers ..... 63
- What passphrases are ..... 64

**Managing clusters** ..... **65**
- Understanding quorum and epsilon ..... 65
- Dashboard window ..... 66

Monitoring a cluster using the dashboard....67
Applications....67
Provisioning a basic template....68
Adding applications to System Manager....69
Application provisioning settings....69
Configuration update....72
Configuring the administration details of an SVM....72
Configuration Updates window....73
Service Processors....74
Assigning IP addresses to Service Processors....74
Editing Service Processor settings....74
Understanding the Service Processor....75
Service Processors window....75
Cluster peers....76
Modifying the cluster peer passphrase....76
Modifying the peer network parameters....77
Deleting cluster peer relationships....77
Peers window....77
High availability....78
High Availability window....78
Licenses....79
Deleting licenses....79
License types and entitlement risk....80
Licenses window....81
Cluster Expansion....83
Adding nodes to a cluster....84
Configuring the network details of the nodes....84
Updating clusters....85
Obtaining ONTAP software images....85
Updating single-node clusters disruptively....86
Updating a cluster nondisruptively....87
How to update a cluster nondisruptively....89
Cluster Update window....89
Date and time settings of a cluster....91
Date and Time window....91
SNMP....91
Enabling or disabling SNMP....92
Setting SNMP information....92
Enabling or disabling SNMP traps....92
Testing the trap host configuration....93
SNMP window....93
LDAP....94
Viewing the LDAP client configuration....94
Using LDAP services....94
LDAP window....95
Users....95
Adding a cluster user account....96
Editing a cluster user account....96
Changing passwords for cluster user accounts....96
Locking or unlocking cluster user accounts....96
User accounts (cluster administrators only)....97
Roles....97
Users window....97
Roles....98

Adding roles....98
Editing roles....98
Roles and permissions....99
Roles window....99

**Managing the network....100**
IPspaces....100
Editing IPspaces....100
Deleting IPspaces....100
Broadcast domains....100
Editing broadcast domain settings....101
Deleting broadcast domains....101
Subnets....102
Editing subnet settings....102
Deleting subnets....102
Network interfaces....103
Creating network interfaces....103
Editing network interface settings....105
Deleting network interfaces....105
Migrating a LIF....106
Ethernet ports....106
Creating interface groups....106
Creating VLAN interfaces....107
Editing Ethernet port settings....107
Editing interface group settings....107
Modifying the MTU size of a VLAN....108
Deleting VLANs....108
Ports and adapters....108
FC/FCoE and NVMe adapters....109
Editing the FC/FCoE and NVMe adapter speed settings....109
Network window....109

**Managing physical storage....117**
Storage tiers....117
Editing aggregates....117
Deleting aggregates....118
Changing the RAID configuration when creating an aggregate....118
Provisioning cache by adding SSDs....119
Adding capacity disks....121
Changing the RAID group when adding capacity disks....122
Moving FlexVol volumes....123
Mirroring aggregates....123
Viewing aggregate information....124
Installing a CA certificate if you use StorageGRID Webscale....124
How moving a FlexVol volume works....124
How you can use effective ONTAP disk type for mixing HDDs....125
What compatible spare disks are....125
How System Manager works with hot spares....126
Rules for displaying disk types and disk RPM....126
How mirrored aggregates work....126
What a FabricPool is....127
Storage recommendations for creating aggregates....128

Storage Tiers window.... 129
Configuring and managing external capacity tiers.... 130
Adding an external capacity tier.... 131
Attaching an aggregate to an external capacity tier.... 132
Provisioning storage by creating a FabricPool-enabled aggregate manually.... 133
Changing the tiering policy of a volume.... 134
Editing an external capacity tier.... 134
Deleting an external capacity tier.... 134
What external capacity tiers and tiering policies are.... 135
What inactive (cold) data is.... 135
External Capacity Tier window.... 135
Aggregates.... 136
Aggregates window .... 136
Storage pools.... 139
Creating a storage pool.... 139
Adding disks to a storage pool.... 140
Deleting storage pools.... 140
How to use SSD storage pools.... 141
Requirements and best practices for using SSD storage pools.... 141
Considerations for when to use SSD storage pools.... 142
Considerations for adding SSDs to an existing storage pool versus creating a new one.... 143
Why you add disks to storage pools.... 143
How storage pool works.... 143
Storage Pools window.... 144
Disks.... 145
Reassigning disks to nodes.... 145
Viewing disk information.... 146
How ONTAP reports disk types.... 146
Minimum number of hot spares required for disks.... 147
Spare requirements for multi-disk carrier disks.... 148
Shelf configuration requirements for multi-disk carrier storage shelves.... 148
How to determine when it is safe to remove a multi-disk carrier.... 148
Considerations for sizing RAID groups.... 148
Disks window .... 149
Array LUNs.... 151
Assigning array LUNs.... 151
Reassigning spare array LUNs to nodes.... 152
Zeroing spare array LUNs.... 152
About disks and array LUNs.... 153
Array LUNs window.... 153
Nodes.... 154
Initializing the ComplianceClock time.... 154
Nodes window.... 154
Flash Cache.... 155
Enabling or disabling Flash Cache.... 155
How Flash Cache works.... 156
Flash Cache window.... 156
Events.... 156
Events window.... 157
System alerts.... 158
Acknowledging system health alerts.... 158
Suppressing system health alerts.... 158
Deleting system health alerts.... 158
Available cluster health monitors.... 159

Ways to respond to system health alerts....159
System Alerts window....160
AutoSupport notifications....161
Setting up AutoSupport notifications....161
Enabling or disabling AutoSupport settings....161
Adding AutoSupport email recipients....161
Testing AutoSupport settings....162
Generating AutoSupport data....162
Viewing AutoSupport summary....162
AutoSupport severity types....163
AutoSupport window....163
Jobs....164
Jobs....164
Job window....164
Flash Pool statistics....165
Flash Pool aggregate Statistics window....165

**Managing logical storage....166**
Storage Virtual Machines....166
SVM Dashboard window....166
Monitoring SVMs....167
Editing SVM settings....167
Deleting SVMs....168
Starting SVMs....169
Stopping SVMs....169
Managing SVMs....169
Types of SVMs....170
Why you use SVMs....170
How ONTAP name service switch configuration works....171
Storage Virtual Machines window....172
Volumes....174
Editing volume properties....174
Editing data protection volumes....175
Deleting volumes....176
Creating FlexClone volumes....176
Creating FlexClone files....177
Splitting a FlexClone volume from its parent volume....178
Viewing the FlexClone volume hierarchy....178
Changing the status of a volume....179
Viewing the list of saved Snapshot copies....179
Creating Snapshot copies outside a defined schedule....180
Setting the Snapshot copy reserve....180
Hiding the Snapshot copy directory....181
Scheduling automatic creation of Snapshot copies....181
Restoring a volume from a Snapshot copy....181
Extending the expiry date of Snapshot copies....182
Renaming Snapshot copies....183
Deleting Snapshot copies....183
Resizing volumes....184
Enabling storage efficiency on a volume....184
Changing the deduplication schedule....185
Running deduplication operations....186
Moving FlexVol volumes between aggregates or nodes....186

Assigning volumes to Storage QoS.... 187
Creating a mirror relationship from a source SVM.... 189
Creating a vault relationship from a source SVM.... 191
Creating a mirror and vault relationship from a source SVM.... 192
Creating an NFS datastore for VMware.... 193
Changing the tiering policy of a volume.... 193
Creating FlexGroup volumes.... 193
Editing FlexGroup volumes.... 194
Resizing FlexGroup volumes.... 195
Changing the status of a FlexGroup volume.... 196
Deleting FlexGroup volumes.... 196
Viewing FlexGroup volume information.... 197
What NetApp Volume Encryption is.... 197
Snapshot configuration.... 198
How volume guarantees work for FlexVol volumes.... 198
FlexClone volumes and space guarantees.... 198
Thin provisioning for greater efficiencies using FlexVol volumes.... 199
Using space reservations with FlexVol volumes.... 199
Benefits of storage efficiency.... 200
Data compression and deduplication.... 201
Guidelines for using deduplication.... 201
Options for resizing volumes.... 202
Considerations when moving volumes.... 202
Volumes window.... 203
Junction Path.... 209
Mounting volumes.... 209
Unmounting FlexVol volumes.... 210
Changing export policies.... 210
Junction Path window.... 210
Shares.... 211
Creating a CIFS share.... 211
Stopping share access.... 212
Creating home directory shares.... 212
Editing share settings.... 212
How ONTAP enables dynamic home directories.... 213
Shares window.... 214
LUNs.... 216
Creating FC SAN optimized LUNs.... 216
Application-specific LUN settings.... 217
Creating LUNs.... 220
Deleting LUNs.... 222
Creating initiator groups.... 223
Deleting initiator groups.... 223
Adding initiators.... 223
Deleting initiators from an initiator group.... 224
Creating port sets.... 224
Deleting port sets.... 224
Cloning LUNs.... 225
Editing LUNs.... 225
Bringing LUNs online.... 226
Taking LUNs offline.... 226
Moving LUNs.... 226
Assigning LUNs to storage QoS.... 227
Editing initiator groups.... 229

Editing initiators.... 229
Editing port sets.... 229
Viewing LUN information.... 230
Viewing initiator groups.... 230
Guidelines for working with FlexVol volumes that contain LUNs.... 230
Understanding space reservations for LUNs.... 231
Guidelines for using LUN multiprotocol type.... 231
Understanding LUN clones .... 233
Initiator hosts.... 233
igroup name.... 233
igroup type.... 233
igroup ostype.... 233
LUNs window.... 233
Qtrees.... 237
Creating qtrees.... 238
Deleting qtrees.... 238
Editing qtrees.... 239
Assigning export policies to qtrees.... 239
Viewing qtree information.... 239
Qtree options.... 240
Qtrees window.... 240
Quotas.... 241
Creating quotas.... 241
Deleting quotas.... 242
Editing quota limits.... 242
Activating or deactivating quotas.... 243
Resizing quotas.... 243
Viewing quota information.... 243
Types of quotas.... 244
Quota limits.... 244
Quota management.... 245
Quotas window .... 245
CIFS protocol.... 247
Setting up CIFS.... 247
Editing the general properties for CIFS.... 248
Adding home directory paths.... 248
Deleting home directory paths.... 249
Resetting CIFS domain controllers.... 249
Updating the CIFS group policy configuration.... 250
Enabling or disabling a CIFS group policy configuration.... 250
Reloading CIFS group policy.... 250
Setting up BranchCache.... 250
Modifying the BranchCache settings.... 251
Deleting the BranchCache configuration.... 252
Adding preferred domain controllers.... 252
Editing preferred domain controllers.... 252
Deleting preferred domain controllers.... 252
Viewing CIFS domain information.... 253
CIFS window.... 253
NFS protocol.... 255
Editing NFS settings.... 255
NFS window.... 255
NVMe protocol.... 256
Setting up NVMe.... 256

Creating an NVMe namespace....257
Editing an NVMe namespace....258
Cloning an NVMe namespace....258
Starting and stopping the NVMe service....259
What NVMe is....259
NVMe namespaces....260
NVMe subsystems....260
NVMe namespaces window....260
iSCSI protocol....262
Creating iSCSI aliases....262
Enabling or disabling the iSCSI service on storage system interfaces....262
Adding the security method for iSCSI initiators....263
Editing default security settings....264
Editing initiator security....264
Changing the default iSCSI initiator authentication method....265
Setting the default security for iSCSI initiators....265
Starting or stopping the iSCSI service....265
Viewing initiator security information....266
iSCSI window....266
FC/FCoE protocol....267
Starting or stopping the FC or FCoE service....267
Changing an FC or FCoE node name....268
The FCoE protocol....268
FC/FCoE window....268
Export policies....269
Creating an export policy....269
Renaming export policies....269
Deleting export policies....270
Adding rules to an export policy....270
Modifying export policy rules....271
Deleting export policy rules....271
How export policies control client access to volumes or qtrees....271
Export Policies window....272
Efficiency policies....273
Adding efficiency policies....273
Editing efficiency policies....273
Deleting efficiency policies....274
Enabling or disabling efficiency policies....274
What an efficiency policy is....274
Understanding predefined efficiency policies....274
Efficiency Policies window....275
Protection policies....276
Creating protection policies....276
Deleting protection policies....276
Editing protection policies....277
Protection Policies window....277
QoS policy groups....278
Creating QoS policy groups....278
Deleting QoS policy groups....278
Editing QoS policy groups....279
Managing workload performance by using Storage QoS....279
How Storage QoS works....280
How the maximum throughput limit works....281
Rules for assigning storage objects to policy groups....281

QoS Policy Groups window.... 282
NIS services.... 283
Adding NIS domains.... 283
Editing NIS domains.... 284
NIS window .... 284
LDAP client services.... 284
Adding an LDAP client configuration.... 284
Deleting an LDAP client configuration.... 285
Editing an LDAP client configuration.... 285
LDAP Client window.... 286
LDAP configuration services.... 286
Editing active LDAP clients.... 287
Deleting active LDAP clients.... 287
LDAP Configuration window.... 287
Kerberos realm services.... 288
Creating a Kerberos realm configuration.... 288
Editing a Kerberos realm configuration.... 289
Deleting Kerberos realm configurations.... 289
Using Kerberos with NFS for strong security.... 289
Kerberos authentication for CIFS.... 289
Kerberos Realm window.... 290
Kerberos interface services.... 290
Editing Kerberos configuration.... 290
Kerberos Interface window.... 291
DNS/DDNS Services.... 291
Enabling or disabling DDNS.... 292
Editing DNS and DDNS settings.... 292
DNS/DDNS Services window .... 293
Users.... 293
Adding SVM user accounts.... 293
Changing the password for SVM user accounts.... 293
Editing SVM user accounts.... 294
Locking or unlocking SVM user accounts.... 294
Users window .... 294
Roles.... 295
Adding roles.... 295
Editing roles.... 295
Roles window.... 296
UNIX.... 296
UNIX window.... 296
Windows.... 297
Creating a local Windows group.... 297
Editing local Windows group properties.... 299
Adding user accounts to a Windows local group.... 299
Renaming a local Windows group.... 300
Deleting a local Windows group.... 301
Creating a local Windows user account.... 301
Editing the local Windows user properties.... 302
Assigning group memberships to a user account.... 303
Renaming a local Windows user.... 303
Resetting the password of a Windows local user.... 304
Deleting a local Windows user account.... 305
Windows window.... 305
Name mapping.... 308

Name mapping conversion rules....308
Name Mapping window....308

**Managing data protection....310**
Mirror relationships....310
Creating a mirror relationship from a destination SVM....310
Deleting mirror relationships....311
Editing mirror relationships....312
Initializing mirror relationships....313
Updating mirror relationships....313
Quiescing mirror relationships....314
Resuming mirror relationships....314
Breaking SnapMirror relationships....315
Resynchronizing mirror relationships....315
Reverse resynchronizing mirror relationships....316
Aborting a mirror transfer....317
Restoring a volume in a mirror relationship....317
How SnapMirror relationships work....318
Vault relationships....319
Creating a vault relationship from a destination SVM....319
Deleting vault relationships....320
Editing vault relationships....321
Initializing a vault relationship....322
Updating a vault relationship....323
Quiescing a vault relationship....323
Resuming a vault relationship....324
Aborting a Snapshot copy transfer....324
Restoring a volume in a vault relationship....325
What a SnapVault backup is....326
Mirror and vault relationships....327
Creating a mirror and vault relationship from a destination SVM....327
Deleting mirror and vault relationships....328
Editing mirror and vault relationships....329
Initializing mirror and vault relationships....330
Updating mirror and vault relationships....330
Quiescing mirror and vault relationships....331
Resuming mirror and vault relationships....331
Breaking mirror and vault relationships....331
Resynchronizing mirror and vault relationships....332
Reverse resynchronizing mirror and vault relationships....332
Aborting mirror and vault relationships....333
Restoring a volume in a mirror and vault relationship....333
What lag time is....334
Types of data protection relationships....335
SnapMirror licensing....335
Protection window....336
Protection policies....339
Creating protection policies....339
Protection Policies window....340
Snapshot policies....341
Creating Snapshot policies....341
Editing Snapshot policies....341
Deleting Snapshot policies....341

About Snapshot policies.....342
Snapshot Policies window.....342
Schedules.....343
Creating schedules.....343
Editing schedules.....343
Deleting schedules.....344
Schedules.....344
Schedules window.....344

**Copyright, trademark, and machine translation.....345**
Copyright.....345
Trademark.....345
Machine translation.....345

# Understanding System Manager

System Manager is a graphical management interface that enables you to manage storage systems and storage objects (such as disks, volumes, and aggregates) and perform common management tasks related to storage systems from a web browser. As a cluster administrator, you can use System Manager to administer the entire cluster and its resources.

**Important:** System Manager is no longer available as an executable file and is now included with ONTAP software as a web service, enabled by default, and accessible by using a browser.

System Manager enables you to perform many common tasks such as the following:

- Create a cluster, configure a network, and set up support details for the cluster.
- Configure and manage storage objects such as disks, aggregates, volumes, qtrees, and quotas.
- Configure protocols such as CIFS and NFS, and provision file sharing.
- Configure protocols such as FC, FCoE, NVMe, and iSCSI for block access.
- Create and configure network components such as subnets, broadcast domains, data and management interfaces, and interface groups.
- Set up and manage mirroring and vaulting relationships.
- Perform cluster management, storage node management, and storage virtual machine (SVM) management operations.
- Create and configure SVMs, manage storage objects associated with SVMs, and manage SVM services.
- Monitor and manage HA configurations in a cluster.
- Configure Service Processors to remotely log in, manage, monitor, and administer the node, regardless of the state of the node.

# Icons used in the application interface

You can view the icons in the interface to get quick information about systems and operations.

### Dashboard window icons

You might see the following icons when viewing the dashboard for the storage system:

| Icon | Name | Description |
|---|---|---|
| | Warning | There are minor issues, but none that require immediate attention. |
| | Error | Problems that might eventually result in downtime and therefore require attention. |
| | Critical | The storage system is not serving data or cannot be contacted. Immediate attention is required. |
| | Link arrow | If this is displayed next to a line item in a dashboard pane, clicking it links to another page from which you can get more information about the line item or make changes to the line item. |

# System Manager window layouts

Understanding the typical window layouts helps you to navigate and use System Manager effectively. Most System Manager windows are similar to one of two general layouts: object list or details.

### Object list window layout

### Details window layout

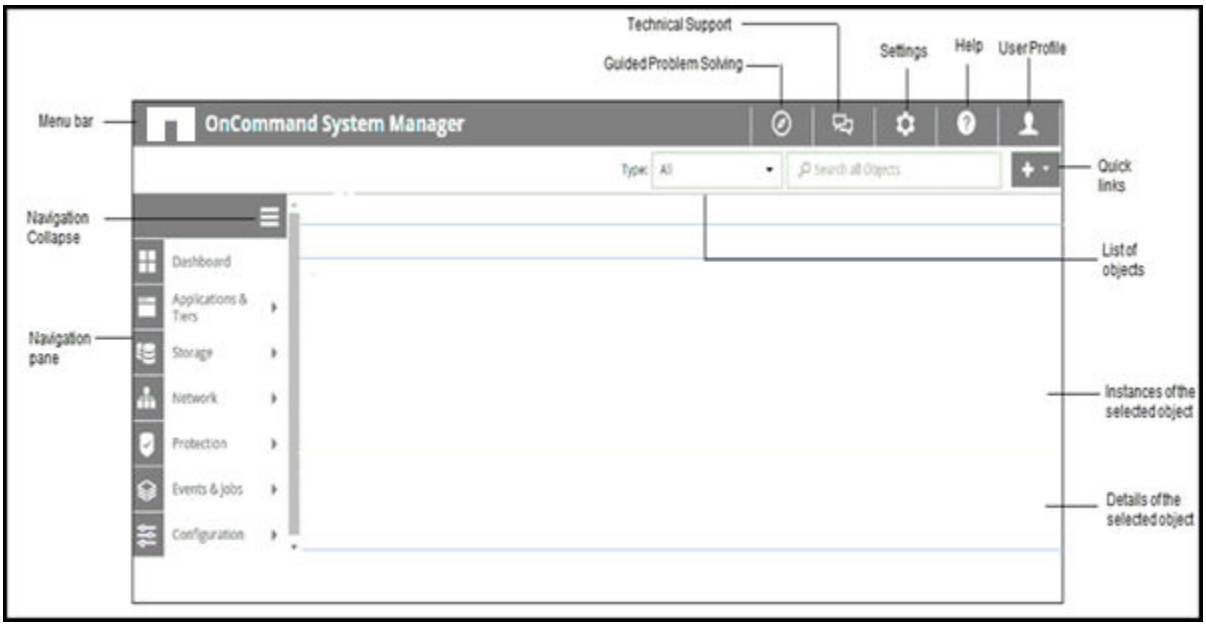

# Window layout customization

System Manager enables you to customize the window layouts. By customizing the windows, you can control which data is viewable and how the data is displayed.

**Sorting**

You can click the column header to change the sort order of the column entries. When you click the column header, the sort arrows ( and ) appear for that column.

**Filtering**

You can click the filter icon ( ) to display only those entries that match the conditions that are provided. You can then use the character filter (?) or string filter (*) to narrow your search. The filter icon is displayed when you move the mouse pointer over the column headings.

You can apply filters to one or more columns.

**Note:** When you apply filters to the physical size field or the usable size field, any value that you enter without the unit suffix in these fields is considered to be in bytes. For example, if you enter a value of 1000 without specifying the unit in the physical size field, the value is automatically considered as 1000 bytes.

**Hiding or redisplaying the columns**

You can click the column display icon ( ) to select which columns you want to display.

**Customizing the layout**

You can drag the bottom of the list of objects area up or down to resize the main areas of the window. You can also display or hide the list of related objects and list of views panels. You can drag the vertical dividers to resize the width of the columns or other areas of the window.

**Searching**

You can use the search box to search for volumes, LUNs, qtrees, network interfaces, storage virtual machines (SVMs), aggregates, disks, or Ethernet ports, or all of these objects. You can click the results to navigate to the exact location of the object.

**Notes:**

- When you search for objects that contain one or more of the { \ ? ^ > | characters, the results are displayed correctly, but they do not navigate to the correct row in the page.
- You must not use the question mark (?) character to search for an object.

# System Manager enhancements

You should become familiar with the features that have been added or changed in this release of System Manager.

### Features and enhancements added in ONTAP 9.3

- Support for SAML authentication for web services

  Beginning with ONTAP 9.3, you can configure multifactor authentication (MFA) for web services by using Security Assertion Markup Language (SAML) authentication. You can use SAML authentication for Service Processor Infrastructure (SPI), ONTAP APIs, and System Manager.
- Application Aware Data Management

  Application aware data management simplifies storage setup and enables you to serve data in minutes for key applications by providing inputs relevant to the application.
- Modified GUI and navigation

  The graphical user interface (GUI) has been revamped to provide users with a more intuitive experience.
- Support for breaking protection relationships between ONTAP and SolidFire systems

  Beginning with ONTAP 9.3, you can use System Manager to break SnapMirror relationships between ONTAP systems and SolidFire storage systems.
- Support for simplified cluster peering and SVM peering

  System Manager offers enhancements that simplify how you configure peer relationships between clusters and between SVMs.
- Support for provisioning an SVM by using a preconfigured template

  Beginning with ONTAP 9.3, you can create and provision an SVM by using a preconfigured template.
- Enhanced cluster dashboard

  The cluster dashboard has been enhanced and made responsive for handheld devices to view different information.

### Features and enhancements added in ONTAP 9.4

- NVMe protocol

  The NVM Express (NVMe) protocol is now supported by ONTAP and can be configured in System Manager. NVMe is an alternative protocol for block access, similar to the existing iSCSi or FC protocols.
- Aggregate recommender

  You can create an aggregate based on storage recommendations. System Manager analyzes the configuration of your storage system and provides storage recommendations such as the number of aggregates that will be created, the available nodes, and the available spare disks.
- FabricPool-enabled aggregates enhancements

  FabricPool-enabled aggregates have been enhanced to support the following features and functionalities:

  - New UI navigation for the external capacity tier menu
  - New "Auto" caching policy
  - Support for inactive (cold) data
  - Support for object store certificate for StorageGRID external capacity tier
  - Support for Microsoft Azure Blob storage external capacity tier
  - More information in the capacity tab of the cluster dashboard
  - Support ONTAP Select

  - Support for viewing external capacity tier, other than StorageGRID, Amazon AWS S3, and Microsoft Azure Blob storage, created using the command-line interface (CLI).
- FlexGroup volumes enhancements

  FlexGroup volumes include the following enhancements and new features:

  - Support for advanced options such as volume encryption, storage efficiency, and QoS.
  - Protect volumes
  - More information in the protection tab of the cluster dashboard
- Support for updating single-node clusters disruptively

  Beginning with ONTAP 9.4, you can update a cluster non-disruptively or you can disruptively update the individual nodes in a high-availability (HA) pair.
- Support for configuring Snapshot copies

  You can configure Snapshot copies by setting a schedule to an existing Snapshot policy. Beginning with ONTAP 9.4, you can have fewer than 1024 Snapshot copies of a FlexVol volume.
- Storage efficiency enhancements

  The percentage of logical space used and the status of logical space reporting is now displayed in the System Manager Volumes window.
- Support for SMB Multichannel

  You can enable SMB protocol to establish multiple channels between a SMB3.0 session and transport connections, specifically for higher performance and fault tolerance and resiliency.

# Setting up your cluster environment

You can create a cluster by using System Manager or the command-line interface (CLI). To create a cluster by using System Manager, you must set up the node management IP address on any node in the cluster network. If you have created a cluster by using the CLI, you can configure the cluster by using System Manager.

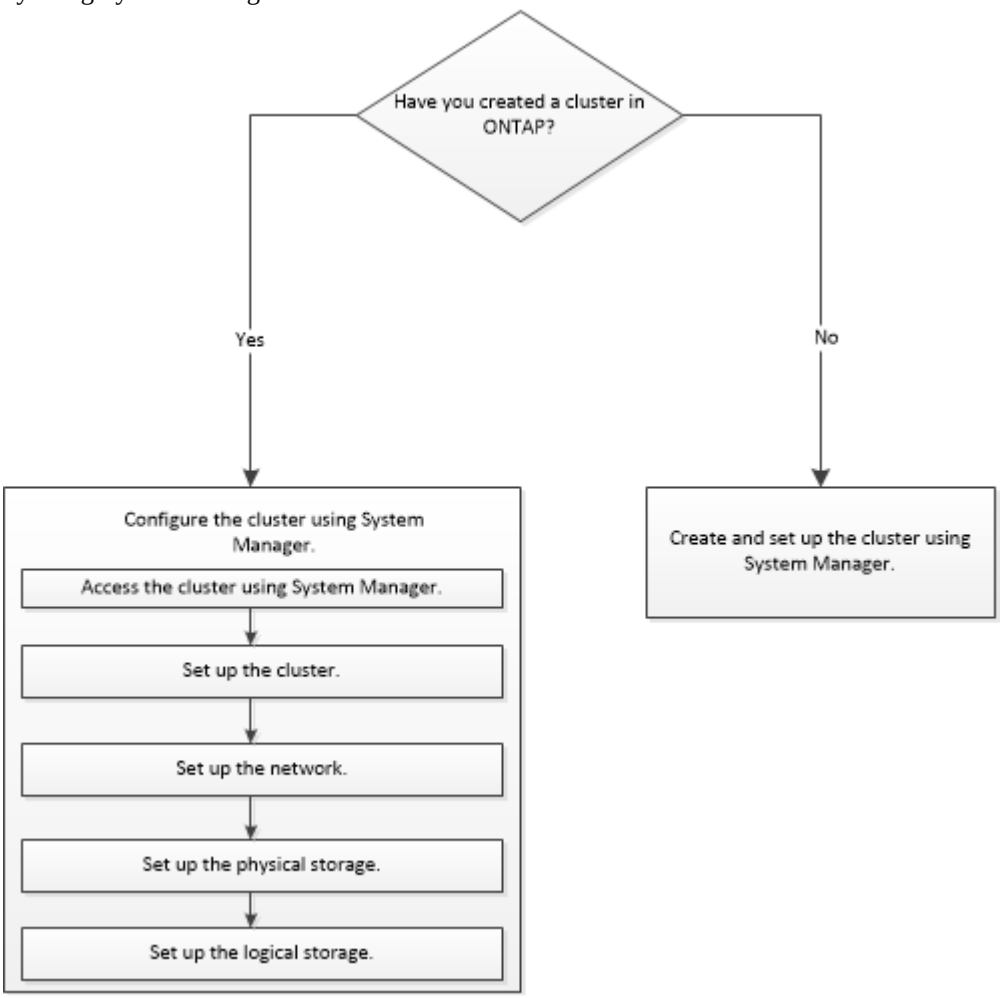


## Setting up the cluster by using OnCommand System Manager

Beginning with ONTAP 9.1, you can use OnCommand System Manager to set up a cluster by creating a cluster, setting up the node management network and cluster management network, and then setting up AutoSupport messages and event notifications.

### Before you begin

- You must have configured the node management IP addresses for at least one node.
- Nodes must be in the default mode of HA.
- Nodes must be running ONTAP 9.1 or later.
- Nodes must be of the same version.
- All of the nodes must be healthy, and cabling for the nodes must be set up.
- Cabling and connectivity must be in place for your cluster configuration.

- You must have sufficient cluster management, node management, Service Processor IP addresses, and gateway and netmask details.
- If the cluster interface is present on a port, then that port must be present in the cluster IPspace.

### About this task

To create a cluster, you have to log in through the console, and configure the node management IP address on any node in the cluster network. After you have configured the node management IP address on a node, you can add other nodes and create a cluster by using OnCommand System Manager.

The cluster setup operation is not supported on MetroCluster configurations for ONTAP software.

You can set up the cluster by using a template file or by manually entering the values in the cluster setup wizard.

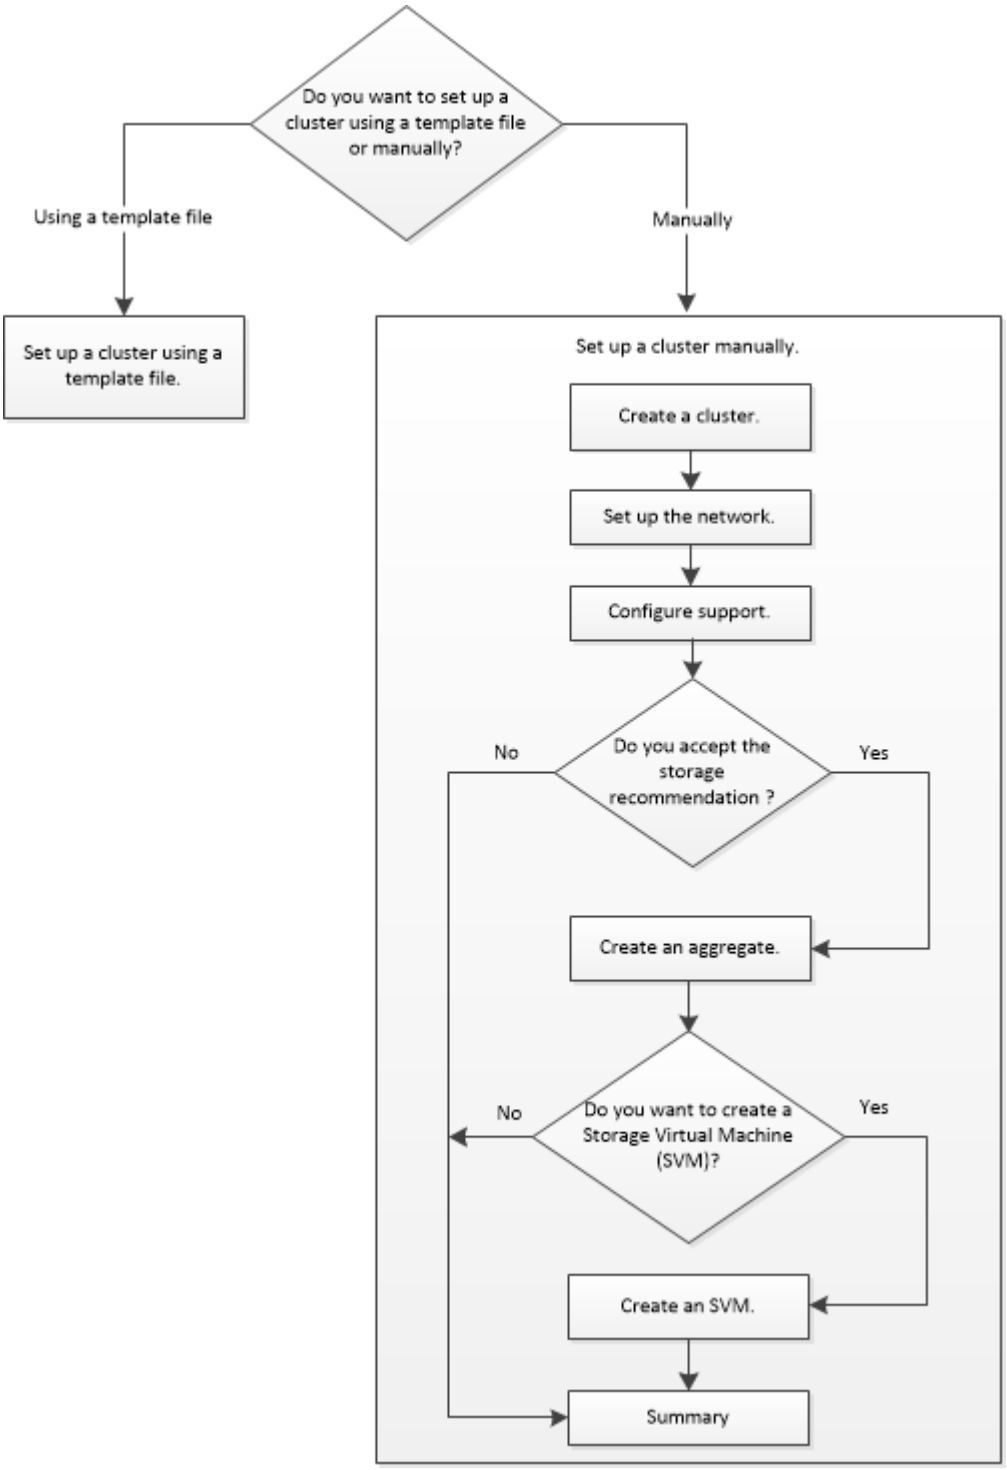


### Choices

- *Setting up a cluster by using the template file* on page 21
- *Setting up the cluster manually* on page 23

## Setting up a cluster by using the template file

You can use the template file that is provided in System Manager to set up a cluster by creating a cluster, setting up the node management and cluster management networks, and then setting up the

AutoSupport messages and event notifications. You can download the template file in `.xlsx` format or `.csv` format.

### About this task

- If the cluster supports ONTAP 9.1 or later, you can add only storage systems that are running ONTAP 9.1 or later.
- All fields are not auto populated when you upload the file.

  You must manually enter the value of some fields such as password and cluster management port.

### Steps

1. Open the web browser, and then enter the node management IP address that you have configured: **`https://node-management-IP`**
   - If you have set up the credentials for the cluster, the Login page is displayed.

     You must enter the credentials to log in.
   - If you have not set up the credentials for the cluster, the Guided Setup window is displayed.
2. Download the `.xlsx` template file or the `.csv` template file.
3. Provide all the required values in the template file, and save the file.

   **Note:**

   - Do not edit any other column in the template other than Value.
   - Do not change the version of the template file.
4. Click **Browse**, and select the updated template file.
   - You can upload the template file only in the `.csv` format. If you have downloaded the template file in `.xlsx` format, you must save the file as a `.csv` file, and then upload the file.
   - You must ensure that the encoding used for this file is UTF8. If not, the values will not be read.
5. Click **Upload**.

   The details that you have provided in the template file are used to complete the cluster setup process.
6. Click the **Guided Setup** icon to view the details for the cluster.
7. Verify the details in the **Cluster** window, and then click **Submit and Continue**.

   You can edit the cluster details, if required.

   If you log in to the Cluster window for the second time, the **Feature Licenses** field is enabled by default. You can add new feature license keys or retain the pre-populated license keys.
8. Verify the details in the **Network** window, and then click **Submit and Continue**.

   You can edit the network details, if required.
9. Verify the details in the **Support** window, and then click **Submit and Continue**.

   You can edit the support details, if required.
10. Verify the details in the **Storage** window, and then create aggregates or exit the cluster setup:

| If you want to... | Then... |
|---|---|
| Exit cluster setup without provisioning storage and creating an SVM | Click **Skip this step**. |

| If you want to... | Then... |
|---|---|
| Provision storage using aggregates and create an SVM | Click **Submit and Continue**. |

You can edit the support details, if required.

11. If you have clicked **Skip this step** on the **Storage** window, view the details on the **Summary** window, and then click **Manage your Cluster** to launch System Manager.
12. If you have clicked **Submit and Continue** on the **Storage** window, verify the details in the SVM window, and then click **Submit and Continue**.

    You can edit the SVM details, if required.
13. Verify all the details in the **Summary** window, and then click **Provision an Application** to provision storage for applications, or click **Manage your Cluster** to complete the cluster setup process and launch System Manager, or click **Export Configuration** to download the configuration file.

### Related information

*System Manager Cluster Guided Setup Templates*

## Setting up the cluster manually

You can use System Manager to manually setup the cluster by creating a cluster, setting up the node management and cluster management networks, and setting up the AutoSupport messages and event notifications.

### Creating a cluster

You can use OnCommand System Manager to create and set up a cluster in your data center.

#### About this task

If the cluster supports ONTAP 9.1 or later, you can add only those storage systems that are running ONTAP 9.1 or later.

#### Steps

1. Open the web browser, and then enter the node management IP address that you have configured: `https://node-management-IP`
   - If you have set up the credentials for the cluster, the Login page is displayed.
     You must enter the credentials to log in.
   - If you have not set up the credentials for the cluster, the Guided Setup window is displayed.
     Click the **Guided Setup** icon to set up a cluster.
2. In the **Cluster** page, enter a name for the cluster.

   **Note:** If all the nodes are not discovered, click **Refresh**.

   The nodes in that cluster network are displayed in the Nodes field.
3. Optional: If desired, update the node names in the **Nodes** field.
4. Enter the password for the cluster.
5. Optional: Enter the feature license keys.
6. Click **Submit**.

#### After you finish

Enter the network details in the Network page to continue with the cluster setup.

### Related reference

*Licenses window* on page 81
Your storage system arrives from the factory with preinstalled software. If you want to add or remove a software license after you receive the storage system, you can use the Licenses window.

*Configuration Updates window* on page 73
You can use the Configuration Updates window to update the configuration details of the cluster, storage virtual machine (SVM), and nodes.

## Setting up a network

By setting up a network, you can manage your cluster, nodes, and Service Processors. You can also set up DNS and NTP details by using the network window.

### Before you begin

You must have set up the cluster.

### About this task

- Only those nodes that are up and running are listed for cluster creation.
  You can create LIFs for those nodes.
- You can disable IP address range and enter individual IP addresses for cluster management, node management, and Service Processor management networks.

## Setting up a network when an IP address range is enabled

You can set up a network by enabling an IP address range. The IP address range enables you to enter IP addresses that are in the same netmask range or in the different netmask range.

### Steps

1. Enter a range of IP addresses in the **IP Address Range** field, and then click **Apply**.

| Option | Description |
|---|---|
| You have a range of IP addresses in the same netmask | Enter the IP address range, and then click **Apply**.<br>IP addresses are applied to cluster management, node management, and Service Processor management networks sequentially. |
| You have a range of IP addresses in different netmasks | Enter the IP address range in rows, and then click **Apply**.<br>The first IP address applied to cluster management and other IP addresses are applied to node management and Service Processor management networks sequentially. |

   **Note:** After entering the IP address range for cluster management, node management, and Service Processor management, you must not manually modify the IP address values in these fields. You must ensure that all the IP addresses are IPv4 addresses.

2. Enter the netmask and gateway details.
3. Select the port for cluster management in the **Port** field.
4. If the **Port** field in the node management is not populated with **e0M**, enter the port details.

   **Note:** By default, the Port field displays e0M.

5. For Service Processor management, if you are overriding the default values, ensure that you have entered the mandatory gateway details.
6. If you have enabled the **DNS Details** field, enter the DNS server details.
7. If you have enabled the **NTP Details** field, enter the NTP server details.

**Note:** Providing alternative NTP server details is optional.

8. Click **Submit**.

### After you finish

Enter AutoSupport message details and event notifications in the Support page to continue with the cluster setup.

#### Related information

*What is a Service Processor and how do I use it?*

*How to configure and troubleshoot NTP on clustered Data ONTAP 8.2 and later using CLI*

*NetApp Documentation: ONTAP 9*

## Setting up a network when an IP address range is disabled

You can set up a network by disabling an IP address range and entering individual IP addresses for cluster management, node management, and service provider networks.

### About this task

In the Networks page, if the **IP Address Range** is disabled, enter individual IP addresses for cluster management, node management, and service processor networks.

### Steps

1. Enter the cluster management IP address in the **Cluster Management IP Address** field.
2. Enter the netmask details for cluster management.
3. Optional: Enter the gateway details for cluster management.
4. Select the port for cluster management in the **Port** field.
5. If you want to provide netmask and gateway details to manage your nodes, clear the **Retain Netmask and Gateway configuration of the Cluster Management** check box, and then enter the netmask and gateway details.
6. Enter the node management IP addresses in the **Node Management** field.
7. If the **Port** field in the node management is not populated with **e0M**, enter the port details.

   **Note:** By default, the Port field displays e0M.

8. Enter the Service Processor management netmask and gateway details.
9. Enter the Service Processor IP management addresses in the **Service Processor Management** field.
10. If you have enabled the **DNS Details** field, enter the DNS server details.
11. If you have enabled the **NTP Details** field, enter the NTP server details.

    **Note:** Providing alternative NTP server details is optional.

12. Click **Submit**.

### After you finish

Enter AutoSupport message details and event notifications in the Support page to continue with the cluster setup.

#### Related reference

*Network window* on page 109

You can use the Network window to view the list of network components, such as subnets, network interfaces, Ethernet ports, broadcast domains, FC/FCoE and NVMe adapters, and IPspaces, and to create, edit, or delete these components in your storage system.

*Configuration Updates window* on page 73

You can use the Configuration Updates window to update the configuration details of the cluster, storage virtual machine (SVM), and nodes.

*Date and Time window* on page 91
The Date and Time window enables you to view the current date and time settings for your storage system and to modify the settings when required.

*Service Processors window* on page 75
You can use the Service Processors window to view and modify Service Processors attributes, such as the IP address, network mask (IPv4) or prefix-length (IPv6), and gateway, and to configure the IP source for a Service Processor.

### Related information

*What is a Service Processor and how do I use it?*

*How to configure and troubleshoot NTP on clustered Data ONTAP 8.2 and later using CLI*

*NetApp Documentation: ONTAP 9*

## Setting up a support page

Setting up the support page completes the cluster setup, and involves setting up the AutoSupport messages and event notifications, and for single-node clusters, configuring system backup.

### Before you begin

You must have set up the cluster and network.

### About this task

If you have enabled the AutoSupport button, all the nodes in that cluster are enabled to send AutoSupport messages. If you have disabled the AutoSupport button, then all the nodes in that cluster are disabled to send AutoSupport messages.

### Steps

1. If you have enabled the AutoSupport button, set up the AutoSupport messages by entering the proxy URL in the **Proxy URL** field.

   **Note:** The format of the proxy URL must be `username:password@proxyUrl:port`.

2. Set up the event notifications by using the mailhost, or SNMP trap host, or Syslog server.

   **Note:** You must set up at least one event notification system.

3. If you have a single-node cluster, configure a system backup on an FTP server or on an HTTP server.

   **Note:** System backup is applicable only for single-node clusters.

4. Click **Submit and continue**.

### After you finish

View the storage recommendations and create SVMs to continue with the cluster setup.

### Related reference

*AutoSupport window* on page 163
The AutoSupport window enables you to view the current AutoSupport settings for your system. You can also change your system's AutoSupport settings.

### Related information

*NetApp AutoSupport*

## Reviewing storage recommendations

Using the Storage window, you can review the storage recommendations that are provided for creating aggregates.

### Before you begin

You must have set up the cluster, network, and the support details.

### About this task

You can create data aggregates per the storage recommendations or you can skip this step and create data aggregates at a later time using System Manager.

### Choices

- To create data aggregates as per the storage recommendations, click **Submit and Continue**.
- To create data aggregates at a later time using System Manager, click **Skip this step**.

### After you finish

If you opted to create aggregates per the storage recommendations, you must create a storage virtual machine (SVM) to continue with the cluster setup.

## Creating an SVM

You can use the Storage Virtual Machine (SVM) window to create fully configured SVMs. The SVMs serve data after storage objects are created on these SVMs.

### Before you begin

- You must have created an aggregate and the aggregate must be online
- You must have ensured that the aggregate has sufficient space for the SVM root volume

### Steps

1. Enter a name for the SVM.
2. Select data protocols for the SVM:

| If you want to... | Then... |
|---|---|
| Enable CIFS protocol by configuring the CIFS server using an Active Directory | **a.** Select the **Active Directory** box.<br>**b.** Enter the Active Directory administrator name.<br>**c.** Enter the Active Directory administrator password.<br>**d.** Enter a name for the CIFS server.<br>**e.** Enter a name for the Active Directory domain.<br>**f.** Depending on your requirements, select the **One data LIF on this SVM** or **One data LIF per node on this SVM** box.<br>**g.** Provide data LIF details such as IP address, netmask, gateway, and port.<br>**h.** Provide DNS details. |
| Enable CIFS protocol by configuring the CIFS server using a workgroup | **a.** Select the **Workgroup** box.<br>**b.** Enter a name for the workgroup.<br>**c.** Enter a name for the CIFS server.<br>**d.** Depending on your requirements, select the **One data LIF on this SVM** or **One data LIF per node on this SVM** check box.<br>**e.** Provide data LIF details such as IP address, netmask, gateway, and port. |

| If you want to... | Then... |
|---|---|
| Enable NFS protocol | **a.** Select the **NFS** box.<br>**b.** Depending on your requirements, select the **One data LIF on this SVM** or **One data LIF per node on this SVM** check box.<br>**c.** Provide data LIF details such as IP address, netmask, gateway, and port. |
| Enable iSCSI protocol | **a.** Select the **iSCSI** box.<br>**b.** Provide data LIF details such as IP address, netmask, gateway, and port. |
| Enable FC/FCoE protocol | **a.** Select the **FC/FCoE** box.<br>**b.** Select the FC/FCoE ports for FC or FCoE protocols.<br>**Note:** Each node must have at least one correctly configured port for each protocol (FC and FCoE). |
| Enable NVMe protocol | **a.** Select the **NVMe** box.<br>**b.** Select the NVMe ports for NVMe protocols.<br>**Note:** At least on NVMe capable adapter must be available in one of the nodes to configure NVMe. |

3. Optional: Click the **Advanced Options** icon and provide details to configure advanced options such as the default language, security style, CIFS server details, and NFS details.
4. Click **Submit and Continue** to create the SVM.

### After you finish

If you have clicked **Submit and Continue**, you must verify the details that you have provided in the Summary window, and then click **Manage your Cluster** to launch System Manager, or click **Provision an Application** to provision storage applications, or click **Export Configuration** to download the configuration file.

## Accessing a cluster by using the OnCommand System Manager browser-based graphic interface

If you prefer to use a graphic interface instead of the command-line interface for accessing and managing a cluster, you can do so by using OnCommand System Manager, which is included with ONTAP as a web service, is enabled by default, and is accessible by using a browser.

### Before you begin

- You must have a cluster user account that is configured with the **`admin`** role and the **`http`**, **`ontapi`**, and **`console`** application types.
- You must have enabled cookies and site data in the browser.

### About this task

You can use a cluster management LIF or node management LIF to access OnCommand System Manager. For uninterrupted access to OnCommand System Manager, you should use a cluster management LIF.

### Steps

1. Point the web browser to the IP address of the cluster management LIF:
   - If you are using IPv4: **`https://`*`cluster-mgmt-LIF`***

- If you are using IPv6: **`https://[cluster-mgmt-LIF]`**

Only HTTPS is supported for browser access of OnCommand System Manager.

If the cluster uses a self-signed digital certificate, the browser might display a warning indicating that the certificate is not trusted. You can either acknowledge the risk to continue the access or install a Certificate Authority (CA) signed digital certificate on the cluster for server authentication.

2. Optional: If you have configured an access banner by using the CLI, then read the message that is displayed in the **Warning** dialog box, and choose the required option to proceed.

   This option is not supported on systems on which Security Assertion Markup Language (SAML) authentication is enabled.

   - If you do not want to continue, click **Cancel**, and close the browser.
   - If you want to continue, click **OK** to navigate to the OnCommand System Manager login page.

3. Log in to OnCommand System Manager by using your cluster administrator credentials.

### Related tasks

*Enabling SAML authentication* on page 58
You can use System Manager to configure Security Assertion Markup Language (SAML) authentication so that remote users can log in by using a secure identity provider (IdP).

*Disabling SAML authentication* on page 59
You can disable Security Assertion Markup Language (SAML) authentication if you want to disable remote access to System Manager, or to edit the SAML configuration.

## Configuring System Manager options

You can enable logging and specify the inactivity timeout value for System Manager.

### About this task

You can configure the options from the System Manager login window. However, you must log in to the application to specify the inactivity timeout value.

### Steps

1. Click ⚙.
2. In the **Setup** pane, click **General**.
3. Specify a log level.
4. Specify the inactivity timeout value in minutes.

   **Note:** This option is not available if you have enabled Security Assertion Markup Language (SAML) authentication.

5. Click **OK**.

## Viewing OnCommand System Manager log files

If you encounter any issues when using System Manager, you can send the log files to technical support to help troubleshoot the issues. The System Manager log files are located in the `mlog` directory along with the ONTAP log files.

### Steps

1. Identify the node that hosts the cluster management LIF.
2. Enter the following URL in a web browser:

**https://*cluster-mgmt-LIF*/spi**

*cluster-mgmt-LIF* is the IP address of the cluster management LIF.

3. Type your cluster administrator credentials, and then click **OK**.
4. In the **Data ONTAP - Root Volume File Access** window, click the **logs** link for the node that hosts the cluster management LIF.
5. Navigate to the `mlog` directory to access the System Manager log files.

   You might require the following log files, depending on the type of issue that you encountered:

   - `sysmgr.log`

     This file contains the latest logs for System Manager.
   - `mgwd.log`
   - `php.log`
   - `apache_access.log`
   - `messages.log`

### How system logging works

System logging is an essential tool for application troubleshooting. You should enable system logging so that if there is a problem with an application, the problem can be located. You can enable System Manager logging at runtime without modifying the application binary.

Log output can be voluminous and therefore can become difficult to manage. System Manager enables you to refine the logging output by selecting which type of log statements are output. By default, system logging is set to INFO. You can choose one of the following log levels:

- OFF
- ERROR
- WARN
- INFO
- DEBUG

These levels function hierarchically. A log level set to OFF indicates no logging of messages.

## Configuring a cluster by using System Manager

Certain prerequisites must be met before you configure a cluster using System Manager.

- You must have created a cluster.
- You must have not configured the cluster.

### Accessing a cluster by using the OnCommand System Manager browser-based graphic interface

If you prefer to use a graphic interface instead of the command-line interface for accessing and managing a cluster, you can do so by using OnCommand System Manager, which is included with ONTAP as a web service, is enabled by default, and is accessible by using a browser.

#### Before you begin

- You must have a cluster user account that is configured with the **`admin`** role and the **`http`**, **`ontapi`**, and **`console`** application types.
- You must have enabled cookies and site data in the browser.

### About this task

You can use a cluster management LIF or node management LIF to access OnCommand System Manager. For uninterrupted access to OnCommand System Manager, you should use a cluster management LIF.

### Steps

1. Point the web browser to the IP address of the cluster management LIF:
   - If you are using IPv4: **`https://cluster-mgmt-LIF`**
   - If you are using IPv6: **`https://[cluster-mgmt-LIF]`**

   Only HTTPS is supported for browser access of OnCommand System Manager.

   If the cluster uses a self-signed digital certificate, the browser might display a warning indicating that the certificate is not trusted. You can either acknowledge the risk to continue the access or install a Certificate Authority (CA) signed digital certificate on the cluster for server authentication.
2. Optional: If you have configured an access banner by using the CLI, then read the message that is displayed in the **Warning** dialog box, and choose the required option to proceed.

   This option is not supported on systems on which Security Assertion Markup Language (SAML) authentication is enabled.
   - If you do not want to continue, click **Cancel**, and close the browser.
   - If you want to continue, click **OK** to navigate to the OnCommand System Manager login page.
3. Log in to OnCommand System Manager by using your cluster administrator credentials.

## Setting up the cluster

Setting up the cluster involves gathering the configuration information, creating cluster-management and node-management interfaces, adding licenses, setting up the cluster time, and monitoring HA pairs.

### Updating the cluster name

You can use System Manager to modify the name of a cluster when required.

#### Steps

1. Click **Configuration** > **Configuration Updates**.
2. In the **Cluster Details** pane, click **Update Cluster Name**.
3. In the **Update Cluster Name** dialog box, specify a new name for the cluster, and then click **Submit**.

### Changing the cluster password

You can use System Manager to reset the password of a cluster.

#### Steps

1. Click **Configuration** > **Configuration Updates**.
2. In the **Cluster Details** pane, click **Change Password**.
3. In the **Change Password** dialog box, specify a new password, confirm the new password, and then click **Change**.

## Editing DNS configurations

You can use System Manager to add host information to centrally manage DNS configurations. You can modify the DNS details when you want to change the domain names or IP addresses.

### Steps

1. Click **Configuration** > **Configuration Updates**.
2. In the **Cluster Details** pane, click **Edit DNS Configuration**.
3. In the **DNS Domains** area, add or modify the DNS domain names.
4. In the **Name Servers** area, add or modify the IP addresses.
5. Click **OK**.

## Creating a cluster management interface

You can use System Manager to create a cluster management interface or LIF to provide a single management interface for a cluster. You can use this LIF to manage all of the activities of the cluster.

### Steps

1. Click **Configuration** > **Configuration Updates**.
2. In the **Cluster Details** pane, click **Create Cluster-management LIF**.
3. In the **Create Cluster-Management LIF** dialog box, specify a name for the cluster management LIF.
4. Assign an IP address to the cluster management LIF:

| If you want to... | Then... |
|---|---|
| Specify the IP address by using a subnet | **a.** Select **Using a subnet**.<br>**b.** In the Add Details dialog box, select the subnet from which the IP address should be assigned.<br>For an intercluster LIF, only the subnets that are associated with the selected IPspace are displayed.<br>**c.** If you want to assign a specific IP address to the LIF, select **Use a specific IP address**, and then type the IP address.<br>The IP address that you specify is added to the subnet if the IP address is not already present in the subnet range.<br>**d.** Click **OK**. |
| Specify the IP address manually without using a subnet | **a.** Select **Without a subnet**.<br>**b.** In the Add Details dialog box, perform the following steps:<br>**i.** Specify the IP address and the network mask or prefix.<br>**ii.** Optional: Specify the gateway.<br>**iii.** If you do not want to use the default value for the Destination field, specify a new destination value.<br>If you do not specify a value, the Destination field is populated with the default value based on the family of the IP address.<br>If a route does not exist, a new route is automatically created based on the gateway and destination.<br>**c.** Click **OK**. |

5. Select the required ports from the **Port details** area.
6. Click **Create**.

## Editing the node name

You can use System Manager to modify the name of a node when required.

### Steps

1. Click **Configuration** > **Configuration Updates**.
2. In the **Nodes** tab, select the node that you want to rename, and then click **Edit Node Name**.
3. In the **Edit Node Name** dialog box, type the new name for the node, and then click **Submit**.

## Creating a node management LIF

You can use System Manager to create a dedicated IP address for managing a particular node in a cluster. You can use this LIF to manage the system maintenance activities of the node.

### Steps

1. Click **Configuration** > **Configuration Updates**.
2. In the **Nodes** tab, select the node for which you want to create a node management LIF, and then click **Create Node-management LIF**.
3. In the **Create Node-Management LIF** dialog box, specify a name for the node management LIF.
4. Assign the IP address to the node management LIF:

| If you want to... | Then... |
|---|---|
| Specify the IP address by using a subnet | **a.** Select **Using a subnet**.<br>**b.** In the Add Details dialog box, select the subnet from which the IP address should be assigned.<br>For an intercluster LIF, only the subnets that are associated with the selected IPspace are displayed.<br>**c.** If you want to assign a specific IP address to the LIF, select **Use a specific IP address**, and then type the IP address.<br>The IP address that you specify is added to the subnet if the IP address is not already present in the subnet range.<br>**d.** Click **OK**. |
| Specify the IP address manually without using a subnet | **a.** Select **Without a subnet**.<br>**b.** In the Add Details dialog box, perform the following steps:<br>**i.** Specify the IP address and the network mask or prefix.<br>**ii.** Optional: Specify the gateway.<br>**iii.** If you do not want to use the default value for the Destination field, specify a new destination value.<br>If you do not specify a value, the Destination field is populated with the default value based on the family of the IP address.<br>If a route does not exist, a new route is automatically created based on the gateway and destination.<br>**c.** Click **OK**. |

5. Select the required ports from the **Ports details** area.
6. Click **Create**.

### After you finish

If you want to delete an existing node management interface or LIF, you must use the command-line interface (CLI).

## Editing AutoSupport settings

You can use System Manager to modify your AutoSupport settings to specify an email address from which email notifications are sent and to add multiple email host names.

### Steps

1. Click ⚙ > **AutoSupport**.
2. Select the node for which you want to modify AutoSupport settings, and then click **Edit**.
3. In the **Email Recipient** tab, type the email address from which email notifications are sent, specify the email recipients and the message content for each email recipient, and then add the mail hosts.

   You can add up to five email addresses for each host.
4. In the **Others** tab, select a transport protocol for delivering the email messages, and then specify the HTTP or HTTPS proxy server details.
5. Click **OK**.

## Adding licenses

If your storage system software was installed at the factory, System Manager automatically adds the software to its list of licenses. If the software was not installed at the factory or if you want to add additional software licenses, you can add the software license by using System Manager.

### Before you begin

The software license code for the specific ONTAP service must be available.

### About this task

- When you add a new license in a MetroCluster configuration, it is a best practice to add the license on the surviving site cluster as well.
- You cannot use System Manager to add the Cloud Volumes ONTAP license.

  The Cloud Volumes ONTAP license is not listed in the license page. System Manager does not raise any alert about the entitlement risk status of the Cloud Volumes ONTAP license.
- You can upload only capacity based licenses.

  The capacity based licenses are of "json" type.

### Steps

1. Click **Configuration** > **Licenses**.
2. Click **Add**.
3. In the **Add License** dialog box, perform the appropriate steps:

| If you want to... | Do this... |
|---|---|
| Add a license for a specific ONTAP service | **a.** Enter the software license key.<br>You can add multiple licenses by entering the software license keys separated by commas.<br>**b.** Click **Add**. |
| Add a capacity based license | **a.** Click Browse, and then select the capacity based license file.<br>**b.** Click **Add**. |

| If you want to... | Do this... |
|---|---|
| Add a license for a specific ONTAP service and add a capacity based license | **a.** Enter the software license key.<br>You can add multiple licenses by entering the software license keys separated by commas.<br>**b.** Click Browse, and then select the capacity based license file.<br>**c.** Click **Add**. |

The new license is added.

The Add License Status dialog box displays the list of licenses that were added successfully. The window also displays the license keys of the licenses that were not added and the reason why the licenses were not added.

4. Click **Close**.

### Result

The software license is added to your storage system and is displayed in the list of licenses in the Licenses window.

### Related reference

*Licenses window* on page 81

Your storage system arrives from the factory with preinstalled software. If you want to add or remove a software license after you receive the storage system, you can use the Licenses window.

## Setting the time zone for a cluster

You can manually set or modify the time zone for a cluster by using the Edit Date and Time dialog box in System Manager. You can also add time servers to the cluster.

### About this task

Network Time Protocol (NTP) is always enabled on a cluster. You can disable NTP by contacting technical support. However, disabling NTP is not recommended.

You can add the IP addresses of the NTP server at your site. This server is used to synchronize the time across the cluster.

You can specify either an IPv4 address or an IPv6 address for the time server.

### Steps

1. Click ⚙.
2. In the **Setup** panel, click **Date and Time**.
3. Click **Edit**.
4. In the **Edit Date and Time** dialog box, select the time zone.
5. Specify the IP address of the time servers, and then click **Add**.
6. Click **OK**.
7. Verify the changes that you made to the time settings in the **Date and Time** window.

### Related tasks

*Creating a Kerberos realm configuration* on page 288

If you want to use Kerberos authentication for client access, you must configure the storage virtual machine (SVM) to use an existing Kerberos realm. You can use System Manager to create a Kerberos realm configuration, which enables SVMs to use Kerberos security services for NFS.

### Related reference

*Date and Time window* on page 91

The Date and Time window enables you to view the current date and time settings for your storage system and to modify the settings when required.

### Monitoring HA pairs

You can use System Manager to monitor the node status and interconnect status of all of the high-availability (HA) pairs in a cluster. You can also verify whether takeover or giveback is enabled or has occurred, and view the reasons why takeover or giveback is not currently possible.

#### Steps

1. Click **Configuration** > **High Availability**.
2. In the **High Availability** window, click the HA pair image to view details such as the cluster HA status, node status, interconnect status, and hardware model of each node.

   If the cluster management LIF or the data LIFs of a node are not in their home node, a warning message is displayed indicating that the node has some LIFs that are not in the home node.

#### Related reference

*High Availability window* on page 78

The High Availability window provides a pictorial representation of the high-availability (HA) state, interconnect status, and takeover or giveback status of all of the HA pairs in ONTAP. You can also manually initiate a takeover operation or giveback operation by using the High Availability window.

## Setting up the network

Setting up the network consists of creating IPspaces, a broadcast domain, and subnets.

### Creating IPspaces

You can create an IPspace by using System Manager to configure a single ONTAP cluster for client access from more than one administratively separate network domain, even when the clients use the same IP address subnet range. This enables you to separate client traffic for privacy and security.

#### About this task

All of the IPspace names must be unique within a cluster and must not consist of names that are reserved by the system, such as "local" or "localhost."

#### Steps

1. Click the **Network** tab.
2. In the **IPspaces** tab, click **Create**.
3. In the **Create IPspaces** dialog box, specify a name for the IPspace that you want to create.
4. Click **Create**.

### Creating broadcast domains

You can create a broadcast domain by using System Manager to provide a logical division of a computer network. In a broadcast domain, all associated nodes can be reached through broadcast at the datalink layer.

#### Steps

1. Click the **Network** tab.
2. In the **Broadcast Domains** tab, click **Create**.
3. In the **Create Broadcast Domain** dialog box, specify the name, MTU size, IPspace, and ports for the broadcast domain that you want to create.
4. Click **Create**.

#### Related reference

*Network window* on page 109

You can use the Network window to view the list of network components, such as subnets, network interfaces, Ethernet ports, broadcast domains, FC/FCoE and NVMe adapters, and IPspaces, and to create, edit, or delete these components in your storage system.

### Creating subnets

You can create a subnet by using System Manager to provide a logical subdivision of an IP network to pre-allocate the IP addresses. A subnet enables you to create interfaces more easily by specifying a subnet instead of an IP address and network mask values for each new interface.

#### Before you begin

You must have created the broadcast domain on which the subnet is used.

#### About this task

If you specify a gateway when creating a subnet, a default route to the gateway is added automatically to the SVM when a LIF is created using that subnet.

#### Steps

1. Click the **Network** tab.
2. In the **Subnets** tab, click **Create**.
3. In the **Create Subnet** dialog box, specify subnet details, such as the name, subnet IP address or subnet mask, range of IP addresses, gateway address, and broadcast domain.

   You can specify the IP addresses as a range, as comma-separated multiple addresses, or as a mix of both.
4. Click **Create**.

#### Related reference

*Network window* on page 109

You can use the Network window to view the list of network components, such as subnets, network interfaces, Ethernet ports, broadcast domains, FC/FCoE and NVMe adapters, and IPspaces, and to create, edit, or delete these components in your storage system.

## Setting up physical storage

Setting up the physical storage consists of assigning disks to nodes, zeroing the spare disks, and creating aggregates.

### Assigning disks to nodes

You can use System Manager to assign ownership of an unassigned disk to a specific node to increase the capacity of an aggregate or storage pool.

#### About this task

- You can assign disks if the following conditions are true:
  - The container type of the selected disks must be "unassigned".
  - The disks must be connected to nodes in an HA pair.
  - The disks must be visible to the node.
- For MetroCluster configurations, you cannot use System Manager to assign disks.
  You must use the command-line interface instead.

#### Steps

1. Click **Storage** > **Aggregates & Disks** > **Disks**.
2. In the **Disks** window, select the **Inventory** tab.

3. Select the disks that you want to assign, and then click **Assign**.
4. In the **Assign Disks** dialog box, select the node to which you want to assign the disks.
5. Click **Assign**.

## Zeroing spare disks

You can use System Manager to erase all the data and to format the spare disks by writing zeros to the disk. These disks can then be used in new aggregates.

### About this task

When you zero the spare disks, all the spares in the cluster, including array LUNs, are zeroed. You can zero the spare disks for a specific node or for the entire cluster.

### Steps

1. Click **Storage** > **Aggregates & Disks** > **Disks**.
2. In the **Disks** window, select the **Inventory** tab.
3. Click **Zero Spares**.
4. In the **Zero Spares** dialog box, select a node or "All nodes" from which you want to zero the disks.
5. Select the **Zero all non-zeroed spares** check box to confirm the zeroing operation.
6. Click **Zero Spares**.

### Related concepts

*Storage recommendations for creating aggregates* on page 128

Starting with System Manager 9.4, you can create aggregates based on storage recommendations. However, you must determine whether create aggregates based on storage recommendations is supported in your environment. If it is not, you must decide the RAID policy and disk configuration, and then create the aggregates manually.

## Provisioning storage through aggregates

You can create an aggregate based on storage recommendations or manually depending on your requirement. You can create Flash Pool aggregates, SnapLock aggregates, and a FabricPool-enabled aggregates to provide storage for one or more volumes by using System Manager.

### Before you begin

You must have enough spare disks to create an aggregate.

### About this task

You cannot perform the following actions by using System Manager:

- Combine disks of different sizes even if there are enough spare disks of different sizes.
  You can initially create an aggregate with disks of the same size and then add disks of a different size later.
- Combine disks with different checksum types.
  You can initially create an aggregate with a single checksum type and add storage of a different checksum type later.

### Related reference

*Aggregates window* on page 136

You can use the Aggregates window to create, display, and manage information about aggregates.

*Storage Tiers window* on page 129

You can use the Storage Tiers window to view cluster-wide space details and to add and view aggregate details.

### Choices


- *Provisioning storage by creating an aggregate based on storage recommendations* on page 39
- *Provisioning storage by creating an aggregate manually* on page 39
- *Provisioning storage by creating a Flash Pool aggregate manually* on page 40
- *Provisioning storage by creating a SnapLock aggregate manually* on page 41
- *Provisioning storage by creating a FabricPool-enabled aggregate manually* on page 42


## Provisioning storage by creating an aggregate based on storage recommendations

You can use System Manager to create an aggregate based on storage recommendations. System Manager analyzes the configuration of your storage system and provides storage recommendations such as the number of aggregates that will be created, the available nodes, and the available spare disks.

### About this task

- You cannot create an aggregate based on storage recommendations in Cloud Volumes ONTAP, ONTAP Select, and MetroCluster configurations.
- Errors, if any, are displayed on the screen.
  You can fix these errors and then create an aggregate based on the storage recommendations, or you can create an aggregate manually.

### Steps

1. Create an aggregate by using one of the following methods:
   - Click **Applications & Tiers** > **Storage Tiers** > **Add Aggregate**.
   - Click **Storage** > **Aggregate & Disks** > **Aggregates** > **Create**.
2. Review the storage recommendations, and then click **Submit**.
   The Information dialog box displays the status of the aggregates.
3. Optional: Click **Run in Background** to navigate to the **Aggregates** window.
4. Click **Refresh** to view the aggregates that are created.

## Provisioning storage by creating an aggregate manually

You can manually create an aggregate that consists of only HDDs or only SSDs by using System Manager.

### Before you begin

All of the disks must be of the same size.

### About this task

- If you are creating an aggregate on a four-node cluster in ONTAP Select, the mirrored aggregate option is selected by default.
- Starting with ONTAP 9.0, you can create aggregates with disk size equal to or larger than 10 TB.
- If the disk type of the aggregate disks is FSAS or MSATA, and the disk size is equal to or larger than 10 TB, then RAID-TEC is the only option available for RAID type.

### Steps

1. Create an aggregate by using one of the following methods:
   - Click **Applications & Tiers** > **Storage Tiers** > **Add Aggregate**.
   - Click **Storage** > **Aggregate & Disks** > **Aggregates** > **Create**.
2. Enable the **Manually Create Aggregate** option to create an aggregate.

3. To create an aggregate:
   a. Specify the name of the aggregate, the disk type, and the number of disks or partitions to include in the aggregate.

      The minimum hot spare rule is applied to the disk group that has the largest disk size.
   b. Optional: Modify the RAID configuration of the aggregate:

      **i.** Click **Change**.

      **ii.** In the Change RAID Configuration dialog box, specify the RAID type and the RAID group size.

      RAID-DP is the only supported RAID type for shared disks.

      **iii.** Click **Save**.
   c. If you want to mirror the aggregate, select the **Mirror this aggregate** check box.

      For MetroCluster configurations, creating unmirrored aggregates is restricted. Therefore, the mirroring option is enabled by default for MetroCluster configurations.
4. Click **Create**.

### Result

The aggregate is created with the specified configuration, and is added to the list of aggregates in the Aggregates window.

## Provisioning storage by creating a Flash Pool aggregate manually

You can use System Manager to create a Flash Pool aggregate manually, or to convert an existing HDD aggregate to a Flash Pool aggregate by adding SSDs. When you create a new HDD aggregate, you can provision an SSD cache to it and create a Flash Pool aggregate.

### Before you begin

- You must be aware of the platform-specific best practices and workload-specific best practices for the Flash Pool aggregate SSD tier size and configuration.
- All of the HDDs must be in the zeroed state.
- If you want to add SSDs to the aggregate, all of the existing SSDs and dedicated SSDs must be of the same size.

### About this task

- You cannot use partitioned SSDs while creating a Flash Pool aggregate.
- You cannot mirror the aggregates if the cache source is storage pools.
- Starting with ONTAP 9.0, you can create aggregates with disk size equal to or larger than 10 TB.
- If the disk type of the aggregate disks is FSAS or MSATA, and the disk size is equal to or larger than 10 TB, then RAID-TEC is the only option available for RAID type.

### Steps

1. Create an aggregate by using one of the following methods:
   - Click **Applications & Tiers** > **Storage Tiers** > **Add Aggregate**.
   - Click **Storage** > **Aggregate & Disks** > **Aggregates** > **Create**.
2. Enable the **Manually Create Aggregate** option to create an aggregate.
3. In the **Create Aggregate** window, specify the name of the aggregate, the disk type, and the number of disks or partitions to include for the HDDs in the aggregate.
4. If you want to mirror the aggregate, select the **Mirror this aggregate** check box.

   For MetroCluster configurations, creating unmirrored aggregates is restricted. Therefore, the mirroring option is enabled by default for MetroCluster configurations.
5. Click **Use Flash Pool Cache with this aggregate**.

6. Specify the cache source:

| If you want to select the cache source as... | Then... |
|---|---|
| Storage pools | **a.** Select **Storage pools** as the Cache Source.<br>**b.** Select the storage pool from which the cache can be obtained, and then specify the cache size.<br>**c.** Modify the RAID type, if required. |
| Dedicated SSDs | **a.** Select **Dedicated SSDs** as the Cache Source.<br>**b.** Select the SSD size and the number of SSDs to include in the aggregate.<br>**c.** Modify the RAID configuration, if required:<br>**i.** Click **Change**.<br>**ii.** In the Change RAID Configuration dialog box, specify the RAID type and the RAID group size.<br>**iii.** Click **Save**. |

7. Click **Create**.

### Result

The Flash Pool aggregate is created with the specified configuration, and is added to the list of aggregates in the Aggregates window.

#### Related concepts

*How storage pool works* on page 143

A *storage pool* is a collection of SSDs. You can combine SSDs to create a storage pool, which enables you to share the SSDs and SSD spares across multiple Flash Pool aggregates, at the same time.

#### Related information

*NetApp Technical Report 4070: Flash Pool Design and Implementation Guide*

## Provisioning storage by creating a SnapLock aggregate manually

You can use System Manager to create a SnapLock Compliance aggregate or a SnapLock Enterprise aggregate manually. You can create SnapLock volumes on these aggregates, which provide "write once, read many" (WORM) capabilities.

### Before you begin

The SnapLock license must have been added.

### About this task

- In MetroCluster configurations, you can create only SnapLock Enterprise aggregates.
- For array LUNs, only SnapLock Enterprise aggregates are supported.
- Starting with ONTAP 9.0, you can create aggregates with disk size equal to or larger than 10 TB.
- If the disk type of the aggregate disks is FSAS or MSATA, and the disk size is equal to or larger than 10 TB, then RAID-TEC is the only option available for RAID type.
- Starting with ONTAP 9.1, you can create a SnapLock aggregate on an AFF platform.

### Steps

1. Create a SnapLock aggregate by using one of the following methods:
   - Click **Applications & Tiers** > **Storage Tiers** > **Add Aggregate**.

- Click **Storage** > **Aggregate & Disks** > **Aggregates** > **Create**.

2. Enable the **Manually Create Aggregate** option to create an aggregate.
3. To create a SnapLock aggregate:
   a. Specify the name of the aggregate, the disk type, and the number of disks or partitions to include in the aggregate.

      You cannot change the name of a SnapLock Compliance aggregate after you create the aggregate.

      The minimum hot spare rule is applied to the disk group that has the largest disk size.
   b. Optional: Modify the RAID configuration of the aggregate:

      **i.** Click **Change**.

      **ii.** In the Change RAID Configuration dialog box, specify the RAID type and the RAID group size.

      Shared disks support two RAID types: RAID-DP and RAID-TEC.

      **iii.** Click **Save**.
   c. Specify the SnapLock type.
   d. If you have not initialized the system ComplianceClock, select the **Initialize ComplianceClock** check box.

      This option is not displayed if the ComplianceClock is already initialized on the node.

      **Note:** You must ensure that the current system time is correct. The ComplianceClock is set based on the system clock. Once the ComplianceClock is set, you cannot modify or stop the ComplianceClock.
   e. Optional: If you want to mirror the aggregate, select the **Mirror this aggregate** check box.

      For MetroCluster configurations, creating unmirrored aggregates is restricted. Therefore, the mirroring option is enabled by default for MetroCluster configurations.

      By default, the mirroring option is disabled for SnapLock Compliance aggregates.
4. Click **Create**.

## Provisioning storage by creating a FabricPool-enabled aggregate manually

You can use System Manager to create a FabricPool-enabled aggregate manually or to convert an existing SSD aggregate to a FabricPool-enabled aggregate by attaching an external capacity tier to the SSD aggregate.

### Before you begin

- You must have created an external capacity tier and attached the external capacity tier to the cluster in which the SSD aggregate resides.
- An on-premises external capacity tier must have been created.
- A dedicated network connection must exist between the external capacity tier and the aggregate.

### About this task

- The supported external capacity tiers are StorageGRID, Amazon AWS S3, and Microsoft Azure Blob storage.

  **Note:**

  - Azure Stack, which is an on-premises Azure services, is not supported.
  - If you want to use Amazon AWS S3 or Microsoft Azure Blob storage as an external capacity tier, you must have the FabricPool capacity license.

- FabricPool-enabled aggregates are not supported on MetroCluster configurations.

### Steps

1. Create a FabricPool-enabled aggregate by using one of the following methods:
   - Click **Applications & Tiers** > **Storage Tiers** > **Add Aggregate**.
   - Click **Storage** > **Aggregate & Disks** > **Aggregates** > **Create**.
2. Enable the **Manually Create Aggregate** option to create an aggregate.
3. To create a FabricPool-enabled aggregate:
   a. Specify the name of the aggregate, the disk type, and the number of disks or partitions to include in the aggregate.

      **Note:** Only all flash (all SSD) aggregates support FabricPool-enabled aggregates.

      The minimum hot spare rule is applied to the disk group that has the largest disk size.
   b. Optional: Modify the RAID configuration of the aggregate:
      **i.** Click **Change**.
      **ii.** In the Change RAID Configuration dialog box, specify the RAID type and the RAID group size.
      RAID-DP is the only supported RAID type for shared disks.
      **iii.** Click **Save**.
4. Select the **FabricPool** checkbox, and then select an external capacity tier from the list.
5. Click **Create**.

## Setting up logical storage

Setting up the logical storage consists of creating storage virtual machines (SVMs) and volumes.

### Creating SVMs

You can use System Manager to create fully configured storage virtual machines (SVMs) that can serve data immediately. A cluster can have one or more SVMs.

#### Before you begin

- The cluster must have at least one non-root aggregate in the online state.
- The aggregate must have sufficient space for the SVM root volume.
- You must have synchronized the time across the cluster by configuring and enabling NTP to prevent CIFS creation and authentication failures.
- The protocols that you want to configure on the SVM must be licensed.
- You must have configured the CIFS protocol for secure DDNS to work.

#### About this task

- While creating SVMs, you can perform the following tasks:
  - Create and fully configure SVMs.
  - Configure the volume type that is allowed on SVMs.
  - Create and configure SVMs with minimal network configuration.
  - Delegate the administration to SVM administrators.
- To name the SVM, you can use alphanumeric characters and the following special characters: "." (period), "-" (hyphen), and "_" (underscore).
  The SVM name should start with an alphabet or "_" (underscore) and must not contain more than 47 characters.

  **Note:** You should use unique fully qualified domain names (FQDNs) for the SVM name such as vs0.example.com.

- You can establish SnapMirror relationships only between volumes that have the same language settings.

  The language of the SVM determines the character set that is used to display file names and data for all NAS volumes in the SVM.
- You cannot use a SnapLock aggregate as the root aggregate of SVMs.

### Steps

1. Click **Storage** > **SVMs**.
2. Click **Create**.
3. In the **Storage Virtual Machine (SVM) Setup** window, specify the following details:
   - SVM name
   - IPspace allocated to the SVM
   - Volume type allowed
   - Protocols allowed
   - SVM language
   - Security style of the root volume
   - Root aggregate

   The default language setting for any SVM is C.UTF-8.

   By default, the aggregate with the maximum free space is selected as the container for the root volume of the SVM. Based on the protocols selected, the default security style and the root aggregate are selected.

   The security style is set to NTFS if you select CIFS protocol or a combination of CIFS protocol with the other protocols. The security style is set to UNIX if you select NFS, iSCSI, NVMe, or FC/FCoE, or a combination of these protocols.

   **Note:** NVMe does not allow the combination of protocols.

   In a MetroCluster configuration, only the aggregates that are contained in the cluster are displayed.
4. Specify the DNS domain names and the name server IP addresses to configure the DNS services.

   The default values are selected from the existing SVM configurations.
5. Optional: When configuring a data LIF to access data using a protocol, specify the target alias, subnets, and the number of LIFs per node.

   You can select the **Review or Modify LIFs configuration (Advanced Settings)** checkbox to modify the number of portsets in the LIF.

   You can edit the details of the portset in a particular node by selecting the node from the nodes list in the details area.
6. Optional: Enable host-side applications such as SnapDrive and SnapManager for the SVM administrator by providing the SVM credentials.
7. Optional: Create a new LIF for SVM management by clicking **Create a new LIF for SVM management**, and then specify the portsets and the IP address with or without a subnet for the new management LIF.

   For CIFS and NFS protocols, data LIFs have management access by default. You must create a new management LIF only if required. For iSCSI, FC and NVMe protocols, a dedicated SVM management LIF is required because data protocols and management protocols cannot share the same LIF.
8. Click **Submit & Continue**.
   The SVM is created with the specified configuration.

### Result

The SVM that you created is started automatically. The root volume name is automatically generated as *`SVM name_root`*. By default, the `vsadmin` user account is created and is in the locked state.

### After you finish

You must configure at least one protocol on the SVM to allow data access.

## Configuring CIFS and NFS protocols on SVMs

You can use System Manager to configure CIFS and NFS protocols on a storage virtual machine (SVM) to provide file-level data access for NAS clients. To enable the CIFS protocol, you must create data LIFs and the CIFS server. To enable the NFS protocol, you can specify the NIS details and the data LIFs.

### Before you begin

- The protocols that you want to configure or enable on the SVM must be licensed.
  If the protocol that you want to configure is not enabled on the SVM, you can use the Edit Storage Virtual Machine window to enable the protocol for the SVM.
- You must have the Active Directory, organizational unit, and administrative account credentials for configuring the CIFS protocol.

### About this task

SnapLock aggregates are not considered for automatically creating volumes.

### Steps

1. If you have not configured the protocols while creating the SVM, click **Storage** > **SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **Protocols** pane, click the protocol that you want to configure.
4. In the **Data LIF Configuration** section, if you want to retain the same data LIF configuration for both CIFS and NFS, select the **Retain the CIFS data LIF's configuration for NFS client** check box.

   If you do not retain the same data LIF configuration for both CIFS and NFS, you must specify the IP address and ports separately for CIFS and NFS.
5. Specify the IP address by choosing one of the following options:

| If you want to... | Then... |
|---|---|
| Specify the IP address by using a subnet | **a.** Select **Using a subnet**.<br>**b.** In the Add Details dialog box, select the subnet from which the IP address must be assigned.<br>For intercluster LIFs, only the subnets that are associated with the selected IPspace are displayed.<br>**c.** If you want to assign a specific IP address to the interface, select **Use a specific IP address**, and then type the IP address.<br>The IP address that you specify is added to the subnet if the IP address is not already present in the subnet range.<br>**d.** Click **OK**. |

| If you want to... | Then... |
|---|---|
| Specify the IP address manually without using a subnet | **a.** Select **Without a subnet**.<br>**b.** In the Add Details dialog box, perform the following steps:<br>**i.** Specify the IP address and the network mask or prefix.<br>**ii.** Optional: Specify the gateway.<br>**iii.** If you do not want to use the default value for the Destination field, specify a new destination value.<br>If you do not specify a destination value, the Destination field is populated with the default value based on the family of the IP address.<br>If a route does not exist, a new route is automatically created based on the gateway and destination.<br>**c.** Click **OK**. |

6. Specify a port to create a data LIF:
   a. Click **Browse**.
   b. In the **Select Network Port or Adapter** dialog box, select a port.
   c. Click **OK**.
7. Configure the CIFS server by performing the following steps:
   a. Specify the following information to create a CIFS server:
      - CIFS server name
      - Active Directory to associate with the CIFS server
      - Organizational unit (OU) within the Active Directory domain to associate with the CIFS server
        By default, this parameter is set to CN=Computers.
      - Credentials of an administrative account that has sufficient privileges to add the CIFS server to the OU
   b. Optional: Select **Encrypt Data while accessing all shares of this SVM** to enable SMB 3.0 encryption for all of the shares of the SVM.
   c. Provision a volume for CIFS storage when configuring the protocol by specifying the share name, size of the share, and access permissions.
   d. Optional: Select **Encrypt Data while accessing this share** to enable SMB 3.0 encryption for a particular share.
8. Optional: Configure NIS services:
   a. Specify the IP addresses of the NIS servers and NIS domain name to configure NIS services on the SVM.
   b. Select the appropriate database type for which you want to add the "nis" name service source.
   c. Provision a volume for NFS storage by specifying the export name, size, and permission.
9. Click **Submit & Continue**.

### Result

The CIFS server and NIS domain are configured with the specified configuration, and the data LIFs are created. By default, the data LIFs have management access. You can view the configuration details on the Summary page.

## Configuring iSCSI protocol on SVMs

You can configure the iSCSI protocol on a storage virtual machine (SVM) to provide block-level data access by using System Manager. You can create iSCSI LIFs and portsets and then add the

LIFs to the portsets. LIFs are created on the most suitable adapters and are assigned to portsets to ensure data path redundancy.

### Before you begin

- The iSCSI license must be enabled on the cluster.
  If the iSCSI protocol is not enabled on the SVM, you can use the Edit Storage Virtual Machine window to enable the protocol for the SVM.
- All of the nodes in the cluster must be healthy.
- Each node must have at least two data ports, and the port state must be **up**.

### About this task

- You can configure the iSCSI protocol while creating the SVM or you can do so at a later time.
- SnapLock aggregates are not considered for automatically creating volumes.

### Steps

1. If you have not configured the iSCSI protocol while creating the SVM, click **Storage** > **SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **Protocols** pane, click **iSCSI**.
4. Optional: In the **Network Access** section, specify an alias for the iSCSI target.

   The maximum number of characters for an alias name is 128. If you do not specify a target alias, the SVM name is used as an alias.
5. Specify the number of iSCSI LIFs that can be assigned to a single node.

   The minimum number of LIFs per node is one. The maximum number is the minimum of all the ports in the **up** state across the nodes. If the maximum value is an odd number, the previous even number is considered as the maximum value. You can choose any even number in the minimum and maximum value range.

   A 4-node cluster has node1, node2, and node3 with six ports each in the **up** state, and node4 with seven ports in the **up** state. The effective maximum value for the cluster is 6.

   If the number of LIFs that you want to assign to the node is more than two, you must assign at least one portset to each LIF.
6. Specify the network details, including the subnet details, to create iSCSI LIFs:

| If you want to... | Then... |
|---|---|
| Specify the IP address by using a subnet | **a.** Select **Using a subnet**.<br>**b.** In the Add Details dialog box, select the subnet from which the IP address must be assigned.<br>For intercluster LIFs, only the subnets that are associated with the selected IPspace are displayed.<br>**c.** If you want to assign a specific IP address to the interface, select **Use a specific IP address**, and then type the IP address.<br>The IP address that you specify is added to the subnet if the IP address is not already present in the subnet range.<br>**d.** Click **OK**. |

| If you want to... | Then... |
|---|---|
| Specify the IP address manually without using a subnet | **a.** Select **Without a subnet**.<br>**b.** In the Add Details dialog box, perform the following steps:<br>**i.** Specify the IP address and the network mask or prefix.<br>**ii.** Optional: Specify the gateway.<br>**iii.** If you do not want to use the default value for the Destination field, specify a new destination value.<br>If you do not specify a destination value, the Destination field is populated with the default value based on the family of the IP address.<br>If a route does not exist, a new route is automatically created based on the gateway and destination.<br>**c.** Click **OK**. |

7. Select the broadcast domain.
8. Select the adapter type.

   If you have NIC cards configured in your cluster, you should select **NIC**.

   If you have CNS cards configured in your cluster, you should select **CNA**.

   If you have ifgrps configured in your cluster, you should select **Interface Group**.

   **Note:** The ifgrp port must be added in the broadcast domain.
9. Optional: Provision a LUN for iSCSI storage when configuring the iSCSI protocol by specifying the LUN size, OS type for the LUN, and host initiator details.
10. If you want to verify or modify the configuration of the automatically generated iSCSI LIFs, select **Review or Modify LIFs configuration (Advanced Settings)**.

    You can modify only the LIF name and the home port. By default, the portsets are set to the minimum value. You must specify unique entries. If you specify duplicate LIF names, System Manager appends numeric values to the duplicate LIF name.

    Based on the selected portset, the LIFs are distributed across the portsets by using a round-robin method to ensure redundancy in case of node failure or port failure.
11. Click **Submit & Continue**.

### Result

The iSCSI LIFs and portsets are created with the specified configuration. The LIFs are distributed among the portsets based on the selected portset. The iSCSI service is started if all of the LIFs are successfully created.

If LIF creation fails, you can create the LIFs by using the Network Interfaces window, attach the LIFs to the portsets by using the LUNs window, and then start the iSCSI service by using the iSCSI window.

## Configuring FC protocol and FCoE protocol on SVMs

You can configure the FC protocol and the FCoE protocol on the storage virtual machine (SVM) for SAN hosts. LIFs are created on the most suitable adapters and are assigned to port sets to ensure data path redundancy. Based on your requirements, you can configure either the FC protocol or the FCoE protocols, or both the protocols by using System Manager.

### Before you begin

- The FCP license must be enabled on the cluster.

- All of the nodes in the cluster must be healthy.
- Each node must have at least two correctly configured ports for each protocol (FC and FCoE).

### About this task

- You can configure the FC protocol and the FCoE protocol while creating the SVM or you can configure the protocols at a later time.

  If the protocols are not allowed on the SVM, you can use the Edit Storage Virtual Machine window to enable the protocols for the SVM.
- SnapLock aggregates are not considered for automatically creating volumes.

### Steps

1. If you have not configured the protocols while creating the SVM, click the **Storage** > **SVMs** tab.
2. Select the SVM, and then click **SVM Settings**.
3. In the **Protocols** pane, click **FC/FCoE**.
4. In the **Data Interface Configuration** section, select the corresponding option to configure data LIFs for the FC protocol and the FCoE protocol.
5. Specify the number of data LIFs per node for each protocol.

   The minimum number of LIFs per node is one. The maximum number is the minimum of all the ports in the `up` state across the nodes. If the maximum value is an odd number, the previous even number is considered as the maximum value. You can choose any even number in the minimum and maximum value range.

   A four-node cluster has node1, node2, and node3 with six ports each in the `up` state, and node4 with seven ports in the `up` state. The effective maximum value for the cluster is six.

   If the number of LIFs that you want to assign to the node is more than two, you must assign at least one port set to each LIF.
6. If you want to verify or modify the automatically generated LIFs configuration, select **Review or Edit the Interface Association**.

   You can modify only the LIF name and home port. You must ensure that you do not specify duplicate entries.
7. Optional: Provision a LUN for the FC storage or FCoE storage when configuring the protocol by providing the LUN size, OS type for the LUN, and host initiator details.
8. Click **Submit & Continue**.

### Result

The data LIFs and port sets are created with the specified configuration. The LIFs are distributed accordingly among the port sets. The FCP service is started if all of the LIFs are successfully created for at least one protocol.

If LIF creation fails, you can create the LIFs and start the FCP service from the FC/FCoE window.

#### Related information

*NetApp Documentation: ONTAP 9*

## Configuring NVMe protocol on SVMs

You can configure the NVMe protocol on a storage virtual machine (SVM) using System Manager. You can then create namespaces and assign them to an NVMe subsystem and host.

### About this task

The SVM with NVMe should not have any other protocol. If you select NVMe, then the rest of the protocols will be disabled. You can also configure NVMe while creating the SVM.

### Steps

1. If you did not configure the NVMe protocol when creating the SVM, click **Storage** > **SVMs**
2. Select the SVM, and then click **SVM settings**.
3. In the **Protocols** pane, click **NVMe**.
4. Click the link to configure the protocol, as required.

   **Note:** If there are any other protocols enabled, you must deselect these to make NVMe available to select. NVMe cannot be combined with any other protocol.
5. In the **Edit Storage Virtual Machine** pane, click on **Resource Allocation**.
6. In the **Resource Allocation** tab, you can choose not to delegate volume creation or you can select an aggregate to provision the volumes automatically.
7. Click on the **Services** tab to configure the Name Service Switch details.
8. Click **Save and Close**
   The NVMe protocol is configured on the SVM. After the protocol has been configured, you can start or stop the service using **SVM Settings**

### Related concepts

*Setting up NVMe* on page 256

You can set up the NVMe protocol for an SVM using System Manager. When the NVMe protocol is enabled on the SVM, you can then provision a namespace or namespaces and assign them to a host and a subsystem.

## Delegating administration to SVM administrators

After setting up a functional storage virtual machine (SVM) or SVMs with basic network configuration, you can optionally delegate the administration of the SVM to SVM administrators.

### About this task

SVM administrators cannot use System Manager to manage delegated SVMs. Administrators can manage them only by using the command-line interface (CLI).

### Steps

1. In the **Administrator Details** section, set up a password for the `vsadmin` user account.
2. If you want a dedicated LIF for SVM management, select **Create a LIF for SVM management**, and then specify the network details.

   A dedicated SVM management LIF is required for SAN protocols, where data and management protocols cannot share the same LIF. SVM management LIFs can be created only on data ports.
3. Specify the network details, including subnet details, for creating iSCSI LIFs:

| If you want to… | Then… |
|---|---|
| Specify the IP address by using a subnet | **a.** Select **Using a subnet**.<br>**b.** In the Add Details dialog box, select the subnet from which the IP address must be assigned.<br>For intercluster LIFs, only the subnets that are associated with the selected IPspace are displayed.<br>**c.** If you want to assign a specific IP address to the interface, select **Use a specific IP address**, and then type the IP address.<br>The IP address that you specify is added to the subnet if the IP address is not already present in the subnet range.<br>**d.** Click **OK**. |

| If you want to… | Then… |
|---|---|
| Specify the IP address manually without using a subnet | **a.** Select **Without a subnet**.<br>**b.** In the Add Details dialog box, perform the following steps:<br>**i.** Specify the IP address and the network mask or prefix.<br>**ii.** Optional: Specify the gateway.<br>**iii.** If you do not want to use the default value for the Destination field, specify a new destination value.<br>If you do not specify a custom value, the Destination field is populated with the default value based on the family of the IP address.<br>If a route does not exist, a new route is automatically created based on the gateway and destination.<br>**c.** Click **OK**. |

4. Specify a port for creating a data LIF:
   a. Click **Browse**.
   b. Select a port from the Select Network Port or Adapter dialog box.
   c. Click **OK**.

### Result

The `vsadmin` account is unlocked and configured with the password.

The default access methods for the `vsadmin` account are ONTAP API `(ontapi)` and SSH `(ssh)`. The SVM administrator can log in to the storage system by using the management IP address.

### After you finish

You must assign aggregates to the SVM by using the Edit Storage Virtual Machine dialog box.

**Note:** If the SVM does not have any assigned aggregates, the SVM administrator cannot create volumes.

## Creating FlexVol volumes

You can create a FlexVol volume for your data by using the Create Volume dialog box in System Manager. You should always create a separate volume for your data rather than storing data in the root volume.

### Before you begin

- The cluster must contain a non-root aggregate and a storage virtual machine (SVM).
- If you want to create read/write (rw) volumes, you must have configured the protocols for the SVM, and you must have installed either the SnapMirror license or the SnapVault license.

  If you have not configured the protocols but have installed any one of these licenses, you can create only data protection (DP) volumes.
- For creating an encrypted volume, you must have installed the volume encryption license by using System Manager and enabled "key-manager setup" by using the command-line interface (CLI).

  You must refresh your web browser after enabling "key-manager setup".

### About this task

- You can enable storage Quality of Service (QoS) only for a read/write (rw) volume.
- When you create a DP volume on the sync-source SVM in a MetroCluster configuration, the volume is not replicated on the sync-destination SVM.

- When you create a DP volume in a MetroCluster configuration, the source volume is not replicated (mirrored or vaulted) in the destination SVM.
- In a MetroCluster configuration, System Manager displays only the following aggregates for creating volumes:
  - In normal mode, when you create volumes on sync-source SVMs or data-serving SVMs in the primary site, only those aggregates that belong to the cluster in the primary site are displayed.
  - In switched-over mode, when you create volumes on sync-destination SVMs or data-serving SVMs in the surviving site, only switched-over aggregates are displayed.
- You cannot encrypt a volume in Cloud Volumes ONTAP.
- If encryption is enabled on the source volume and if the destination cluster is running a version of ONTAP software earlier than ONTAP 9.3, then encryption is disabled on the destination volume by default.

### Steps

1. Click **Storage** > **Volumes**.
2. Click **Create** > **Create FlexVol**.
3. Browse and select the SVM in which you want to create the volume.
4. In the **Create Volume** dialog box, specify a name for the volume.
5. Select the containing aggregate for the volume.
6. Select the **Encrypted** check box to enable encryption for the volume.

   This option is available only if you have enabled the Volume Encryption license and if the corresponding platform is capable of supporting encryption.
7. Select the type of storage for which you are creating this volume.

   You must select **Data Protection** if you are creating a SnapMirror destination volume. You are provided read-only access to this volume.
8. Specify the tiering policy for the volume.
9. Specify the size of the volume and the percentage of the total volume size that you want to reserve for Snapshot copies.

   The default space reserved for Snapshot copies is zero percent for SAN and VMware volumes. For NAS volumes, the default is 5 percent.
10. Select **Default**, **Thin provisioned**, or **Thick provisioned** for the volume.

    When thin provisioning is enabled, space is allocated to the volume from the aggregate only when data is written to the volume.

    **Note:**

    - For AFF storage systems, the value of thin provisioning is "Default", and for other storage systems, the value of thick provisioning is "Default".
    - For FabricPool-enabled aggregates, the value of thin provisioning is "Default".
11. If you want to enable deduplication on the volume, make the required changes in the **Storage Efficiency** tab.

    System Manager uses the default deduplication schedule. If the specified volume size exceeds the limit that is required for running deduplication, the volume is created and deduplication is not enabled.

    For systems with All Flash Optimized personality, inline compression and the `auto` deduplication schedule are enabled by default.
12. If you want to enable storage QoS for the FlexVol volume to manage workload performance, select the **Manage Storage Quality of Service** check box in the **Quality of Service** tab.

**13.** Create a new storage QoS policy group or select an existing policy group to control the input/output (I/O) performance of the FlexVol volume:

| If you want to... | Do this... |
|---|---|
| Create a new policy group | **a.** Select **New Policy Group**.<br>**b.** Specify the policy group name.<br>**c.** Specify the minimum throughput limit.<br>• You can set the minimum throughput limit only on an AFF platform.<br>• You cannot set the minimum throughput limit for volumes on a FabricPool-enabled aggregate.<br>• If you do not specify the minimum throughput value or if the minimum throughput value is set to 0, the system automatically displays "None" as the value.<br>This value is case-sensitive.<br>**d.** Specify the maximum throughput limit to ensure that the workload of the objects in the policy group does not exceed the specified throughput limit.<br>• The minimum throughput limit and the maximum throughput limit must be of the same unit type.<br>• If you do not specify the minimum throughput limit, then you can set the maximum throughput limit in IOPS, B/s, KB/s, MB/s, and so on.<br>• If you do not specify the maximum throughput value, the system automatically displays "Unlimited" as the value.<br>This value is case-sensitive. The unit that you specify does not affect the maximum throughput. |
| Select an existing policy group | **a.** Select **Existing Policy Group**, and then click **Choose** to select an existing policy group from the Select Policy Group dialog box.<br>**b.** Specify the minimum throughput limit.<br>• You can set the minimum throughput limit only on an AFF platform.<br>• You cannot set the minimum throughput limit for volumes on a FabricPool-enabled aggregate.<br>• If you do not specify the minimum throughput value or if the minimum throughput value is set to 0, the system automatically displays "None" as the value.<br>This value is case-sensitive.<br>**c.** Specify the maximum throughput limit to ensure that the workload of the objects in the policy group does not exceed the specified throughput limit.<br>• The minimum throughput limit and the maximum throughput limit must be of the same unit type.<br>• If you do not specify the minimum throughput limit, then you can set the maximum throughput limit in IOPS, B/s, KB/s, MB/s, and so on.<br>• If you do not specify the maximum throughput value, the system automatically displays "Unlimited" as the value.<br>This value is case-sensitive. The unit that you specify does not affect the maximum throughput.<br>If the policy group is assigned to more than one object, the maximum throughput that you specify is shared among the objects. |

**14.** Perform the following steps in the **Protection** tab to protect the volume:

a. Enable **Volume Protection**.
b. Select a relationship for the volume by using the **Volume Relationship Type** list.
c. Select a cluster and an SVM for the destination volume.

   If the selected cluster is running a version of ONTAP software earlier than ONTAP 9.3, then only peered SVMs are listed. If the selected cluster is running ONTAP 9.3 or later, peered SVMs and permitted SVMs are listed.
d. Enter the volume name suffix.

**15.** Click **Create**.

**16.** Verify that the volume that you created is included in the list of volumes in the **Volume** window.
The volume is created with UNIX-style security and UNIX 700 "read write execute" permissions for the owner.

### Related reference

*Volumes window* on page 203

You can use the Volumes window to manage your volumes, and to display information about these volumes.

## Creating SnapLock volumes

You can use System Manager to create a SnapLock Compliance volume or a SnapLock Enterprise volume. When you create a volume, you can also set retention times, and choose whether to automate setting the WORM state on data in the volume.

### Before you begin

- The SnapLock license must have been installed.
- The SnapLock aggregate must be online.
- For creating an encrypted volume, you must have installed the volume encryption license by using System Manager and enabled "key-manager setup" by using the command-line interface (CLI).

  You must refresh your web browser after enabling "key-manager setup".

### About this task

- You can delete a complete SnapLock Enterprise volume or a file in a SnapLock Enterprise volume; however, you cannot delete only the data within a file in a SnapLock Enterprise volume.
- You cannot delete a SnapLock Compliance volume if data is committed to the volume.
- You cannot encrypt a volume in Cloud Volumes ONTAP.
- If encryption is enabled on the source volume and if the destination cluster is running a version of ONTAP software earlier than ONTAP 9.3, then encryption is disabled on the destination volume by default.

### Steps

**1.** Click **Storage** > **Volumes**.

**2.** Click **Create** > **Create FlexVol**.

**3.** Browse and select the storage virtual machine (SVM) in which you want to create the volume.

**4.** In the **Create Volume** dialog box, specify a new name if you want to change the default name of the volume.

   You cannot change the name of a SnapLock Compliance volume after you create the volume.

**5.** Select the container aggregate for the volume.

You must select a SnapLock Compliance aggregate or SnapLock Enterprise aggregate to create a SnapLock volume. The volume inherits the SnapLock type from the aggregate, and the SnapLock type cannot be changed after the volume is created; therefore, you must select the correct aggregate.

6. Select the **Encrypted** checkbox to enable encryption for the volume.

   This option is available only if you have enabled the Volume Encryption license and if the corresponding platform is capable of supporting encryption.

7. Select the type of storage for which you are creating this volume.

   If you are creating a SnapMirror destination volume, you must select **Data Protection**. You are provided read-only access to this volume.

8. Specify the size of the volume and the percentage of the total volume size that you want to reserve for Snapshot copies.

   The default space that is reserved for Snapshot copies is zero percent for SAN and VMware volumes. For NAS volumes, the default is 5 percent.

9. Optional: Select **Thin Provisioned** to enable thin provisioning for the volume.

   When thin provisioning is enabled, space is allocated to the volume from the aggregate only when data is written to the volume.

10. Optional: Make the required changes in the **Storage Efficiency** tab to enable deduplication on the volume.

    System Manager uses the default deduplication schedule. If the specified volume size exceeds the limit that is required for running deduplication, the volume is created, and deduplication is not enabled.

11. Select the **SnapLock** tab, and then perform the following steps:

    a. Optional: Specify the autocommit period.

       The file in the volume will remain unchanged for the period that you specify before the file is committed to the WORM state. To set files to the WORM state manually, you must select **Not specified** as the autocommit setting.

       The values must be in the range of 5 minutes to 10 years.

    b. Specify the minimum retention period and maximum retention period.

       The values must be in the range of 1 day through 70 years or Infinite.

    c. Select the default retention period.

       The default retention period must be within the specified minimum retention period and maximum retention period.

12. Optional: Select the **Manage Storage Quality of Service** checkbox in the **Quality of Service** tab to enable storage QoS for the FlexVol volume in order to manage workload performance.

13. Create a new storage QoS policy group or select an existing policy group to control the input/output (I/O) performance of the FlexVol volume.

| If you want to... | Do this... |
|---|---|
| Create a new policy group | **a.** Select **New Policy Group**.<br>**b.** Specify the policy group name.<br>**c.** Specify the minimum throughput limit.<br>• You can set the minimum throughput limit only on an AFF platform.<br>• You cannot set the minimum throughput limit for volumes on a FabricPool-enabled aggregate.<br>• If you do not specify the minimum throughput value or if the minimum throughput value is set to 0, the system automatically displays "None" as the value.<br>This value is case-sensitive.<br>**d.** Specify the maximum throughput limit to ensure that the workload of the objects in the policy group does not exceed the specified throughput limit.<br>• The minimum throughput limit and the maximum throughput limit must be of the same unit type.<br>• If you do not specify the minimum throughput limit, then you can set the maximum throughput limit in IOPS, B/s, KB/s, MB/s, and so on.<br>• If you do not specify the maximum throughput value, the system automatically displays "Unlimited" as the value.<br>This value is case-sensitive. The unit that you specify does not affect the maximum throughput. |
| Select an existing policy group | **a.** Select **Existing Policy Group**, and then click **Choose** to select an existing policy group from the Select Policy Group dialog box.<br>**b.** Specify the minimum throughput limit.<br>• You can set the minimum throughput limit only on an AFF platform.<br>• You cannot set the minimum throughput limit for volumes on a FabricPool-enabled aggregate.<br>• If you do not specify the minimum throughput value or if the minimum throughput value is set to 0, the system automatically displays "None" as the value.<br>This value is case-sensitive.<br>**c.** Specify the maximum throughput limit to ensure that the workload of the objects in the policy group does not exceed the specified throughput limit.<br>• The minimum throughput limit and the maximum throughput limit must be of the same unit type.<br>• If you do not specify the minimum throughput limit, then you can set the maximum throughput limit in IOPS, B/s, KB/s, MB/s, and so on.<br>• If you do not specify the maximum throughput value, the system automatically displays "Unlimited" as the value.<br>This value is case-sensitive. The unit that you specify does not affect the maximum throughput.<br>If the policy group is assigned to more than one object, the maximum throughput that you specify is shared among the objects. |

**14.** Perform the following steps in the **Protection** tab to protect the volume:

a. Enable **Volume Protection**.

b. Select a relationship for the volume by using the **Volume Relationship Type** list.

c. Select a cluster and an SVM for the destination volume.

   If the selected cluster is running a version of ONTAP software earlier than ONTAP 9.3, then only peered SVMs are listed. If the selected cluster is running ONTAP 9.3 or later, peered SVMs and permitted SVMs are listed.

d. Specify the volume name suffix.

**15.** Click **Create**.

**16.** Verify that the volume that you created is included in the list of volumes in the **Volume** window.

### Result

The volume is created with UNIX-style security and UNIX 700 "read write execute" permissions for the owner.

# Setting up SAML authentication

You can set up Security Assertion Markup Language (SAML) authentication so that remote users are authenticated through a secure identity provider (IdP) before they log in to System Manager.

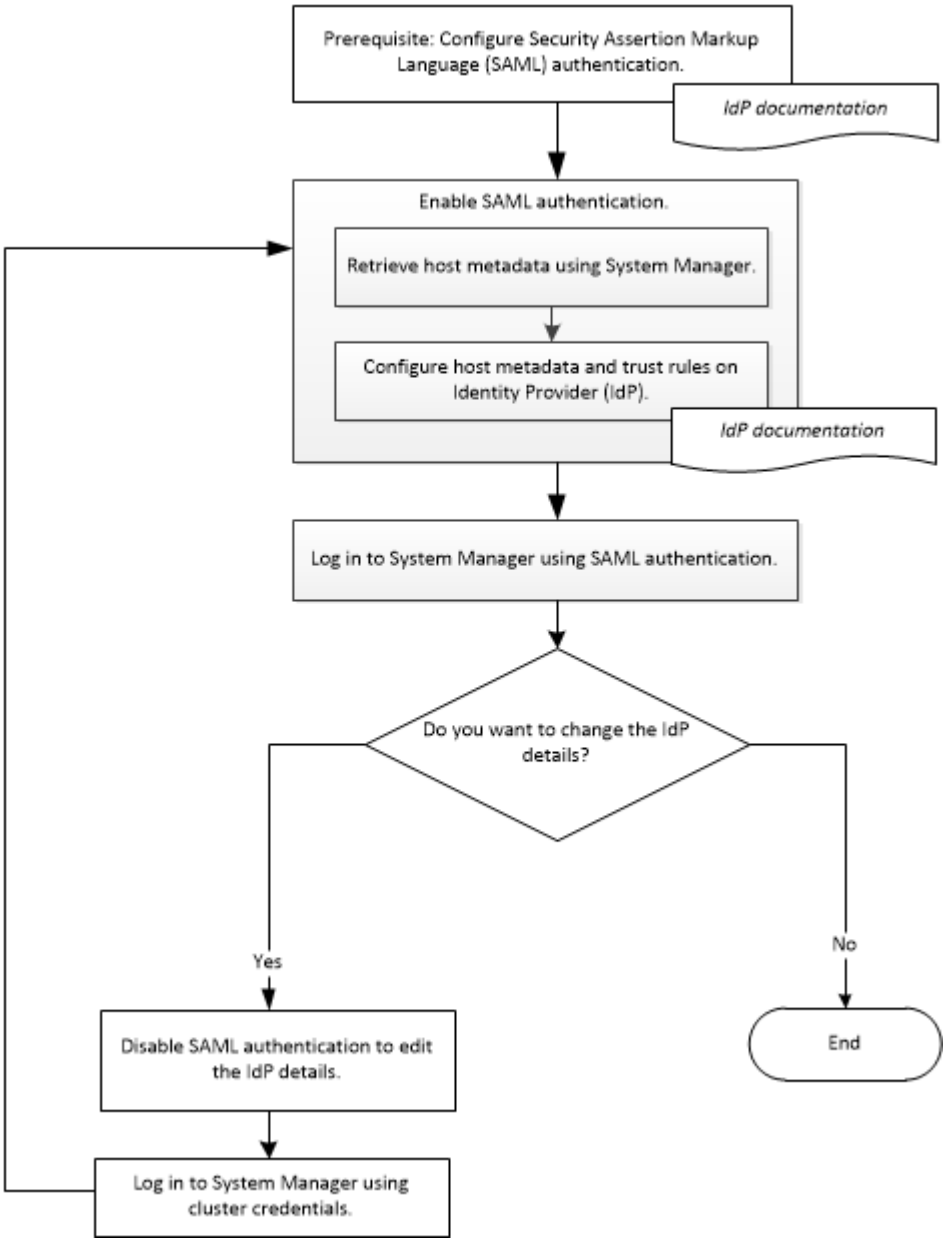


## Enabling SAML authentication

You can use System Manager to configure Security Assertion Markup Language (SAML) authentication so that remote users can log in by using a secure identity provider (IdP).

### Before you begin

- The IdP that you plan to use for remote authentication must be configured.

  **Note:** See the documentation that is provided by the IdP that you have configured.

- You must have the URI of the IdP.

### About this task

The IdPs that have been validated with System Manager are Shibboleth and Active Directory Federation Services.

**Note:** After SAML authentication is enabled, only remote users can access the System Manager GUI. Local users cannot access the System Manager GUI after SAML authentication is enabled.

### Steps

1. Click **Configuration** > **Authentication**.
2. Select the **Enable SAML authentication** check box.
3. Configure System Manager to use SAML authentication:
   a. Enter the URI of the IdP.
   b. Enter the IP address of the host system.
   c. Optional: If required, change the host system certificate.
4. Click **Retrieve Host Metadata** to retrieve the host URI and host metadata information.
5. Copy the host URI or host metadata details, access your IdP, and then specify the host URI or host metadata details and the trust rules in the IdP window.

   **Note:** See the documentation that is provided by the IdP that you have configured.
6. Click **Save**.
   The IdP login window is displayed.
7. Log in to System Manager by using the IdP login window.

   After the IdP is configured, if the user tries to log in by using the fully qualified domain name (FQDN), IPv6, or a cluster management LIF, then the system automatically changes the IP address to the IP address of the host system that was specified during the IdP configuration.

### Related tasks

*Accessing a cluster by using the OnCommand System Manager browser-based graphic interface* on page 28

If you prefer to use a graphic interface instead of the command-line interface for accessing and managing a cluster, you can do so by using OnCommand System Manager, which is included with ONTAP as a web service, is enabled by default, and is accessible by using a browser.

## Disabling SAML authentication

You can disable Security Assertion Markup Language (SAML) authentication if you want to disable remote access to System Manager, or to edit the SAML configuration.

### About this task

Disabling SAML authentication does not delete SAML configuration.

### Steps

1. Click **Configuration** > **Authentication**.
2. Clear the **Enable SAML authentication** check box.
3. Click **Save**.
   System Manager restarts.
4. Log in to System Manager by using the cluster credentials.

### Related tasks

*Accessing a cluster by using the OnCommand System Manager browser-based graphic interface* on page 28

If you prefer to use a graphic interface instead of the command-line interface for accessing and managing a cluster, you can do so by using OnCommand System Manager, which is included with ONTAP as a web service, is enabled by default, and is accessible by using a browser.

# Setting up peering

Setting up peering involves creating intercluster logical interfaces (LIFs) on each node, creating cluster peering, and creating SVM peering.

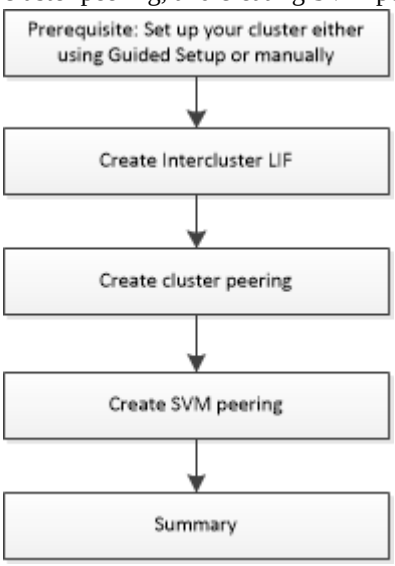


## Prerequisites for cluster peering

Before you set up cluster peering, you should confirm that the connectivity, port, IP address, subnet, firewall, and cluster-naming requirements are met.

### Connectivity requirements

Intercluster LIFs must have *pair-wise full-mesh connectivity:* Every intercluster LIF on the local cluster must be able to communicate with every intercluster LIF on the remote cluster.

Although it is not required, it is typically simpler to configure the IP addresses used for intercluster LIFs in the same subnet. The IP addresses can reside in the same subnet as data LIFs, or in a different subnet. The subnet used in each cluster must meet the following requirements:

- The subnet must belong to the broadcast domain that contains the ports that are used for intercluster communication.

Intercluster LIFs can have an IPv4 address or an IPv6 address.

### Port requirements

You can use dedicated ports for intercluster communication, or share ports used by the data network. Ports must meet the following requirements:

- All ports that are used to communicate with a given remote cluster must be in the same IPspace.

  You can use multiple IPspaces to peer with multiple clusters. Pair-wise full-mesh connectivity is required only within an IPspace.
- The broadcast domain that is used for intercluster communication must include at least two ports per node so that intercluster communication can fail over from one port to another port.

  Ports added to a broadcast domain can be physical network ports, VLANs, or interface groups (ifgrps).
- All ports must be cabled.
- All ports must be in a healthy state.

- The MTU settings of the ports must be consistent.

### Firewall requirements

Firewalls and the intercluster firewall policy must allow the following protocols:

- ICMP service
- TCP to the IP addresses of all the intercluster LIFs over the ports 10000, 11104, and 11105
- HTTPS

The default **`intercluster`** firewall policy allows access through the HTTPS protocol and from all IP addresses (0.0.0.0/0). You can modify or replace the policy if necessary.

### Cluster requirements

Clusters must meet the following requirements:

- The time on the clusters in a cluster peering relationship must be synchronized within 300 seconds (5 minutes).
  Cluster peers can be in different time zones.

### Related information

*NetApp Documentation: ONTAP 9*

## Creating intercluster LIFs

Creating intercluster logical interfaces (LIFs) enables the cluster network to communicate with a node. You must create an intercluster LIF on each node in each cluster for which you want to create a peer relationship.

### Steps

1. Click **Configuration** > **Advanced Cluster Setup**.
2. In the **Setup Advanced Cluster Features** window, click **Proceed** next to the **Cluster Peering** option.
3. Select an IPspace from the **IPspace** list.
4. Enter the IP address, port, network mask, and gateway details of each node.
5. Click **Submit and Continue**.

### After you finish

You should enter the cluster details in the Cluster Peering window to continue with cluster peering.

## Creating cluster peer relationships

You can create an authenticated cluster peer relationship to connect clusters so that the clusters in the peer relationship can communicate securely with each other.

### Before you begin

- You must have reviewed and completed the requirements for performing this task.
  *Prerequisites for cluster peering*
- You must have created intercluster logical interfaces (LIFs).

### About this task

- If you want to create a peer relationship with a cluster running Data ONTAP 8.2.2 or earlier, you must use the command-line interface (CLI).
- In a MetroCluster configuration, when you create a peer relationship between the primary cluster and an external cluster, it is a best practice to create a peer relationship between the surviving site cluster and the external cluster as well.

- You can create a custom passphrase or you can use the system-generated passphrase to authenticate the cluster peer relationship.

### Steps

1. Click **Configuration** > **Advanced Cluster Setup**.
2. In the **Target Cluster Intercluster LIF IP addresses** filed, enter the IP addresses of the remote cluster's intercluster LIFs.
3. In the **Passphrase** field, specify a passphrase for the cluster peer relationship.

   If you specify a custom passphrase, the passphrase that you specify will be validated against the passphrase of the peered cluster to ensure an authenticated cluster peer relationship.

   If the names of the local cluster and remote cluster are identical, and if you are using a custom passphrase, an alias is created for the remote cluster.
4. Optional: To generate a passphrase from the remote cluster, enter the management IP address of the remote cluster.
5. Initiate cluster peering.

| If you want to... | Do this... |
|---|---|
| Initiate cluster peering from the initiator cluster | Click **Initiate Cluster Peering**. |
| Initiate cluster peering from the remote cluster<br>(Applicable if you have created a custom passphrase) | **a.** Enter the management address of the remote cluster.<br>**b.** Click the **Management URL** link to access the remote cluster.<br>**c.** Click **Create Cluster Peering**.<br>**d.** Specify the intercluster LIF IP addresses and passphrase of the initiator cluster.<br>**e.** Click **Initiate Peering**.<br>**f.** Access the initiator cluster, and then click **Validate Peering**. |

### After you finish

You should specify the SVM details in the SVM Peering window to continue with the peering process.

## Creating SVM peers

SVM peering enables you to establish a peer relationship between two storage virtual machines (SVMs) for data protection.

### Before you begin

You must have created a peer relationship between the clusters in which the SVMs that you plan to peer reside.

### About this task

- The clusters that you can select as target clusters are listed when you create SVM peers by using the **Configuration** > **SVM Peers** window.
- If the target SVM resides on a cluster in a system running ONTAP 9.2 or earlier, SVM peering cannot be accepted by using System Manager.

  **Note:** In such a scenario, you can use the command-line interface (CLI) to accept SVM peering.

### Steps

1. Select the initiator SVM.
2. Select the target SVM from the list of permitted SVMs.
3. Specify the name of the target SVM in the **Enter an SVM** field.

   **Note:** If you have navigated from the **Configuration** > **SVM Peers** window, you should select the target SVM from the list of peered clusters.

4. Initiate SVM peering.

| If you want to... | Do this... |
|---|---|
| Initiate SVM peering from the initiator cluster | Click Initiate SVM Peering. |
| Accept SVM peering from the remote cluster | **Note:** Applicable for non-permitted SVMs<br>**a.** Specify the management address of the remote cluster.<br>**b.** Click the **Management URL** link to access the SVM Peer window of the remote cluster.<br>**c.** On the remote cluster, accept the **Pending SVM Peer** request.<br>**d.** Access the initiator cluster, and then click **Validate Peering**. |

5. Click **Continue**.

### After you finish

You can view the intercluster LIFs, cluster peer relationship, and SVM peer relationship in the Summary window.

## What passphrases are

You can use a passphrase to authorize peering requests. You can use a custom passphrase or a system-generated passphrase for cluster peering.

- You can generate a passphrase on the remote cluster.
- The minimum required length for a passphrase is eight characters.
- The passphrase is generated based on the IPspace.
- If you are using a system-generated passphrase for cluster peering, after you enter the passphrase in the initiator cluster, peering is authorized automatically.
- If you are using a custom passphrase for cluster peering, you have to navigate to the remote cluster to complete the peering process.

# Managing clusters

You can use System Manager to manage clusters.

**Related information**

*ONTAP concepts*

## Understanding quorum and epsilon

Quorum and epsilon are important measures of cluster health and function that together indicate how clusters address potential communications and connectivity challenges.

*Quorum* is a precondition for a fully functioning cluster. When a cluster is in quorum, a simple majority of nodes are healthy and can communicate with each other. When quorum is lost, the cluster loses the ability to accomplish normal cluster operations. Only one collection of nodes can have quorum at any one time because all of the nodes collectively share a single view of the data. Therefore, if two non-communicating nodes are permitted to modify the data in divergent ways, it is no longer possible to reconcile the data into a single data view.

Each node in the cluster participates in a voting protocol that elects one node *master*; each remaining node is a *secondary*. The master node is responsible for synchronizing information across the cluster. When quorum is formed, it is maintained by continual voting. If the master node goes offline and the cluster is still in quorum, a new master is elected by the nodes that remain online.

Because there is the possibility of a tie in a cluster that has an even number of nodes, one node has an extra fractional voting weight called *epsilon*. If the connectivity between two equal portions of a large cluster fails, the group of nodes containing epsilon maintains quorum, assuming that all of the nodes are healthy. For example, the following illustration shows a four-node cluster in which two of the nodes have failed. However, because one of the surviving nodes holds epsilon, the cluster remains in quorum even though there is not a simple majority of healthy nodes.

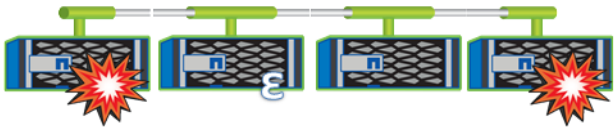

Epsilon is automatically assigned to the first node when the cluster is created. If the node that holds epsilon becomes unhealthy, takes over its high-availability partner, or is taken over by its high-availability partner, then epsilon is automatically reassigned to a healthy node in a different HA pair.

Taking a node offline can affect the ability of the cluster to remain in quorum. Therefore, ONTAP issues a warning message if you attempt an operation that will either take the cluster out of quorum or else put it one outage away from a loss of quorum. You can disable the quorum warning messages by using the `cluster quorum-service options modify` command at the advanced privilege level.

In general, assuming reliable connectivity among the nodes of the cluster, a larger cluster is more stable than a smaller cluster. The quorum requirement of a simple majority of half the nodes plus epsilon is easier to maintain in a cluster of 24 nodes than in a cluster of two nodes.

A two-node cluster presents some unique challenges for maintaining quorum. Two-node clusters use *cluster HA*, in which neither node holds epsilon; instead, both nodes are continuously polled to ensure that if one node fails, the other has full read-write access to data, as well as access to logical interfaces and management functions.

# Dashboard window

The Dashboard window contains multiple panels that provide cumulative at-a-glance information about your system and its performance.

You can use the Dashboard window to view information about important alerts and notifications, the efficiency and capacity of aggregates and volumes, the nodes that are available in a cluster, the status of the nodes in a high-availability (HA) pair, the most active applications and objects, and the performance metrics of a cluster or a node.

**Alerts and Notifications**

Displays all alerts in red, such as emergency EMS events, offline node details, broken disk details, license entitlements that are in high risk, and offline network port details. Displays all notifications in yellow, such as health monitor notifications that occurred in the past 24 hours at the cluster level, license entitlements that are in medium risk, unassigned disk details, the number of migrated LIFs, volume move operations that failed, and volume move operations that required administrative intervention in the past 24 hours.

The Alerts and Notifications panel displays up to three alerts and notifications beyond which a View-All link is displayed. You can click the View-All link to view more information about the alerts and notifications.

The refresh interval for the Alerts and Notifications panel is one minute.

**Health Overview**

Displays the aggregates and volumes that are nearing capacity, the storage efficiency of a cluster or node, and the protection details of top volumes.

The Capacity tab displays the top online aggregates that are nearing capacity, in descending order of used space.

The Capacity tab provides a link to the number of volumes with the highest capacity utilized when you enter a valid value in the Volumes exceeding used capacity of field. It also displays the amount of inactive (cold) data available in the cluster.

The Efficiency tab displays the storage efficiency savings for a cluster or node. You can view the total logical space used, total physical space used, overall savings from storage efficiency, data reduction ratio, FlexClone volume ratio, and Snapshot copies ratio. You can select a cluster or a specific node to view the storage efficiency savings.

> **Note:** During a takeover operation or giveback operation, the storage efficiency data may not be fully reported. In such cases, the reported storage efficiency data of these operations is corrected after some time, depending on the number of Snapshot copies across all of the volumes in the nodes.

The refresh interval for the Health Overview panel is 15 minutes.

The Protection tab displays information about cluster-wide volumes that are missing protection relationships. Only the FlexVol volumes and FlexGroup volumes that meet the following criteria are displayed:

- The volumes are RW volumes and are online.
- The aggregate containing the volumes is online.
- The volumes have protection relationships and are not yet initialized.

You can navigate to the Volumes window to view the volumes that do not have a protection relationship.

The Protection tab also displays the top five SVMs that have the highest number of volumes that are missing protection relationships.

**Nodes**

Displays a pictorial representation of the number and names of the nodes that are available in the cluster, and the status of the nodes that are in an HA pair. You should position the cursor over the pictorial representation of the nodes to view the status of the nodes in an HA pair.

You can view more information about all of the nodes by using the Nodes link. You can also click the pictorial representation to view the model of the nodes and the number of aggregates, storage pools, shelves, and disks that are available in the nodes. You can manage the nodes by using the Manage Nodes link. You can manage the nodes in an HA pair by using the Manage HA link.

The refresh interval for the Nodes panel is 15 minutes.

**Applications and Objects**

The Applications tab displays information about the top five applications of the cluster. You can view the top five applications based on either IOPS (from low to high or from high to low) or capacity (from low to high or from high to low). You should click the specific bar chart to view more information about the application. For capacity, the total space, used space, and available space are displayed, and for IOPS, the IOPS details are displayed. You can click **View details** to open the Applications window of the specific application.

The Objects tab displays information about the top five active clients and files in the cluster. You can view the top five active clients and files based on IOPS or throughput.

The refresh interval for the Applications and Objects panel is one minute.

**Performance**

Displays the average performance metrics, read performance metrics, and write performance metrics of the cluster based on latency, IOPS, and throughput. The average performance metrics is displayed by default. You can click Read or Write to view the read performance metrics or write performance metrics, respectively. You can view the performance metrics of the cluster or a node.

If the information about cluster performance cannot be retrieved from ONTAP, you cannot view the respective graph. In such cases, System Manager displays the specific error message.

The refresh interval for the charts in the Performance panel is 15 seconds.

## Monitoring a cluster using the dashboard

The dashboard in System Manager enables you to monitor the health and performance of a cluster. You can also identify hardware problems and storage configuration issues by using the dashboard.

### Step

Click the **Dashboard** tab to view the health and performance dashboard panels.

# Applications

You can use predefined application templates in System Manager to create new configurations that are based on existing application templates . You can then provision instances of the application in ONTAP.

You configure applications by clicking **Applications &Tiers** > **Applications**.

The following applications can be configured in System Manager:

**General Applications**

- NAS Container (Volume is exported to NFS or CIFS clients)
- General SAN Application (Set of LUNs exported to the application server)

**Databases**

- MongoDB (over SAN)
- Oracle (over NFS or SAN)
- Oracle (Real Application Cluster over NFS or SAN)
- Microsoft SQL Server (over SAN)

**Virtual Infrastructure**

- Virtual Servers (with VMware, Hyper-V or XEN)

### Related information

*ONTAP concepts*

### Steps

1. *Provisioning a basic template* on page 68
2. *Adding applications to System Manager* on page 69
3. *Application provisioning settings* on page 69

## Provisioning a basic template

You can use System Manager to quickly provision basic templates for SQL, VMware and SAP HANA on AFF, SAN, and NAS optimized clusters.

### About this task

As the cluster administrator, you can provision applications by configuring one of the basic templates provided. The example describes how to configure the **SMB SQL Server**.

### Steps

1. Click **Applications & Tiers** > **Applications**
2. In the **Basic** tab, select the **SMB SQL Server** template.
3. In the **Database Details** section, specify the following:
   - Database name
   - Database size
   - Log size
   - Tempdb size
   - Number of server cores
   - Span HA Controller Notes
4. In the **SQL Server Account Details** section, specify the following:
   - Optional: SQL server installation account
   - SQL server service account
   - Optional: SQL server agent service account
5. Click **Provision Storage**

### Result

The SMB SQL Server application is provisioned.

### Related information

*Refer to Application Provisioning Settings for field descriptions* on page 69

## Adding applications to System Manager

You can use the Enhanced tab to add general applications, databases, and virtual servers to System Manager.

### About this task

The following procedure describes how to add a **Microsoft SQL Server** instance to System Manager.

### Steps

1. Click **Applications & Tiers** > **Applications**
2. In the **Enhanced** tab, click **Add**
3. Select an application type from the menu.

   **Note:** The dropdown list includes a list of all available application types and template types.

   The Add Microsoft SQL Server Instance window is displayed.
4. Specify the following details:
   - Database name
   - Database size and the required ONTAP service level
   - Number of server cores
   - Log size and the required ONTAP service level
   - Tempdb
     Specify if the server should be provisioned for Tempdb
   - Host operating system
   - LUN format
   - Host mapping
5. Click **Add Application**

### Result

The Microsoft SQL Server instance is added to System Manager.

## Application provisioning settings

When setting up the basic template for a database, server, or virtual desktop, you must provide details to System Manager. This section describes the fields in each basic template. Only the fields that are required for the specific application you are provisioning are displayed.

### Database Details

You enter the following information to provision the database applications:

**Database Name**

Mandatory: The name of the database you are configuring; this string is used as a prefix when provisioning storage for each database.

**Database Size**

Mandatory: The size of the database, in units of MB, GB, TB, or PB.

**Log Size**

Mandatory: The size of the database log in units of MB, GB, TB, or PB.

**Tempdb Size**

Mandatory: The size of the tempdb database in units of MB, GB, TB, or PB.

**Number of Server Cores (on the SQL server)**

Indicates the number of CPU cores on the databases server in increments of 2.

**Span HA Controller Nodes**

Specifies if storage objects should be created across a high-availability pair of nodes.

The following fields apply only when you are provisioning a *SAP HANA* database:

**Active SAP HANA Nodes**

The number of active SAP HANA nodes. The maximum number of nodes is 16.

**Memory Size per HANA Node**

The memory size of a single SAP HANA node.

**Data Disk Size per HANA Node**

The data disk size for each node.

**Note:** If set to 0, the memory size field above is used to calculate the size of the data area.

### SQL Server Account Details

You enter the following information to provide full control access to the SQL server accounts:

**Note:** The installation account is granted `SeSecurityPrivilege`.

**SQL Server Installation Account**

Optional: The domain user account used for the new SQL server installations. Enter in the format

**`domain\user`**

.

**SQL Server Service Account**

Mandatory: This is an existing domain account, specify as

**`domain\user`**

.

**SQL Server Agent Service Account**

Optional: This is this domain account if SQL server agent service is configured, specify in the format domain\user.

### Virtual Desktop Details

You enter the following information to provision storage for virtual desktop infrastructure over NAS:

**Select Hypervisor**

The hypervisor used for these volumes; the hypervisor determines the correct datastore protocol. The options are VMware, Hyper-V, or XenServer/KVM.

**Desktop Persistence**

Determines if the desktop is persistent or nonpersistent. Selecting the desktop persistence sets the default values for the volume such as Snapshot schedules and post-process deduplication policies. Inline efficiencies are enabled by default for all volumes.

**Note:** These policies can be modified manually after provisioning.

**Datastore Prefix**

The value entered is used to generate the names of the datastores and, if applicable, the export policy name or share name.

**Number of Desktops**

This number is used to determine the number of volumes created.

**Note:** This is not used to provision the virtual machines.

**Desktop Size**

This is used to determine the size of the volumes which should be provisioned in units of MB, GB, TB, or PB.

**Average Desktop Size (used for the SAN Virtual Desktop)**

This is used to determine the thin-provisioned size of each volume in units of MB, GB, TB, or PB.

### Initiator Details

You enter the following information to set up the initiator:

**Initiator Group**

You can select an existing group or create a new group.

**Initiator Group Name**

The name of the new initiator group.

**Initiators**

This is a comma-separated list of the initiators (WWPN or IQN) in the initiator group.

The following fields apply only to *SAP HANA* provisioning:

**Initiator OS Type**

This is the operating system type of the new initiator group.

**FCP Portset**

The FCP portset that the initiator group is bound to.

### Host Access Configuration

You enter the following information to configure the host access to the volumes:

**Volume Export Configuration**

Select the export policy to apply to the volumes during creation. The options are:

- Allow All
  This option implies that an export rule is created which permits read-write access to any clients.
- Create Custom Policy
  This option allows you to specify a list of host IP addresses to receive read-write access.

**Note:** You can modify the volume export policy later using System Manager workflows.

**Host IP Addresses**

This is a comma-separated list of IP addresses.

**Note:** For NFS-based systems, a new export policy is created using the datastore prefix and a rule is created in it to give access to the list of IP.

### Application Details

When the application is added, you can view the configuration settings in the **Overview** tab of the Application Details window. Other details such as NFS or CIFS Access and Permissions are displayed depending on the type of application that was set up.

**Type**

This is the type of general application, database, or virtual infrastructure that was created.

**SVM**

The name of the server virtual machine that the application was created on.

**Size**

The total size of the volume.

**Available**

The amount of space currently available in the volume.

**Protection**

The type of data protection configured.

You can expand the **Components** and **Volumes** panes for details of the space and IOPs performance.

**Note:** The used size displayed in the Components pane is different than the used size displayed in the CLI.

# Configuration update

You can use System Manager to configure the administration details of storage virtual machines (SVMs).

## Configuring the administration details of an SVM

You can use System Manager to quickly configure the administration details of a storage virtual machine (SVM). You can optionally delegate the administration of the SVM to SVM administrators.

### About this task

As an SVM administrator, you cannot use System Manager to manage delegated SVMs. You can manage the SVMs only by using the command-line interface (CLI).

### Steps

1. Click **Configuration** > **Configuration Updates**.
2. In the **SVMs** tab, select the node, and then click **Configure Administration Details**.
3. In the **Administrator Details** section, set up a password for the `vsadmin` user account.
4. If you want a dedicated LIF for SVM management, select **Create a LIF for SVM management**, and then specify the network details.

   A dedicated SVM management LIF is required for SAN protocols, where data and management protocols cannot share the same LIF. SVM management LIFs can be created only on data ports.
5. Specify the network details:

| If you want to… | Then… |
|---|---|
| Specify the IP address by using a subnet | **a.** Select **Using a subnet**.<br>**b.** In the Add Details dialog box, select the subnet from which the IP address must be assigned.<br>For intercluster LIFs, only the subnets that are associated with the selected IPspace are displayed.<br>**c.** If you want to assign a specific IP address to the interface, select **Use a specific IP address**, and then type the IP address.<br>The IP address that you specify is added to the subnet if that IP address is not already present in the subnet range.<br>**d.** Click **OK**. |
| Specify the IP address manually without using a subnet | **a.** Select **Without a subnet**.<br>**b.** In the Add Details dialog box, perform the following steps:<br>**i.** Specify the IP address and network mask or prefix.<br>**ii.** Optional: Specify the gateway.<br>The destination field is populated with the default value based on the family of the IP address.<br>**iii.** If you do not want the default value, specify a new destination value.<br>If a route does not exist, a new route is automatically created based on the gateway and destination.<br>**c.** Click **OK**. |

**6.** Specify a port to create a data LIF:
   a. Click **Browse**.
   b. In the **Select Network Port or Adapter** dialog box, select a port, and then click **OK**.

## Configuration Updates window

You can use the Configuration Updates window to update the configuration details of the cluster, storage virtual machine (SVM), and nodes.

### Tabs

**Nodes**

Enables you to configure details of the node.

**SVMs**

Enables you to configure details of the SVM.

### Nodes tab

### Command buttons

**Edit Node Name**

Opens the Edit Node Name dialog box, which enables you to modify the name of the node.

**Create Node-management LIF**

Opens the Create Node-management LIF dialog box, which enables you to create a node-management LIF for managing a specific node.

**Edit AutoSupport**

Opens the Edit AutoSupport Settings dialog box, which enables you to specify an email address from which email notifications are sent and to add multiple email addresses of the host names.

### SVMs tab

### Command button

**Configure Administration Details**

Opens the Configure Administration Details dialog box, which enables you configure the administration details of the SVM.

### Related tasks

*Creating a cluster* on page 23
You can use OnCommand System Manager to create and set up a cluster in your data center.

*Setting up a network when an IP address range is disabled* on page 25
You can set up a network by disabling an IP address range and entering individual IP addresses for cluster management, node management, and service provider networks.

# Service Processors

You can use a Services Processor to monitor and manage your storage system parameters such as temperature, voltage, current, and fan speeds through System Manager.

## Assigning IP addresses to Service Processors

You can use System Manager to assign IP addresses to all of your Service Processors at the same time and to use these Service Processors to monitor and manage various system parameters of your storage systems.

### Steps

1. Click **Configuration** > **Configuration Updates**.
2. In the **Service Processor** window, click **Global Settings**.
3. In the **Global Settings** dialog box, choose the source for assigning the IP addresses:

| If you want to... | Select the option... |
|---|---|
| Assign IP addresses automatically from a DHCP server | DHCP |
| Assign IP addresses from a subnet | Subnet |
| Manually provide IP addresses | Manual Assignment |

4. Click **Save**.

## Editing Service Processor settings

You can modify Service Processor attributes such as the IP address, the network mask or the prefix length, and the gateway address by using System Manager. You can also allocate IP addresses to Service Processors that do not have any IP addresses assigned.

### About this task

- You can edit the settings of a Service Processor that was assigned IP addresses manually.

- You cannot edit the settings of a Service Processor that was assigned IP addresses through a DHCP server or through a subnet.

### Steps

1. Click **Configuration** > **Service Processor**.
2. In the **Service Processor** window, select the Service Processor that you want to modify, and then click **Edit**.
3. In the **Edit Service Processor** dialog box, make the required changes, and then click **Save and Close**.

## Understanding the Service Processor

A Service Processor is a system-independent resource in the storage system that helps you to monitor and manage storage system parameters such as temperature, voltage, current, and fan speeds.

When the Service Processor detects an abnormal condition in any of the storage system parameters, the Service Processor logs an event, notifies ONTAP about the issue, and generates AutoSupport messages through email or through SNMP traps.

The Service Processor monitors ONTAP through a watchdog mechanism and can facilitate a quick failover to the partner node. The Service Processor also tracks numerous system events and saves the events in a log file. The events include boot progress, field-replaceable unit (FRU) changes, ONTAP generated events, and user transaction history.

The Service Processor can remotely log in and administer the storage system and can diagnose, shut down, power cycle, or reboot the system, regardless of the state of the storage system. In addition, the Service Processor provides remote diagnostic features.

The combined monitoring and managing capabilities of the Service Processor enables you to evaluate the storage system in the event of an issue, and then immediately perform effective service actions.

## Service Processors window

You can use the Service Processors window to view and modify Service Processors attributes, such as the IP address, network mask (IPv4) or prefix-length (IPv6), and gateway, and to configure the IP source for a Service Processor.

- *Command buttons* on page 75
- *Service processors list* on page 76
- *Details area* on page 76

### Command buttons

**Edit**

Opens the Edit Service Processor dialog box, which enables you to modify the IP address, network mask (IPv4) or prefix-length (IPv6), and gateway information of a Service Processor.

**Global Settings**

Opens the Global Settings dialog box, which allows you to configure the source of IP address for all your Service Processors as one of the following: DHCP, subnet, or manual.

**Refresh**

Updates the information in the window.

### Service processors list

**Node**

Specifies the node on which the Service Processor is located.

**IP Address**

Specifies the IP addresses of the Service Processor.

**Status**

Specifies the status the Service Processor, which can be online, offline, daemon offline, node offline, degraded, rebooted, or unknown.

**MAC Address**

Specifies the MAC address of the Service Processor.

### Details area

The area below the Service Processor list displays detailed information about the Service Processor, including network details, such as the IP address, network mask (IPv4) or prefix-length (IPv6), gateway, IP source, and MAC address, as well as general details, such as the firmware version and whether automatic update of the firmware is enabled.

### Related tasks

*Setting up a network when an IP address range is disabled* on page 25
You can set up a network by disabling an IP address range and entering individual IP addresses for cluster management, node management, and service provider networks.

# Cluster peers

Peered clusters are required for data replication using SnapMirror technology and SnapVault technology, and for data replication using SyncMirror technology in MetroCluster configurations. You can use System Manager to peer two clusters so that the peered clusters can coordinate and share resources between them.

## Modifying the cluster peer passphrase

For security reasons, you can modify the passphrase that is provided during cluster peer creation by using System Manager.

### Steps

1. Click **Configuration** > **Cluster Peers**.
2. Select the peered cluster, and click **Modify Passphrase**.
3. In the **Modify Passphrase** dialog box, enter a new passphrase, and then click **Modify**.

   **Note:** The minimum required length of the passphrase is eight characters.

4. Log in to the remote cluster, and perform steps 1 through 3 to modify the passphrase in the remote cluster.

   The authentication status for the local cluster is displayed as **`ok_and_offer`** until you modify the passphrase in the remote cluster.

## Modifying the peer network parameters

You can use System Manager to modify the IPspace and intercluster logical interfaces (LIFs) that are configured for the remote cluster. You can add new intercluster IP addresses or remove existing IP addresses.

### Before you begin

You must have at least one intercluster IP address to create the cluster peer relationship.

### Steps

1. Click **Configuration** > **Configuration Updates**.
2. In the **Cluster Details** pane, click **Cluster Peers**.
3. Select a peer cluster, and then click **Modify Peer Network Parameters**.
4. In the **Modify Peer Network Parameters** dialog box, select the IPspace, and then add or remove the intercluster IP addresses.

   You can add multiple IP addresses by using comma separators.
5. Click **Modify**.
6. Verify the changes that you made in the **Peers** window.

## Deleting cluster peer relationships

You can use System Manager to delete a cluster peer relationship if the relationship is no longer required. You must delete the cluster peering relationship from each of the clusters in the peer relationship.

### Steps

1. Click **Configuration** > **Cluster Peers**.
2. Select the cluster peer that you want to delete, and then click **Delete**.
3. Select the confirmation check box, and then click **Delete**.
4. Log in to the remote cluster, and perform steps 1 through 3 to delete the peer relationship between the local cluster and the remote cluster.

   The status of the peer relationship is displayed as "unhealthy" until the relationship is deleted from both the local cluster and the remote cluster.

## Peers window

You can use the Peers window to manage peer relationships, which enable you to move data from one cluster to another.

### Command buttons

**Create**

Opens the Create Cluster Peering dialog box, which enables you to create a relationship with a remote cluster.

**Local Cluster Passphrase**

Opens the Local Cluster Passphrase dialog box, which enables you to enter a new passphrase for the local cluster.

**Peer Cluster Network Parameters**

Opens the Peer Cluster Network Parameters dialog box, which enables you to modify the IPspace, add new intercluster IP addresses, or remove existing IP addresses.

You can add multiple IP addresses, separated by commas.

**Delete**

Opens the Delete Cluster Peer Relationship dialog box, which enables you to delete the selected peer cluster relationship.

**Refresh**

Updates the information in the window.

**Manage SVM Permissions**

Enables SVMs to automatically accept SVM peering requests.

**Generate Peering Passphrase**

Enables you to generate a passphrase for cluster peering.

### Peer cluster list

**Peer Cluster**

Specifies the name of the peer cluster in the relationship.

**Availability**

Specifies whether the peer cluster is available for communication.

**Authentication Status**

Specifies whether the peer cluster is authenticated or not.

**Local Cluster IPspace**

Displays IPspace associated to the cluster peer relation.

**Last Updated Time**

Displays the time at which peer cluster was last modified.

### Details area

The details area displays detailed information about the selected peer cluster relationship, including the active IP addresses discovered by the system to set up the intercluster network and the last updated time.

# High availability

You can use System Manager to create high availability (HA) pairs that provide hardware redundancy that is required for nondisruptive operations and fault tolerance.

### Related information

*ONTAP concepts*

## High Availability window

The High Availability window provides a pictorial representation of the high-availability (HA) state, interconnect status, and takeover or giveback status of all of the HA pairs in ONTAP. You can also manually initiate a takeover operation or giveback operation by using the High Availability window.

You can view details such as the takeover or giveback status and the interconnect status by clicking the HA pair image.

The color indicates the HA pair status:

- Green: Indicates that the HA pair and the interconnect are optimally configured and available for takeover or giveback.

  Green also indicates the takeover in progress state, giveback in progress state, and waiting for giveback state.
- Red: Indicates a downgraded state such as a takeover failure.

- Yellow: Indicates that the interconnect status is down.

When multiple HA pairs in a cluster are simultaneously involved in storage failover operations, the cluster status that is displayed is based on the status and severity of the HA pair. The following order of severity is considered while displaying the cluster status: takeover in progress, giveback in progress, waiting for giveback.

### Actions

You can perform tasks such as takeover or giveback based on the status of the nodes in the HA pair.

- Takeover *node_name*

  Enables you to perform a takeover operation when maintenance is required on the partner node.
- Giveback *node_name*

  Enables you to perform a giveback operation when the partner node that has been taken over is waiting for giveback or is in a partial giveback state.
- Enable or Disable automatic giveback

  Enables or disables the automatic giveback operation.

  **Note:** Automatic giveback is enabled by default.

### Command buttons

**Refresh**

Updates the information in the window.

**Note:** The information that is displayed in the High Availability window is automatically refreshed every 60 seconds.

### Related tasks

*Monitoring HA pairs* on page 36

You can use System Manager to monitor the node status and interconnect status of all of the high-availability (HA) pairs in a cluster. You can also verify whether takeover or giveback is enabled or has occurred, and view the reasons why takeover or giveback is not currently possible.

# Licenses

You can use System Manager to view, manage, or delete any software licenses installed on a cluster or node.

### Related information

*System administration*

## Deleting licenses

You can use the Licenses window in System Manager to delete any software license that is installed on a cluster or a node.

### Before you begin

The software license that you want to delete must not be used by any service or feature.

### Steps

1. Click **Configuration** > **Licenses**.
2. In the **Licenses** window, perform the appropriate action:

| If you want to... | Do this... |
|---|---|
| Delete a specific license package on a node or a master license | Click the **Details** tab. |
| Delete a specific license package across all of the nodes in the cluster | Click the **Packages** tab. |

3. Select the software license package that you want to delete, and then click **Delete**.

   You can delete only one license package at a time.
4. Select the confirmation check box, and then click **Delete**.

### Result

The software license is deleted from your storage system. The deleted license is also removed from the list of licenses in the Licenses window.

### Related reference

*Licenses window* on page 81

Your storage system arrives from the factory with preinstalled software. If you want to add or remove a software license after you receive the storage system, you can use the Licenses window.

## License types and entitlement risk

Understanding the various license types and the associated entitlement risk helps you manage the risk that is associated with the licenses in a cluster.

### License types

A package can have one or more of the following types of licenses installed in the cluster:

- Node-locked license or standard license

  A node-locked license is issued for a node with a specific system serial number (also known as a *controller serial number*). This license is valid only for the node that has the matching serial number.

  Installing a node-locked license entitles a node to the licensed functionality. For the cluster to use the licensed functionality, at least one node must be licensed for the functionality. It might be out of compliance to use the licensed functionality on a node that does not have an entitlement for the functionality.

  ONTAP 8.2 and later releases treat a license that was installed prior to Data ONTAP 8.2 as a standard license. Therefore, in ONTAP 8.2 and later releases, all of the nodes in the cluster automatically have the standard license for the package that the previously licensed functionality is part of.
- Master or site license

  A master or site license is not tied to a specific system serial number. When you install a site license, all of the nodes in the cluster are entitled to the licensed functionality.

  If your cluster has a master license and you remove a node from the cluster, the node does not carry the site license with it, and it is no longer entitled to the licensed functionality. If you add a node to a cluster that has a master license, the node is automatically entitled to the functionality that is granted by the site license.
- Demo or temporary license

  A demo or temporary license expires after a certain period of time. This license enables you to try certain software functionality without purchasing an entitlement. A temporary license is a cluster-wide license, and is not tied to a specific serial number of a node.

  If your cluster has a temporary license for a package and you remove a node from the cluster, the node does not carry the evaluation license with it.

- Capacity license (ONTAP Select and FabricPool only)
  An ONTAP Select instance is licensed according to the amount of data that the user wants to manage. For example, the user may buy a 10 TB capacity license to enable ONTAP Select to manage up to 10 TB of data. If more storage capacity is attached to the system than ONTAP Select is licensed to manage, ONTAP Select will not operate. By default, the maximum storage capacity that can be attached to an ONTAP Select instance is 2 TB until a capacity license (for example, a 5 TB capacity license, a 10 TB capacity license, and so on) is purchased and installed.
  Starting with ONTAP 9.2, FabricPool requires a capacity license to be used with a third-party storage tier (for example, AWS). The FabricPool capacity license defines the amount of data that can be stored in the external tiered storage.

### Entitlement risk

An entitlement risk arises because of the non-uniform installation of a node-locked license. If the node-locked license is installed on all the nodes, there is no entitlement risk.

The entitlement risk level can be high risk, medium risk, no risk, or unknown risk depending on certain conditions:

- High risk
  - If there is usage on a particular node, but the node-locked license is not installed on that node
  - If the demo license that was installed on the cluster expires, and there is usage on any node

  **Note:** If a master license is installed on a cluster, the entitlement risk is never high.
- Medium risk
  - If there is usage on the nodes, and only the site license is installed on the cluster
  - If there is usage on the nodes, but the node-locked license is not installed on these nodes
  - If the site license is not installed, and the node-locked license is non-uniformly installed on the nodes in a cluster
- No risk
  There is no entitlement risk if a node-locked license is installed on all the nodes, irrespective of the usage.
- Unknown
  The risk is unknown if the API is sometimes unable to retrieve the data related to entitlement risk that is associated with the cluster or the nodes in the cluster.

## Licenses window

Your storage system arrives from the factory with preinstalled software. If you want to add or remove a software license after you receive the storage system, you can use the Licenses window.

**Note:** System Manager does not monitor evaluation licenses and does not provide any warning when an evaluation license is nearing expiry. An evaluation license is a temporary license that expires after a certain period of time.

- *Command buttons*
- *Packages tab*
- *Packages details area*
- *Details tab*

### Command buttons

**Add**

Opens the Add License window, which enables you to add new software licenses.

**Delete**

Deletes the software license that you select from the software license list.

**Refresh**

Updates the information in the window.

### Packages tab

Displays information about the license packages that are installed on your storage system.

**Package**

Displays the name of the license package.

**Entitlement Risk**

Indicates the level of risk as a result of license entitlement issues for a cluster. The entitlement risk level can be high risk ( ), medium risk ( ), no risk ( ), unknown ( ), or unlicensed (-).

**Description**

Displays the level of risk as a result of license entitlement issues for a cluster.

### License Package details area

The area below the license packages list displays additional information about the selected license package. This area includes information about the cluster or node on which the license is installed, the serial number of the license, usage in the previous week, whether the license is installed, the expiration date of the license, and whether the license is a legacy one.

### Details tab

Displays additional information about the license packages that are installed on your storage system.

**Package**

Displays the name of the license package.

**Cluster/Node**

Displays the cluster or node on which the license package is installed.

**Serial Number**

Displays the serial number of the license package that is installed on the cluster or node.

**Type**

Displays the type of the license package, which can be the following:

- Temporary: Specifies that the license is a temporary license, which is valid only during the demonstration period.
- Master: Specifies that the license is a master license, which is installed on all the nodes in the cluster.
- Node Locked: Specifies that the license is a node-locked license, which is installed on a single node in the cluster.
- Capacity:
  - For ONTAP Select, specifies that the license is a capacity license, which defines the total amount of data capacity that the instance is licensed to manage.
  - For FabricPool, specifies that the license is a capacity license, which defines the amount of data that can be managed in the attached third-party storage (for example, AWS).

**State**

Displays the state of the license package, which can be the following:

- Evaluation: Specifies that the installed license is an evaluation license.
- Installed: Specifies that the installed license is a valid purchased license.
- Warning: Specifies that the installed license is a valid purchased license and is approaching maximum capacity.
- Enforcement: Specifies that the installed license is a valid purchased license and has exceeded the expiry date.
- Waiting for License: Specifies that the license has not yet been installed.

**Legacy**

Displays whether the license is a legacy license.

**Maximum Capacity**

- For ONTAP Select, displays the maximum amount of storage that can be attached to the ONTAP Select instance.
- For FabricPool, displays the maximum amount of third-party object store storage that can be used as external tiered storage.

**Current Capacity**

- For ONTAP Select, displays the total amount of storage that is currently attached to the ONTAP Select instance.
- For FabricPool, displays the total amount of third-party object store storage that is currently used as external tiered storage.

**Expiration Date**

Displays the expiration date of the software license package.

### Related tasks

*Adding licenses* on page 34
If your storage system software was installed at the factory, System Manager automatically adds the software to its list of licenses. If the software was not installed at the factory or if you want to add additional software licenses, you can add the software license by using System Manager.

*Deleting licenses* on page 79
You can use the Licenses window in System Manager to delete any software license that is installed on a cluster or a node.

*Creating a cluster* on page 23
You can use OnCommand System Manager to create and set up a cluster in your data center.

## Cluster Expansion

You can use System Manager to increase the size and capabilities of your storage by adding compatible nodes to the cluster and configuring the node network details. You can also view the summary of the nodes.

When you log in to System Manager, System Manager automatically detects compatible nodes that have been cabled but have not been added to the cluster and prompts you to add the nodes. You can add compatible nodes as and when System Manager detects the nodes or you can manually add the nodes at a later time.

### Steps

1. *Adding nodes to a cluster* on page 84
2. *Configuring the network details of the nodes* on page 84

## Adding nodes to a cluster

You can use System Manager to increase the size and capabilities of your storage system by adding nodes to an existing cluster.

### Before you begin

- New compatible nodes must be cabled to the cluster.
  Only the ports that are in the default broadcast domain will be listed in the Network window.
- All of the nodes in the cluster must be up and running.
- All of the nodes must be of the same version.

### Step

Add the new compatible nodes to the cluster:

| If you are... | Do this... |
|---|---|
| Not logged in to System Manager | **a.** Log in to System Manager.<br>**Note:** The new compatible nodes are automatically detected by System Manager at login. System Manager prompts you to add the new compatible nodes to the cluster.<br>**b.** Click **Add Nodes to Cluster**.<br>**c.** Modify the name of the nodes.<br>**d.** Specify the node licenses.<br>**e.** Click **Submit and Proceed**. |
| Logged in to System Manager | **a.** Click **Configurations** > **Cluster Expansion**.<br>System Manager searches for newly added nodes. If any warnings are displayed, you must fix them before proceeding. If new compatible nodes are discovered, proceed to the next step.<br>**b.** Modify the name of the nodes.<br>**c.** Specify the node licenses.<br>**d.** Click **Submit and Proceed**. |

## Configuring the network details of the nodes

You can use System Manager to configure the node management LIF and Service Processor settings for the newly added nodes.

### Before you begin

- Sufficient number of ports must be present in the default IPspace for LIF creation.
- All the ports must be up and running.

### Steps

1. Configure node management:
   a. Enter the IP address in the **IP Address** field.
   b. Select the port for node management in the **Port** field.
   c. Enter the netmask and gateway details.
2. Configure Service Processor settings:
   a. Select the **Override defaults** check box to override the default values.
   b. Enter the IP address, netmask, and gateway details.
3. Click **Submit and Proceed** to complete the network configuration of the nodes.

4. Verify the details of the nodes in the **Summary** page.

### After you finish

- If your cluster is protected, you should create the required number of intercluster LIFs in the newly added nodes to avoid partial peering and unhealthy protection.
- If SAN data protocols are enabled in your cluster, you should create the required number of SAN Data LIFs for serving data.

### Related tasks

*Creating network interfaces* on page 103

You can use System Manager to create a network interface or LIF to access data from storage virtual machines (SVMs), to manage SVMs, and to provide an interface for intercluster connectivity.

## Updating clusters

You can use System Manager to update a cluster or the individual nodes in a high-availability (HA) pair. To perform an update, you should select an ONTAP image, validate that your cluster or the individual nodes in the HA pair are ready for the update, and then perform the update.

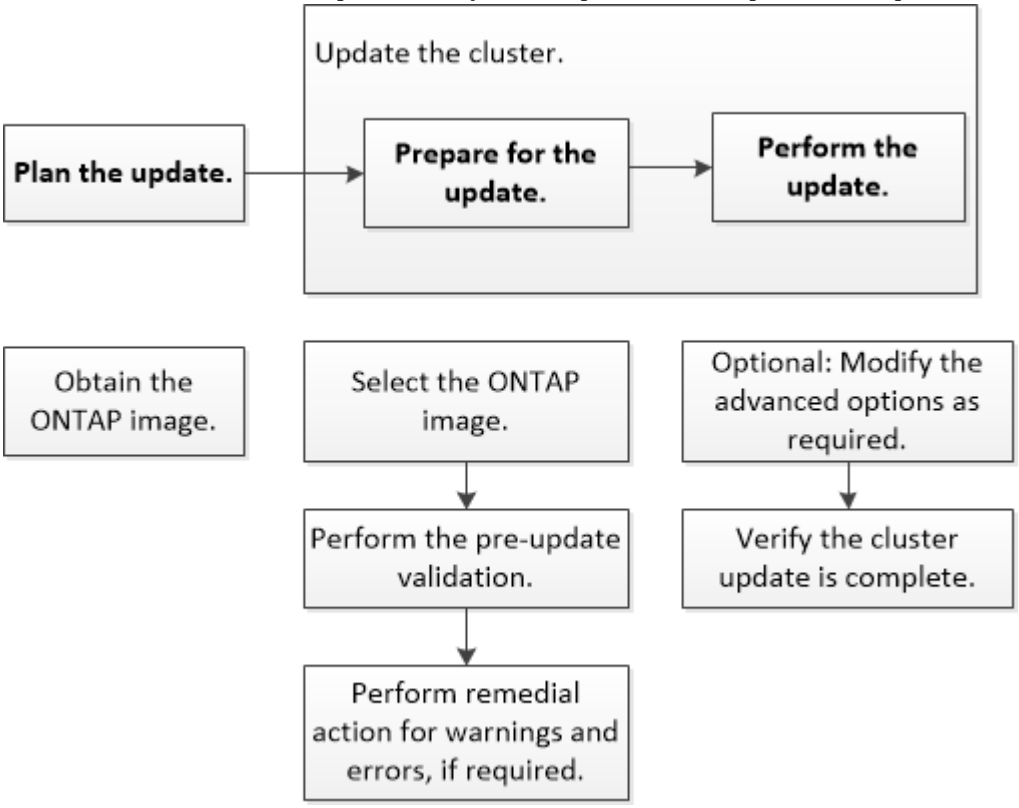


## Obtaining ONTAP software images

Beginning in ONTAP 9.4, you can copy the ONTAP software image from the NetApp Support Site to a local folder. This only applies to ONTAP 9.4 patch updates. For upgrades from ONTAP 9.3 or earlier, you must copy the ONTAP software image to an HTTP or FTP server on your network.

### About this task

To upgrade the cluster to the target release of ONTAP, you require access to software images. Software images, firmware version information, and the latest firmware for your platform model are available on the NetApp Support Site. You should note the following important information:

- Software images are specific to platform models.
  You must obtain the correct image for your cluster.
- Software images include the latest version of system firmware that was available when a given version of ONTAP was released.

### Steps

1. Locate the target ONTAP software in the **Software Downloads** area of the NetApp Support Site.
2. Copy the software image (for example, `94_q_image.tgz`) from the NetApp Support Site to the directory on the HTTP or FTP server from which the image will be served, or for ONTAP 9.4 patch updates, to a local folder.

## Updating single-node clusters disruptively

Starting with System Manager 9.4, you can update single-node clusters. Updating single-node clusters is disruptive, and client data will not be available while the update is in progress.

### Before you begin

- The clusters must be running ONTAP 9.4 or later.
- You must have copied the software image from the NetApp Support Site to an HTTP server on your network, to an FTP server on your network, or to your local system so that the nodes can access the image.
  *Obtaining Data ONTAP software images*

### About this task

If you try to perform other tasks from System Manager while updating the node that hosts the cluster management LIF, an error message might be displayed. You must wait for the update to finish before performing any operations.

### Steps

1. Click **Configuration** > **Cluster Update**.
2. In the **Cluster Update** tab, add a new image or select an available image.

| If you want to… | Then… |
|---|---|
| Add a new software image from the local client | **a.** Click **Add from Local Client**.<br>**b.** Search for the software image, and then click **Open**. |
| Add a new software image from the NetApp Support Site | **a.** Click **Add from Server**.<br>**b.** In the Add a New Software Image dialog box, enter the URL of the HTTP server or FTP server on which you have saved the image that was downloaded from the NetApp Support Site.<br>For anonymous FTP, you must specify the URL in the `ftp://anonymous@ftpserver` format.<br>**c.** Click **Add**. |
| Select an available image | Choose one of the listed images. |

3. Click **Validate** to run the pre-update validation checks to verify whether the cluster is ready for an update.

   The validation operation checks the cluster components to validate that the update can be completed, and then displays any errors or warnings. It also displays any required remedial action that you must perform before updating the software.

   **Important:** You must perform all of the required remedial actions for the errors before proceeding with the update. Although you can ignore the remedial actions for the warnings, the best practice is to perform all of the remedial actions before proceeding with the update.

4. Click **Next**.
5. Click **Update**.

   Validation is performed again.

   - When the validation is complete, a table displays any errors and warnings, along with any required remedial actions to be taken before proceeding.
   - If the validation is completed with warnings, you can choose to select the **Continue update with warnings** checkbox, and then click **Continue**.

   When the validation is complete and the update is in progress, the update might be paused because of errors. You can click the error message to view the details, and then perform the remedial actions before resuming the update.

   After the update is completed successfully, the node reboots, and you are redirected to the System Manager login page. If the node takes a long time to reboot, you must refresh your browser.

6. Log in to System Manager and verify that the cluster is successfully updated to the selected version by clicking **Cluster** > **Cluster Update** > **Update History**, and then viewing the details.

## Updating a cluster nondisruptively

You can use System Manager to update a cluster or individual nodes in high-availability (HA) pairs that are running ONTAP 9.3 or later to a specific version of ONTAP software without disrupting access to client data.

### Before you begin

- All of the nodes must be in HA pairs.
- All of the nodes must be healthy.
- You must have copied the software image from the NetApp Support Site to an HTTP server or FTP server on your network so that the nodes can access the image.

  *Obtaining Data ONTAP software images*

### About this task

- If you try to perform other tasks from System Manager while updating the node that hosts the cluster management LIF, an error message might be displayed.

  You must wait for the update to finish before performing any operations.
- A rolling update is performed for clusters with fewer than eight nodes, and a batch update is performed for clusters with more than eight nodes.

  In a rolling update, the nodes in the cluster are updated one at a time. In a batch update, multiple nodes are updated in parallel.
- You can nondisruptively update ONTAP software from one long-term service (LTS) release to the next LTS release (LTS+1).

  For example, if ONTAP 9.1 and ONTAP 9.3 are LTS releases, you can nondisruptively update your cluster from ONTAP 9.1 to ONTAP 9.3.

### Steps

1. Click **Configuration** > **Cluster Update**.
2. In the **Cluster Update** tab, add a new image or select an available image.

| If you want to… | Then… |
|---|---|
| Add a new software image from the local client | **a.** Click **Add from Local Client**.<br>**b.** Search for the software image, and then click **Open**. |
| Add a new software image from the NetApp Support Site | **a.** Click **Add from Server**.<br>**b.** In the Add a New Software Image dialog box, enter the URL of the HTTP server or FTP server on which you have saved the image that was downloaded from the NetApp Support Site.<br>For anonymous FTP, you must specify the URL in the `ftp://anonymous@ftpserver` format.<br>**c.** Click **Add**. |
| Select an available image | Choose one of the listed images. |

3. Click **Validate** to run the pre-update validation checks to verify whether the cluster is ready for an update.

   The validation operation checks the cluster components to validate that the update can be completed nondisruptively, and then displays any errors or warnings. It also displays any required remedial action that you must perform before updating the software.

   **Important:** You must perform all of the required remedial actions for the errors before proceeding with the update. Although you can ignore the remedial actions for the warnings, the best practice is to perform all of the remedial actions before proceeding with the update.

4. Click **Next**.
5. Click **Update**.

   Validation is performed again.

   - When the validation is complete, a table displays any errors and warnings, along with any required remedial actions to be taken before proceeding.
   - If the validation is completed with warnings, you can choose to select the **Continue update with warnings** checkbox, and then click **Continue**.

   When the validation is complete and the update is in progress, the update might be paused because of errors. You can click the error message to view the details, and then perform the remedial actions before resuming the update.

   After the update is completed successfully, the node reboots, and you are redirected to the System Manager login page. If the node takes a long time to reboot, you must refresh your browser.

6. Log in to System Manager and verify that the cluster is successfully updated to the selected version by clicking **Cluster** > **Cluster Update** > **Update History**, and then viewing the details.

### Related concepts

*How to update a cluster nondisruptively* on page 89

You can use System Manager to update a cluster nondisruptively to a specific ONTAP version. In a nondisruptive update, you have to select an ONTAP image, validate that your cluster is ready for the update, and then perform the update.

## How to update a cluster nondisruptively

You can use System Manager to update a cluster nondisruptively to a specific ONTAP version. In a nondisruptive update, you have to select an ONTAP image, validate that your cluster is ready for the update, and then perform the update.

During a nondisruptive update, the cluster remains online and continues to serve data.

### Planning and preparing for the update

As part of planning and preparing for the cluster update, you have to obtain the version of the ONTAP image to which you want to update the cluster from the NetApp Support Site, select the software image, and then perform a validation. The pre-update validation verifies whether the cluster is ready for an update to the selected version.

If the validation finishes with errors and warnings, you have to resolve the errors and warnings by performing the required remedial actions, and then verify that the cluster components are ready for the update. For example, during the pre-update validation, if a warning is displayed that offline aggregates are present in the cluster, you must navigate to the aggregate page, and then change the status of all of the offline aggregates to online.

### Performing an update

When you update the cluster, either the entire cluster is updated or the nodes in a high-availability (HA) pair are updated. As part of the update, the pre-update validation is run again to verify that the cluster is ready for the update.

A rolling update or batch update is performed, depending on the number of nodes in the cluster.

**Rolling update**

One of the nodes is taken offline and is updated while the partner node takes over the storage of that node.

A rolling update is performed for a cluster that consists of two or more nodes. This is the only update method for clusters with less than eight nodes.

**Batch update**

The cluster is separated into two batches, each of which contains multiple HA pairs.

A batch update is performed for a cluster that consists of eight or more nodes. In such clusters, you can perform either a batch update or a rolling update. This is the default update method for clusters with eight or more nodes.

### Related tasks

*Updating a cluster nondisruptively* on page 87

You can use System Manager to update a cluster or individual nodes in high-availability (HA) pairs that are running ONTAP 9.3 or later to a specific version of ONTAP software without disrupting access to client data.

## Cluster Update window

You can use the Cluster Update window to perform an automated cluster update nondisruptively or you can update single-node clusters disruptively.

- *Tabs*
- *Cluster Update tab*
- *Update History tab*

### Tabs

**Cluster Update**

Enables you to perform an automated cluster update nondisruptively or you can update single-node clusters disruptively.

**Update History**

Displays the details of previous cluster updates.

### Cluster Update tab

The Cluster Update tab enables you perform an automated cluster update nondisruptively or you can update single-node clusters disruptively.

### Command buttons

**Refresh**

Updates the information in the window.

**Select**

You can select the version of the software image for the update.

- Cluster Version Details: Displays the current cluster version in use and the version details of the nodes or high-availability (HA) pairs.
- Available Software Images: Enables you to select an existing software image for the update.
  Alternatively, you can download a software image from the NetApp Support Site and add the image for the update.

**Validate**

You can view and validate the cluster against the software image version for the update. A pre-update validation checks whether the cluster is in a state that is ready for an update. If the validation is completed with errors, a table displays the status of the various components and the required corrective action for the errors.

You can perform the update only when the validation is completed successfully.

**Update**

You can update all of the nodes in the cluster or an HA pair in the cluster to the selected version of the software image. While the update is in progress, you can choose to pause the update, and you can then either cancel or resume the update.

If an error occurs, the update is paused and an error message is displayed with the remedial steps. You can choose to either resume the update after performing the remedial steps or cancel the update. You can view the table with the node name, uptime, state, and ONTAP version when the update is successfully completed.

### Update History tab

Displays details about the cluster update history.

### Update History list

**Image Version**

Specifies the version of the ONTAP image to which the node will be updated.

**Software Updates Installed on**

Specifies the type of disk on which the updates are installed.

**Status**

Specifies the status of the software image update (whether the update is successful or cancelled).

**Start Time**

Specifies the time when the update was started.

**Completion Time**

Specifies the time when the update was completed.

This field is hidden by default.

**Time Taken for the Update**

Specifies the time taken for the update to finish.

**Previous Version**

Specifies the ONTAP version of the node before the update.

**Updated Version**

Specifies the ONTAP version of the node after the update.

# Date and time settings of a cluster

You can use System Manager to manage the date and time settings of a cluster.

### Related information

*System administration*

## Date and Time window

The Date and Time window enables you to view the current date and time settings for your storage system and to modify the settings when required.

### Command buttons

**Edit**

Opens the Edit Date and Time dialog box, which enables you to edit the time servers.

**Refresh**

Updates the information in the window.

### Details area

The details area displays information about the date, time, time zone, NTP service, and time servers for your storage system.

### Related tasks

*Setting the time zone for a cluster* on page 35

*Setting up a network when an IP address range is disabled* on page 25

You can set up a network by disabling an IP address range and entering individual IP addresses for cluster management, node management, and service provider networks.

# SNMP

You can use System Manager to configure SNMP to monitor SVMs in your cluster.

### Related information

*Network management*

## Enabling or disabling SNMP

You can enable or disable SNMP on your clusters by using System Manager. SNMP enables you to monitor the storage virtual machines (SVMs) in a cluster to avoid issues before they can occur and to prevent issues from occurring.

### Steps

1. Click ⚙.
2. In the **Setup** pane, click **SNMP**.
3. In the **SNMP** window, click either **Enable** or **Disable**.

## Setting SNMP information

You can use the Edit SNMP Settings dialog box in System Manager to update information about the storage system location and contact personnel, and to specify the SNMP communities of your system.

### About this task

System Manager uses the SNMP protocols SNMPv1 and SNMPv2c and an SNMP community to discover storage systems.

### Steps

1. Click ⚙.
2. In the **Setup** pane, click **SNMP**.
3. Click **Edit**.
4. In the **General** tab, specify the contact personnel information and location information for the storage system, and the SNMP communities.

   The community name can be of 32 characters and must not contain the following special characters: , / : " ' |.
5. Click **OK**.
6. Verify the changes that you made to the SNMP settings in the **SNMP** window.

### Related reference

*SNMP window* on page 93

The SNMP window enables you to view the current SNMP settings for your system. You can also change your system's SNMP settings, enable SNMP protocols, and add trap hosts.

## Enabling or disabling SNMP traps

SNMP traps enable you to monitor the health and state of the various components of your storage system. You can use the Edit SNMP Settings dialog box in System Manager to enable or disable SNMP traps on your storage system.

### About this task

Although SNMP is enabled by default, SNMP traps are disabled by default.

### Steps

1. Click ⚙.
2. In the **Setup** pane, click **SNMP**.
3. In the **SNMP** window, click **Edit**.
4. In the **Edit SNMP Settings** dialog box, select the **Trap hosts** tab, and then select or clear the **Enable traps** check box to enable or disable SNMP traps, respectively.

5. If you enable SNMP traps, add the host name or IP address of the hosts to which the traps are sent.
6. Click **OK**.

### Related reference

*SNMP window* on page 93

The SNMP window enables you to view the current SNMP settings for your system. You can also change your system's SNMP settings, enable SNMP protocols, and add trap hosts.

## Testing the trap host configuration

You can use System Manager to test whether you have configured the trap host settings correctly.

### Steps

1. Click ⚙.
2. In the **Setup** pane, click **SNMP**.
3. In the **SNMP** window, click **Test Trap Host**.
4. Click **OK**.

## SNMP window

The SNMP window enables you to view the current SNMP settings for your system. You can also change your system's SNMP settings, enable SNMP protocols, and add trap hosts.

### Command buttons

**Enable/Disable**

Enables or disables SNMP.

**Edit**

Opens the Edit SNMP Settings dialog box, which enables you to specify the SNMP communities for your storage system and enable or disable traps.

**Test Trap Host**

Sends a test trap to all the configured hosts to check whether the test trap reaches all the hosts and whether the configurations for SNMP are set correctly.

**Refresh**

Updates the information in the window.

### Details

The details area displays the following information about the SNMP server and host traps for your storage system:

**SNMP**

Displays whether SNMP is enabled or not.

**Traps**

Displays if SNMP traps are enabled or not.

**Location**

Displays the address of the SNMP server.

**Contact**

Displays the contact details for the SNMP server.

**Trap host IP Address**

Displays the IP addresses of the trap host.

**Community Names**

Displays the community name of the SNMP server.

**Security Names**

Displays the security style for the SNMP server.

### Related tasks

*Setting SNMP information* on page 92
You can use the Edit SNMP Settings dialog box in System Manager to update information about the storage system location and contact personnel, and to specify the SNMP communities of your system.

*Enabling or disabling SNMP traps* on page 92
SNMP traps enable you to monitor the health and state of the various components of your storage system. You can use the Edit SNMP Settings dialog box in System Manager to enable or disable SNMP traps on your storage system.

# LDAP

You can use System Manager to configure an LDAP server that centrally maintains user information.

### Related tasks

*Adding an LDAP client configuration* on page 284
You can use System Manager to add an LDAP client configuration at the cluster level or the storage virtual machine (SVM) level if you want to use LDAP services. You must first set up an LDAP client to use LDAP services.

*Deleting an LDAP client configuration* on page 285
You can use System Manager to delete an LDAP client configuration at the cluster level or the storage virtual machine (SVM) level.

*Editing an LDAP client configuration* on page 285
You can use System Manager to edit an LDAP client configuration at the cluster level or the storage virtual machine (SVM) level.

## Viewing the LDAP client configuration

You can use System Manager to view the LDAP clients that are configured for a storage virtual machine (SVM) in a cluster.

### Steps

1. Click ⚙.
2. In the **Setup** pane, click **LDAP**.
   The list of LDAP clients are displayed in the LDAP window.

## Using LDAP services

An LDAP server enables you to centrally maintain user information. If you store your user database on an LDAP server in your environment, you can configure your storage virtual machine (SVM) to look up user information in your existing LDAP database.

### About this task

ONTAP supports LDAP for user authentication, file access authorization, and user lookup and mapping services between NFS and CIFS.

## LDAP window

You can use the LDAP window to view LDAP clients for user authentication, file access authorization, and user search, and to map services between NFS and CIFS at the cluster level.

### Command buttons

**Add**

Opens the Create LDAP Client dialog box, which enables you to create and configure LDAP clients.

**Edit**

Opens the Edit LDAP Client dialog box, which enables you to edit LDAP client configurations. You can also edit active LDAP clients.

**Delete**

Opens the Delete LDAP Client(s) dialog box, which enables you to delete LDAP client configurations. You can also delete an active LDAP client.

**Refresh**

Updates the information in the window.

### LDAP client list

Displays (in tabular format) details about LDAP clients.

**LDAP Client Configuration**

Displays the name of the LDAP client configuration that you specified.

**Storage Virtual Machine**

Displays the name of the storage virtual machine (SVM) for each LDAP client configuration.

**Schema**

Displays the schema for each LDAP client.

**Minimum Bind Level**

Displays the minimum bind level for each LDAP client.

**Active Directory Domain**

Displays the Active Directory domain for each LDAP client configuration.

**LDAP Servers**

Displays the LDAP server for each LDAP client configuration.

**Preferred Active Directory Servers**

Displays the preferred Active Directory server for each LDAP client configuration.

# Users

You can use System Manager to add, edit, and manage a cluster user account, and specify a login user method to access the storage system.

## Adding a cluster user account

You can use System Manager to add a cluster user account and to specify a user login method for accessing the storage system.

### About this task

In clusters on which SAML authentication is enabled, for a particular application, you can add either SAML authentication or password-based authentication, or you can add both types of authentication.

### Steps

1. Click ⚙.
2. In the **Management** pane, click **Users**.
3. Click **Add**.
4. Type a user name for the new user.
5. Type a password for the user to connect to the storage system, and then confirm the password.
6. Add one or more user login methods, and then click **Add**.

## Editing a cluster user account

You can use System Manager to edit a cluster user account by modifying the user login methods for accessing the storage system.

### Steps

1. Click ⚙.
2. In the **Management** pane, click **Users**.
3. In the **Users** window, select the user account that you want to modify, and then click **Edit**.
4. In the **Modify User** dialog box, modify the user login methods, and then click **Modify**.

## Changing passwords for cluster user accounts

You can use System Manager to reset the password for a cluster user account.

### Steps

1. Click ⚙.
2. In the **Management** pane, click **Users**.
3. Select the user account for which you want to modify the password, and then click **Change Password**.
4. In the **Change Password** dialog box, type the new password, confirm the new password, and then click **Change**.

## Locking or unlocking cluster user accounts

You can use System Manager to lock or unlock cluster user accounts.

### Steps

1. Click ⚙.
2. In the **Management** pane, click **Users**.
3. Select the user account for which you want to modify the status, and click either **Lock** or **Unlock**.

## User accounts (cluster administrators only)

You can create, modify, lock, unlock, or delete a cluster user account, reset a user's password, or display information about all user accounts.

You can manage cluster user accounts in the following ways:

- Creating a login method for a user by specifying the user's account name, the access method, the authentication method, and, optionally, the access-control role that the user is assigned
- Displaying users' login information, such as the account name, allowed access method, authentication method, access-control role, and account status
- Modifying the access-control role that is associated with a user’s login method

  **Note:** It is best to use a single role for all the access and authentication methods of a user account.
- Deleting a user's login method, such as the access method or the authentication method
- Changing the password for a user account
- Locking a user account to prevent the user from accessing the system
- Unlocking a previously locked user account to enable the user to access the system again

## Roles

You can use an access-control role to control the level of access a user has to the system. In addition to using the predefined roles, you can create new access-control roles, modify them, delete them, or specify account restrictions for the users of a role.

## Users window

You can use the Users window to manage user accounts, to reset the password of a user, and to view information about all of the user accounts.

### Command buttons

**Add**

Opens the Add User dialog box, which enables you to add user accounts.

**Edit**

Opens the Modify User dialog box, which enables you to modify user login methods.

**Note:** It is a best practice to use a single role for all of the access and authentication methods of a user account.

**Delete**

Enables you to delete a selected user account.

**Change Password**

Opens the Change Password dialog box, which enables you to reset a selected user's password.

**Lock**

Locks the user account.

**Refresh**

Updates the information in the window.

### Users list

The area below the users list displays detailed information about the selected user.

**User**

Displays the name of the user account.

**Account Locked**

Displays whether the user account is locked.

### User Login Methods area

**Application**

Displays the access method that a user can use to access the storage system. The supported access methods include the following:

- System console (console)
- HTTP(S) (http)
- ONTAP API (ontapi)
- Service Processor (service-processor)
- SSH (ssh)

**Authentication**

Displays the default supported authentication method, which is "password".

**Role**

Displays the role of a selected user.

# Roles

You can use System Manager to create access-controlled user roles.

### Related information

*Administrator authentication and RBAC*

## Adding roles

You can use System Manager to add an access-control role and to specify the command or command directory that users of the role can access. You can also control the level of access that the role has to the command or command directory, and you can specify a query that applies to the command or command directory.

### Steps

1. Click ⚙.
2. In the **Management** pane, click **Roles**.
3. In the **Roles** window, click **Add**.
4. In the **Add Role** dialog box, type the role name and add the role attributes.
5. Click **Add**.

## Editing roles

You can use System Manager to modify an access-control role's access to a command or command directory and to restrict a user's access to only a specified set of commands. You can also remove a role's access to the default command directory.

### Steps

1. Click ⚙.
2. In the **Management** pane, click **Roles**.
3. In the **Roles** window, select the role that you want to modify, and then click **Edit**.
4. In the **Edit Role** dialog box, modify the role attributes, and then click **Modify**.
5. Verify the changes that you made in the **Roles** window.

## Roles and permissions

The cluster administrator can restrict a user's access to only a specified set of commands by creating a restricted access-control role and then assigning the role to a user.

You can manage access-control roles in the following ways:

- By creating an access-control role, and then specifying the command or command directory that the role's users can access.
- By controlling the level of access that the role has for the command or command directory, and then specifying a query that applies to the command or command directory.
- By modifying an access-control role's access to a command or command directory.
- By displaying information about access-control roles, such as the role name, the command or command directory that a role can access, the access level, and the query.
- By deleting an access-control role.
- By restricting a user's access to only a specified set of commands.
- By displaying ONTAP APIs and their corresponding command-line interface (CLI) commands.

## Roles window

You can use the Roles window to manage the roles that are associated with user accounts.

### Command buttons

**Add**

Opens the Add Role dialog box, which enables you to create an access-control role and specify the command or command directory that the role's users can access.

**Edit**

Opens the Edit Role dialog box, which enables you to add or modify role attributes.

**Refresh**

Updates the information in the window.

### Roles list

The roles list provides a list of roles that are available to be assigned to users.

### Role Attributes area

The details area displays the role attributes, such as the command or command directory that the selected role can access, the access level, and the query that applies to the command or command directory.

# Managing the network

You can use System Manager to manage the network of your storage system by creating and managing IPspaces, broadcast domains, subnets, network interfaces, Ethernet ports, and FC/FCoE adapters.

## IPspaces

You can use System Manager to create and manage IPspaces.

**Related information**

*Network management*

### Editing IPspaces

You can use System Manager to rename an existing IPspace.

#### About this task

- All IPspace names must be unique within a cluster and must not consist of names that are reserved by the system, such as local or localhost.
- The system-defined "Default" IPspace and "Cluster" IPspace cannot be modified.

#### Steps

1. Click **Network** > **IPspaces**.
2. Select the IPspace that you want to modify, and then click **Edit**.
3. In the **Edit IPspace** dialog box, specify a new name for the IPspace.
4. Click **Rename**.

### Deleting IPspaces

You can use System Manager to delete an IPspace when you no longer require the IPspace.

#### Before you begin

The IPspace that you want to delete must not be associated with any broadcast domains, network interfaces, peer relationships, or storage virtual machines (SVMs).

#### About this task

The system-defined "Default" IPspace and "Cluster" IPspace cannot be deleted.

#### Steps

1. Click **Network** > **IPspaces**.
2. Select the IPspace that you want to delete, and then click **Delete**.
3. Select the confirmation check box, and then click **Yes**.

## Broadcast domains

You can use System Manager to create and manage broadcast domains.

**Related information**

*Network management*

## Editing broadcast domain settings

You can use System Manager to modify the attributes of a broadcast domain such as the name, the MTU size, and the ports that are associated with the broadcast domain.

### About this task

- You must not modify the MTU size of the broadcast domain to which the management port e0M is assigned.
- You cannot use System Manager to edit broadcast domains in the cluster IPspace.
  You must use the command-line interface (CLI) instead.

### Steps

1. Click **Network** > **Broadcast Domains**.
2. Select the broadcast domain that you want to modify, and then click **Edit**.
3. In the **Edit Broadcast Domain** dialog box, modify the broadcast domain attributes as required.
4. Click **Save and Close**.

### Related reference

*Network window* on page 109

You can use the Network window to view the list of network components, such as subnets, network interfaces, Ethernet ports, broadcast domains, FC/FCoE and NVMe adapters, and IPspaces, and to create, edit, or delete these components in your storage system.

## Deleting broadcast domains

You can delete a broadcast domain by using System Manager when you no longer require the broadcast domain.

### Before you begin

No subnets must be associated with the broadcast domain that you want to delete.

### About this task

- When you delete a broadcast domain, the ports that are associated with the broadcast domain are assigned to the default IPspace, and the MTU settings of the ports are not changed.
- You cannot use System Manager to delete broadcast domains that are in the cluster IPspace.
  You must use the command-line interface (CLI) instead.

### Steps

1. Click **Network** > **Broadcast Domains**.
2. Select the broadcast domain that you want to delete, and then click **Delete**.
3. Select the confirmation check box, and then click **Delete**.

### Related reference

*Network window* on page 109

You can use the Network window to view the list of network components, such as subnets, network interfaces, Ethernet ports, broadcast domains, FC/FCoE and NVMe adapters, and IPspaces, and to create, edit, or delete these components in your storage system.

# Subnets

You can use System Manager to manage subnets.

## Editing subnet settings

You can use System Manager to modify subnet attributes such as the name, subnet address, range of IP addresses, and gateway address of the subnet.

### About this task

- You cannot use System Manager to edit subnets in the cluster IPspace.
  You must use the command-line interface (CLI) instead.
- Modifying the gateway address does not update the route.
  You must use the CLI to update the route.

### Steps

1. Click **Network** > **Subnets**.
2. Select the subnet that you want to modify, and then click **Edit**.
   You can modify the subnet even when the LIF in that subnet is still in use.
3. In the **Edit Subnet** dialog box, modify the subnet attributes as required.
4. Click **Save and Close**.

### Related reference

*Network window* on page 109

You can use the Network window to view the list of network components, such as subnets, network interfaces, Ethernet ports, broadcast domains, FC/FCoE and NVMe adapters, and IPspaces, and to create, edit, or delete these components in your storage system.

## Deleting subnets

You can use System Manager to delete a subnet when you no longer require the subnet and you want to reallocate the IP addresses that were assigned to the subnet.

### Before you begin

The subnet that you want to delete must not have any LIFs that are using the IP addresses from the subnet.

### About this task

You cannot use System Manager to delete subnets in the Cluster IPspace. You must use the command-line interface (CLI) instead.

### Steps

1. Click **Network** > **Subnets**.
2. Select the subnet that you want to delete, and then click **Delete**.
3. Select the confirmation check box, and then click **Delete**.

### Related reference

*Network window* on page 109

You can use the Network window to view the list of network components, such as subnets, network interfaces, Ethernet ports, broadcast domains, FC/FCoE and NVMe adapters, and IPspaces, and to create, edit, or delete these components in your storage system.

# Network interfaces

You can use System Manager to create and manage network interfaces.

**Related information**

*ONTAP concepts*

*Network management*

## Creating network interfaces

You can use System Manager to create a network interface or LIF to access data from storage virtual machines (SVMs), to manage SVMs, and to provide an interface for intercluster connectivity.

### Before you begin

The broadcast domain that is associated with the subnet must have allocated ports.

### About this task

- Dynamic DNS (DDNS) is enabled by default when a LIF is created.
  However, DDNS is disabled if you configure the LIF for intercluster communication using iSCSI, NVMe and FC/FCoE protocols, or for management access only.
- You can specify an IP address by using a subnet or without using a subnet.
- You cannot use System Manager to create a network interface if the ports are degraded.
  You must use the command-line interface (CLI) to create a network interface in such cases.
- To create NVMe-oF data LIF the SVM must already be set up, the NVMe service must already exist on the SVM and the NVMe-oF capable adapters should be available.
- NVMe protocol is enabled only if the selected SVM has the NVMe service configured.

### Steps

1. Click **Network** > **Network Interfaces**.
2. Click **Create**.
3. In the **Create Network Interface** dialog box, specify an interface name.
4. Specify an interface role:

| If you want to... | Then... |
|---|---|
| Associate the network interface with a data LIF | **a.** Select **Serves Data**.<br>**b.** Select the SVM for the network interface. |
| Associate the network interface with an intercluster LIF | **a.** Select **Intercluster Connectivity**.<br>**b.** Select the IPspace for the network interface. |

5. Select the appropriate protocols.
   The interface uses the selected protocols to access data from the SVM

   **Note:** If you select the NVMe protocol, the rest of the protocols are disabled. If NAS (CIFS and NFS) protocols are supported then they remain available. The NVMe transports field is displayed when you select the NVMe protocol and FC-NVMe is shown as the transport protocol.

   .

6. If you want to enable management access on the data LIF, select the **Enable Management Access** check box.

   You cannot enable management access for intercluster LIFs or LIFs with FC/FCoE, NVMe or iSCSI protocols.

7. Assign the IP address:

| If you want to... | Then... |
|---|---|
| Specify the IP address by using a subnet | **a.** Select **Using a subnet**.<br>**b.** In the Add Details dialog box, select the subnet from which the IP address must be assigned.<br>For intercluster LIF, only the subnets that are associated with the selected IPspace are displayed.<br>**c.** If you want to assign a specific IP address to the interface, select **Use a specific IP address**, and then type the IP address.<br>The IP address that you specify is added to the subnet if the IP address is not already present in the subnet range.<br>**d.** Click **OK**. |
| Specify the IP address manually without using a subnet | **a.** Select **Without a subnet**.<br>**b.** In the Add Details dialog box, perform the following steps:<br>**i.** Specify the IP address and the network mask or prefix.<br>**ii.** Optional: Specify the gateway.<br>**iii.** If you do not want to use the default value for the Destination field, specify a new destination value.<br>If you do not specify a destination value, the Destination field is populated with the default value based on the family of the IP address.<br>If a route does not exist, a new route is automatically created based on the gateway and destination.<br>**c.** Click **OK**. |

8. Select the required ports from the **Port** details area.
   - For data LIFs, the details area displays all of the ports from the broadcast domain that is associated with the IPspace of the SVM.
   - For intercluster LIFs, the details area displays all of the ports from the broadcast domain that is associated with the required IPspace.
   - The Port details area will display only NVMe capable adapters if the NVMe protocol is selected.
9. Optional: Select the **Dynamic DNS (DDNS)** check box to enable DDNS.
10. Click **Create**.

### Related tasks

*Configuring iSCSI protocol on SVMs* on page 46

You can configure the iSCSI protocol on a storage virtual machine (SVM) to provide block-level data access by using System Manager. You can create iSCSI LIFs and portsets and then add the LIFs to the portsets. LIFs are created on the most suitable adapters and are assigned to portsets to ensure data path redundancy.

*Configuring the network details of the nodes* on page 84

You can use System Manager to configure the node management LIF and Service Processor settings for the newly added nodes.

### Related reference

*Network window* on page 109
You can use the Network window to view the list of network components, such as subnets, network interfaces, Ethernet ports, broadcast domains, FC/FCoE and NVMe adapters, and IPspaces, and to create, edit, or delete these components in your storage system.

## Editing network interface settings

You can use System Manager to modify the network interface to enable management access for a data LIF.

### About this task

- You cannot modify the network settings of cluster LIFs, cluster management LIFs, or node management LIFs through System Manager.
- You cannot enable management access for an intercluster LIF.

### Steps

1. Click **Network** > **Network Interfaces**.
2. Select the interface that you want to modify, and then click **Edit**.
3. In the **Edit Network Interface** dialog box, modify the network interface settings as required.
4. Click **Save and Close**.

### Related reference

*Network window* on page 109
You can use the Network window to view the list of network components, such as subnets, network interfaces, Ethernet ports, broadcast domains, FC/FCoE and NVMe adapters, and IPspaces, and to create, edit, or delete these components in your storage system.

## Deleting network interfaces

You can use System Manager to delete a network interface to free the IP address of the interface and then use the IP address for a different purpose.

### Before you begin

The status of the network interface must be disabled.

### Steps

1. Click **Network** > **Network Interfaces**.
2. Select the interface that you want to delete, and then click **Delete**.
3. Select the confirmation check box, and then click **Delete**.

### Related reference

*Network window* on page 109

You can use the Network window to view the list of network components, such as subnets, network interfaces, Ethernet ports, broadcast domains, FC/FCoE and NVMe adapters, and IPspaces, and to create, edit, or delete these components in your storage system.

## Migrating a LIF

You can use System Manager to migrate a data LIF or a cluster management LIF to a different port on the same node or on a different node within the cluster if the source port is faulty or requires maintenance.

### Before you begin

The destination node and ports must be operational and must be able to access the same network as the source port.

### About this task

- If you are removing the NIC from the node, you must migrate the LIFs that are hosted on the ports belonging to the NIC to other ports in the cluster.
- You cannot migrate iSCSI LIFs or FC LIFs.

### Steps

1. Click **Network** > **Network Interfaces**.
2. Select the interface that you want to migrate, and then click **Migrate**.
3. In the **Migrate Interface** dialog box, select the destination port to which you want to migrate the LIF.
4. Optional: Select the **Migrate Permanently** check box if you want to set the destination port as the new home port for the LIF.
5. Click **Migrate**.

# Ethernet ports

You can use System Manager to create and manage Ethernet ports.

### Related information

*Network management*

*ONTAP concepts*

## Creating interface groups

You can use System Manager to create an interface group—single-mode, static multimode, or dynamic multimode (LACP)—to present a single interface to clients by combining the capabilities of the aggregated network ports.

### Before you begin

Free ports must be available that do not belong to any broadcast domain or interface group, or that host a VLAN.

### Steps

1. Click **Network** > **Ethernet Ports**.
2. Click **Create Interface Group**.
3. In the **Create Interface Group** dialog box, specify the following settings:
   - Name of the interface group
   - Node
   - Ports that you want to include in the interface group
   - Usage mode of the ports: single-mode, static multiple, or dynamic multimode (LACP)

- Network load distribution: IP-based, MAC address-based, sequential, or port
- Broadcast domain for the interface group, if required

4. Click **Create**.

### Related reference

*Network window* on page 109

You can use the Network window to view the list of network components, such as subnets, network interfaces, Ethernet ports, broadcast domains, FC/FCoE and NVMe adapters, and IPspaces, and to create, edit, or delete these components in your storage system.

## Creating VLAN interfaces

You can create a VLAN to maintain separate broadcast domains within the same network domain by using System Manager.

### Steps

1. Click **Network** > **Ethernet Ports**.
2. Click **Create VLAN**.
3. In the **Create VLAN** dialog box, select the node, the physical interface, and the broadcast domain (if required).

   The physical interface list includes only Ethernet ports and interface groups. The list does not display interfaces that are in another interface group or an existing VLAN.
4. Type a VLAN tag, and then click **Add**.

   You must add unique VLAN tags.
5. Click **Create**.

### Related reference

*Network window* on page 109

You can use the Network window to view the list of network components, such as subnets, network interfaces, Ethernet ports, broadcast domains, FC/FCoE and NVMe adapters, and IPspaces, and to create, edit, or delete these components in your storage system.

## Editing Ethernet port settings

You can edit Ethernet port settings such as the duplex mode and speed settings by using System Manager.

### Steps

1. Click **Network** > **Ethernet Ports**.
2. Select the physical port, and then click **Edit**.
3. In the **Edit Ethernet Port** dialog box, modify the duplex mode and speed settings to either `manual` or `automatic`.
4. Click **Edit**.

## Editing interface group settings

You can use System Manager to add ports to an interface group, to remove ports from an interface group, and to modify the usage mode and load distribution pattern of the ports in an interface group.

### About this task

You cannot modify the MTU settings of an interface group that is assigned to a broadcast domain.

### Steps

1. Click **Network** > **Ethernet Ports**.
2. Select an interface group, and then click **Edit**.
3. Modify the interface group settings as required, and then click **Save and Close**.

### Related reference

*Network window* on page 109
You can use the Network window to view the list of network components, such as subnets, network interfaces, Ethernet ports, broadcast domains, FC/FCoE and NVMe adapters, and IPspaces, and to create, edit, or delete these components in your storage system.

## Modifying the MTU size of a VLAN

If you want to modify the MTU size of a VLAN interface that is not part of a broadcast domain, you can use System Manager to change the size.

### About this task

You must not modify the MTU size of the management port e0M.

### Steps

1. Click **Network** > **Ethernet Ports**.
2. Select the VLAN that you want to modify, and then click **Edit**.
3. In the **Edit VLAN** dialog box, modify the MTU size as required, and then click **Save**.

## Deleting VLANs

You can delete VLANs that are configured on network ports by using System Manager. You might have to delete a VLAN before removing a NIC from its slot. When you delete a VLAN, the VLAN is automatically removed from all of the failover rules and groups that use the VLAN.

### Before you begin

No LIFs must be associated with the VLAN.

### Steps

1. Click **Network** > **Ethernet Ports**.
2. Select the VLAN that you want to delete, and then click **Delete**.
3. Select the confirmation check box, and then click **Delete**.

### Related reference

*Network window* on page 109
You can use the Network window to view the list of network components, such as subnets, network interfaces, Ethernet ports, broadcast domains, FC/FCoE and NVMe adapters, and IPspaces, and to create, edit, or delete these components in your storage system.

## Ports and adapters

Ports are grouped under nodes and the nodes are displayed based on the selected protocol category. For example, if the data is served using the FC protocol, then only the nodes with FCP

adapters are displayed. The hosted interface count helps you in choosing a port which is less loaded.

# FC/FCoE and NVMe adapters

You can use System Manager to manage FC/FCoE and NVMe adapters.

### Related information

*Network management*

## Editing the FC/FCoE and NVMe adapter speed settings

You can modify the FC/FCoE and NVMe adapter speed settings by using the Edit FC/FCoE and NVMe Adapter Settings dialog box in System Manager.

### Steps

1. Click **Network** > **FC/FCoE and NVMe Adapters**
2. Select the adapter that you want to edit, and then click **Edit**.
3. In the **Edit FC/FCoE and NVMe adapter settings** dialog box, set the adapter speed to **Manual** or **Automatic**, and then click **Edit**.

### Related reference

*Network window* on page 109

You can use the Network window to view the list of network components, such as subnets, network interfaces, Ethernet ports, broadcast domains, FC/FCoE and NVMe adapters, and IPspaces, and to create, edit, or delete these components in your storage system.

# Network window

You can use the Network window to view the list of network components, such as subnets, network interfaces, Ethernet ports, broadcast domains, FC/FCoE and NVMe adapters, and IPspaces, and to create, edit, or delete these components in your storage system.

- *Tabs*
- *Subnet tab*
- *Network Interfaces tab*
- *Ethernet Ports tab*
- *Broadcast Domain tab*
- *FC/FCoE and NVMe Adapters tab*
- *IPspaces tab*

### Tabs

**Subnet**

Enables you to view a list of subnets, and create, edit, or delete subnets from your storage system.

**Network Interfaces**

Enables you to view a list of network interfaces, create, edit, or delete interfaces from your storage system, migrate the LIFs, change the status of the interface, and send the interface back to the home port.

**Ethernet Ports**

Enables you to view and edit the ports of a cluster, and create, edit, or delete interface groups and VLAN ports.

**Broadcast Domains**

Enables you to view a list of broadcast domains, and create, edit, or delete domains from your storage system.

**FC/FCoE and NVMe Adapters**

Enables you to view the ports in a cluster, and edit the FC/FCoE and NVMe adapter settings.

**IPspaces**

Enables you to view a list of IPspaces and broadcast domains, and create, edit, or delete an IPspace from your storage system.

### Subnet tab

### Command buttons

**Create**

Opens the Create Subnet dialog box, which enables you to create new subnets that contain configuration information for creating a network interface.

**Edit**

Opens the Edit Subnet dialog box, which enables you to modify certain attributes of a subnet such as the name, subnet address, range of IP addresses, and gateway details.

**Delete**

Deletes the selected subnet.

**Refresh**

Updates the information in the window.

### Subnet list

**Name**

Specifies the name of the subnet.

**Subnet IP/Subnet mask**

Specifies the subnet address details.

**Gateway**

Specifies the IP address of the gateway.

**Available**

Specifies the number of IP addresses available in the subnet.

**Used**

Specifies the number of IP addresses used in the subnet.

**Total Count**

Specifies the total number of IP addresses (available and used) in the subnet.

**Broadcast domain**

Specifies the broadcast domain to which the subnet belongs.

**IPspace**

Specifies the IPspace to which the subnet belongs.

### Details area

The area below the subnet list displays detailed information about the selected subnet, including the subnet range and a graph showing the available, used, and total number of IP addresses.

### Limitations of the Network Interfaces tab

- For cluster LIFs and node management LIFs, you cannot use System Manager to perform the following actions:
  - Create, edit, delete, enable, or disable the LIFs
  - Migrate the LIFs or send the LIFs back to the home port
- For cluster management LIFs, you can use System Manager to migrate the LIFs, or send the LIFs back to the home port.

  However, you cannot create, edit, delete, enable, or disable the LIFs.
- For intercluster LIFs, you can use System Manager to create, edit, delete, enable, or disable the LIFs.

  However, you cannot migrate the LIFs, or send the LIFs back to the home port.
- You cannot create, edit, or delete network interfaces in the following configurations:
  - A MetroCluster configuration
  - SVMs configured for disaster recovery (DR).

### Command buttons

**Create**

Opens the Create Network Interface dialog box, which enables you to create network interfaces and intercluster LIFs to serve data and manage SVMs.

**Edit**

Opens the Edit Network Interface dialog box, which you can use to enable management access for a data LIF.

**Delete**

Deletes the selected network interface.

This button is enabled only if the data LIF is disabled.

**Status**

Open the drop-down menu, which provides the option to enable or disable the selected network interface.

**Migrate**

Enables you to migrate a data LIF or a cluster management LIF to a different port on the same node or a different node within the cluster.

**Send to Home**

Enables you to host the LIF back on its home port.

This command button is enabled only when the selected interface is hosted on a non-home port and when the home port is available.

This command button is disabled when any node in the cluster is down.

**Refresh**

Updates the information in the window.

### Interface list

You can move the pointer over the color-coded icon to view the operational status of the interface:

- Green specifies that the interface is enabled.
- Red specifies that the interface is disabled.

**Interface Name**

Specifies the name of the network interface.

**Storage Virtual Machine**

Specifies the SVM to which the interface belongs.

**IP Address/WWPN**

Specifies the IP address or worldwide port name (WWPN) of the interface.

**Current Port**

Specifies the name of the node and port on which the interface is hosted.

**Data Protocol Access**

Specifies the protocol used to access data.

**Management Access**

Specifies whether management access is enabled on the interface.

**Subnet**

Specifies the subnet to which the interface belongs.

**Role**

Specifies the operational role of the interface, which can be data, intercluster, cluster, cluster management, or node management.

### Details area

The area below the interface list displays detailed information about the selected interface: failover properties such as the home port, current port, speed of the ports, failover policy, failover group, and failover state, and general properties such as the administrative status, role, IPspace, broadcast domain, network mask, gateway, and DDNS status.

### Ethernet Ports tab

### Command buttons

**Create Interface Group**

Opens the Create Interface Group dialog box, which enables you create interface groups by choosing the ports, and determining the use of ports and network traffic distribution.

**Create VLAN**

Opens the Create VLAN dialog box, which enables you to create a VLAN by choosing an Ethernet port or an interface group, and adding VLAN tags.

**Edit**

Opens one of the following dialog boxes:

- Edit Ethernet Port dialog box: Enables you to modify Ethernet port settings.
- Edit VLAN dialog box: Enables you to modify VLAN settings.
- Edit Interface Group dialog box: Enables you to modify interface groups.

You can only edit VLANs that are not associated with a broadcast domain.

**Delete**

Opens one of the following dialog boxes:

- Delete VLAN dialog box: Enables you to delete a VLAN.
- Delete Interface Group dialog box: Enables you to delete an interface group.

**Refresh**

Updates the information in the window.

### Ports list

You can move the pointer over the color-coded icon to view the operational status of the port:

- Green specifies that the port is enabled.
- Red specifies that the port is disabled.

**Port**

Displays the port name of the physical port, VLAN port, or the interface group.

**Node**

Displays the node on which the physical interface is located.

**Broadcast Domain**

Displays the broadcast domain of the port.

**IPspace**

Displays the IPspace to which the port belongs.

**Type**

Displays the type of the interface such as interface group, physical interface, or VLAN.

### Details area

The area below the ports list displays detailed information about the port properties.

**Details tab**

Displays administrative details and operational details.

As part of the operational details, the tab displays the health status of the ports. The ports can be healthy or degraded. A degraded port is a port on which continuous network fluctuations occur, or a port that has no connectivity to any other ports in the same broadcast domain.

In addition, the tab also displays the interface name, SVM details, and IP address details of the network interfaces that are hosted on the selected port. It also indicates whether the interface is at the home port or not.

**Performance tab**

Displays performance metrics graphs of the ethernet ports, including error rate and throughput.

Changing the client time zone or the cluster time zone impacts the performance metrics graphs. You should refresh your browser to view the updated graphs.

## Broadcast Domain tab

### Command buttons

**Create**

Opens the Create Broadcast Domain dialog box, which enables you to create new broadcast domains to contain ports.

**Edit**

Opens the Edit Broadcast Domain dialog box, which enables you to modify the attributes of a broadcast domain, such as the name, MTU size, and associated ports.

**Delete**

Deletes the selected broadcast domain.

**Refresh**

Updates the information in the window.

### Broadcast domain list

**Broadcast Domain**

Specifies the name of the broadcast domain.

**MTU**

Specifies the MTU size.

**IPspace**

Specifies the IPspace.

**Combined Port Update Status**

Specifies the status of the port updates when you create or edit a broadcast domain. Any errors in the port updates are displayed in a separate window, which you can open by clicking the associated link.

### Details area

The area below the broadcast domain list displays all the ports in a broadcast domain. In a non-default IPspace, if a broadcast domain has ports with update errors, such ports are not displayed in the details area. You can move the pointer over the color-coded icon to view the operational status of the ports:

- Green specifies that the port is enabled.
- Red specifies that the port is disabled.

## FC/FCoE and NVMe Adapters tab

### Command buttons

**Edit**

Opens the Edit FC/FCoE and NVMe Settings dialog box, which enables you to modify the speed of the adapter.

**Status**

Enables you to bring the adapter online or take it offline.

**Refresh**

Updates the information in the window.

### FC/FCoE and NVMe adapters list

**WWNN**

Specifies the unique identifier of the FC/FCoE and NVMe adapter.

**Node Name**

Specifies the name of the node that is using the adapter.

**Slot**

Specifies the slot that is using the adapter.

**WWPN**

Specifies the FC worldwide port name (WWPN) of the adapter.

**Status**

Specifies whether the status of the adapter is online or offline.

**Speed**

Specifies whether the speed settings are automatic or manual.

### Details area

The area below the FC/FCoE and NVMe adapters list displays detailed information about the selected adapters.

**Details tab**

Displays adapter details such as the media type, port address, data link rate, connection status, operation status, fabric status, and the speed of the adapter.

**Performance tab**

Displays performance metrics graphs of the FC/FCoE and NVMe adapter, including IOPS and response time.

Changing the client time zone or the cluster time zone impacts the performance metrics graphs. You should refresh your browser to see the updated graphs.

### IPspaces tab

### Command buttons

**Create**

Opens the Create IPspace dialog box, which enables you to create a new IPspace.

**Edit**

Opens the Edit IPspace dialog box, which enables you to rename an existing IPspace.

**Delete**

Deletes the selected IPspace.

**Refresh**

Updates the information in the window.

### IPspaces list

**Name**

Specifies the name of the IPspace.

**Broadcast Domains**

Specifies the broadcast domain.

### Details area

The area below the IPspaces list displays the list of storage virtual machines (SVMs) in the selected IPspace.

#### Related tasks

*Creating network interfaces* on page 103
You can use System Manager to create a network interface or LIF to access data from storage virtual machines (SVMs), to manage SVMs, and to provide an interface for intercluster connectivity.

*Editing network interface settings* on page 105
You can use System Manager to modify the network interface to enable management access for a data LIF.

*Deleting network interfaces* on page 105
You can use System Manager to delete a network interface to free the IP address of the interface and then use the IP address for a different purpose.

*Creating subnets* on page 37

You can create a subnet by using System Manager to provide a logical subdivision of an IP network to pre-allocate the IP addresses. A subnet enables you to create interfaces more easily by specifying a subnet instead of an IP address and network mask values for each new interface.

*Editing subnet settings* on page 102
You can use System Manager to modify subnet attributes such as the name, subnet address, range of IP addresses, and gateway address of the subnet.

*Deleting subnets* on page 102
You can use System Manager to delete a subnet when you no longer require the subnet and you want to reallocate the IP addresses that were assigned to the subnet.

*Creating VLAN interfaces* on page 107
You can create a VLAN to maintain separate broadcast domains within the same network domain by using System Manager.

*Creating interface groups* on page 106
You can use System Manager to create an interface group—single-mode, static multimode, or dynamic multimode (LACP)—to present a single interface to clients by combining the capabilities of the aggregated network ports.

*Editing the FC/FCoE and NVMe adapter speed settings* on page 109
You can modify the FC/FCoE and NVMe adapter speed settings by using the Edit FC/FCoE and NVMe Adapter Settings dialog box in System Manager.

*Editing interface group settings* on page 107
You can use System Manager to add ports to an interface group, to remove ports from an interface group, and to modify the usage mode and load distribution pattern of the ports in an interface group.

*Deleting VLANs* on page 108
You can delete VLANs that are configured on network ports by using System Manager. You might have to delete a VLAN before removing a NIC from its slot. When you delete a VLAN, the VLAN is automatically removed from all of the failover rules and groups that use the VLAN.

*Creating broadcast domains* on page 36
You can create a broadcast domain by using System Manager to provide a logical division of a computer network. In a broadcast domain, all associated nodes can be reached through broadcast at the datalink layer.

*Editing broadcast domain settings* on page 101
You can use System Manager to modify the attributes of a broadcast domain such as the name, the MTU size, and the ports that are associated with the broadcast domain.

*Deleting broadcast domains* on page 101
You can delete a broadcast domain by using System Manager when you no longer require the broadcast domain.

*Setting up a network when an IP address range is disabled* on page 25
You can set up a network by disabling an IP address range and entering individual IP addresses for cluster management, node management, and service provider networks.

# Managing physical storage

You can use System Manager to manage physical storage such as aggregates, storage pools, disks, array LUNs, nodes, Flash Cache, events, system alerts, AutoSupport notifications, jobs, and Flash Pool statistics.

## Storage tiers

You can use System Manager to create aggregates to support the different security requirements, backup requirements, performance requirements, and data sharing requirements of your users.

### Related information

*Disk and aggregate management*

## Editing aggregates

You can use System Manager to change the aggregate name, RAID type, and RAID group size of an existing aggregate when required.

### Before you begin

For modifying the RAID type of an aggregate from RAID4 to RAID-DP, the aggregate must contain enough compatible spare disks, excluding the hot spares.

### About this task

- You cannot change the RAID group of ONTAP systems that support array LUNs.
  RAID0 is the only available option.
- You cannot change the RAID type of partitioned disks.
  RAID-DP is the only option that is available for partitioned disks.
- You cannot rename a SnapLock Compliance aggregate.
- If the aggregate consists of SSDs with storage pool, you can modify only the name of the aggregate.
- If the triple parity disk size is 10 TB, and the other disks are smaller than 10 TB in size, then you can select RAID-DP or RAID-TEC as the RAID type.
- If the triple parity disk size is 10 TB, and if even one of the other disks is larger than 10 TB in size, then RAID-TEC is the only available option for RAID type.

### Steps

1. Choose one of the following methods:
   - Click **Applications & Tiers** > **Storage Tiers**.
   - Click **Storage** > **Aggregates & Disks** > **Aggregates**.
2. Select the aggregate that you want to edit, and then click **Edit**.
3. In the **Edit Aggregate** dialog box, modify the aggregate name, the RAID type, and the RAID group size, as required.
4. Click **Save**.

### Related concepts

*What compatible spare disks are* on page 125

In System Manager, compatible spare disks are disks that match the properties of other disks in the aggregate. When you want to increase the size of an existing aggregate by adding HDDs (capacity

disks) or change the RAID type of an aggregate from RAID4 to RAID-DP, the aggregate must contain sufficient compatible spare disks.

### Related reference

*Aggregates window* on page 136
You can use the Aggregates window to create, display, and manage information about aggregates.

*Storage Tiers window* on page 129
You can use the Storage Tiers window to view cluster-wide space details and to add and view aggregate details.

## Deleting aggregates

You can use System Manager to delete aggregates when you no longer require the data in the aggregates. However, you cannot delete the root aggregate because it contains the root volume, which contains the system configuration information.

### Before you begin

- All the volumes and the associated storage virtual machines (SVMs) contained by the aggregate must be deleted.
- The aggregate must be offline.

### Steps

1. Choose one of the following methods:
   - Click **Applications & Tiers** > **Storage Tiers**.
   - Click **Storage** > **Aggregates & Disks** > **Aggregates**.
2. Select one or more aggregates that you want to delete, and then click **Delete**.
3. Select the confirmation check box, and then click **Delete**.

### Related reference

*Aggregates window* on page 136
You can use the Aggregates window to create, display, and manage information about aggregates.

*Storage Tiers window* on page 129
You can use the Storage Tiers window to view cluster-wide space details and to add and view aggregate details.

## Changing the RAID configuration when creating an aggregate

While creating a new aggregate, you can modify the default values of the RAID type and RAID group size options of the aggregate by using System Manager.

### About this task

If the disk type of the aggregate disks is FSAS or MSATA, and the disk size is equal to or larger than 10 TB, then RAID-TEC is the only option available for RAID type.

### Steps

1. Choose one of the following methods:
   - Click **Applications & Tiers** > **Storage Tiers**.
   - Click **Storage** > **Aggregates & Disks** > **Aggregates**.
2. In the **Storage Tiers** window, click **Add Aggregate**.
3. In the **Create Aggregate** dialog box, perform the following steps:
   a. Click **Change**.

b. In the **Change RAID Configuration** dialog box, specify the RAID type and RAID group size.

   RAID-DP is the only supported RAID type for shared disks.

   The recommended RAID group size is 12 disks through 20 disks for HDDs, and 20 disks through 28 disks for SSDs.

c. Click **Save**.

## Provisioning cache by adding SSDs

You can use System Manager to add SSDs as either storage pools or dedicated SSDs to provision cache. By adding SSDs, you can convert a non-root aggregate or a root aggregate that does not contain partitioned disks to a Flash Pool aggregate, or increase the cache size of an existing Flash Pool aggregate.

### About this task

- The added SSD cache does not add to the size of the aggregate, and you can add an SSD RAID group to an aggregate even when it is at the maximum size.
- You cannot use partitioned SSDs when you add cache by using System Manager.

### Related concepts

*How storage pool works* on page 143

A *storage pool* is a collection of SSDs. You can combine SSDs to create a storage pool, which enables you to share the SSDs and SSD spares across multiple Flash Pool aggregates, at the same time.

## Provisioning cache to aggregates by adding SSDs

You can use System Manager to add storage pools or dedicated SSDs to provision cache by converting an existing non-root HDD aggregate or a root aggregate that does not contain partitioned disks to a Flash Pool aggregate.

### Before you begin

- The aggregate must be online.
- There must be sufficient spare SSDs or allocation units in the storage pool that can be assigned as cache disks.
- All of the nodes in the cluster must be running ONTAP 8.3 or later.

  If the cluster is in a mixed-version state, you can use the command-line interface to create a Flash Pool aggregate and then provide SSD cache.
- You must have identified a valid 64-bit non-root aggregate composed of HDDs that can be converted to a Flash Pool aggregate.
- The aggregate must not contain any array LUNs.

### About this task

You must be aware of platform-specific and workload-specific best practices for Flash Pool aggregate SSD tier size and configuration.

### Steps

1. Choose one of the following methods:
   - Click **Applications & Tiers** > **Storage Tiers**.
   - Click **Storage** > **Aggregates & Disks** > **Aggregates**.
2. In the **Storage Tiers** window, select the aggregate, and then click **More Actions > Add Cache**.

   **Note:** Adding cache is not supported on FabricPool-enabled aggregates.

3. In the **Add Cache** dialog box, perform the appropriate action:

| If you selected the cache source as... | Do this... |
|---|---|
| Storage pools | **a.** Select the storage pool from which cache can be obtained.<br>**b.** Specify the cache size.<br>**c.** Modify the RAID type, if required. |
| Dedicated SSDs | Select the SSD size and the number of SSDs to include, and optionally modify the RAID configuration:<br>**a.** Click **Change**.<br>**b.** In the Change RAID Configuration dialog box, specify the RAID type and RAID group size, and then click **Save**. |

4. Click **Add**.
   For mirrored aggregates, an Add Cache dialog box is displayed with the information that twice the number of selected disks will be added.
5. In the **Add Cache** dialog box, click **Yes**.

### Result

The cache disks are added to the selected aggregate.

#### Related information

*NetApp Technical Report 4070: Flash Pool Design and Implementation Guide*

## Increasing the cache for Flash Pool aggregates by adding SSDs

You can add SSDs as either storage pools or dedicated SSDs to increase the size of a Flash Pool aggregate by using System Manager.

### Before you begin

- The Flash Pool aggregate must be online.
- There must be sufficient spare SSDs or allocation units in the storage pool that can be assigned as cache disks.

### Steps

1. Click **Storage** > **Aggregates & Disks** > **Aggregates**.
2. In the **Aggregates** window, select the Flash Pool aggregate, and then click **Add Cache**.
3. In the **Add Cache** dialog box, perform the appropriate action:

| If you selected the cache source as... | Do this... |
|---|---|
| Storage pools | Select the storage pool from which cache can be obtained, and specify the cache size. |
| Dedicated SSDs | Select the SSD size and the number of SSDs to include. |

4. Click **Add**.
   For mirrored aggregates, an Add Cache dialog box is displayed with the information that twice the number of selected disks will be added.
5. In the **Add Cache** dialog box, click **Yes**.

### Result

The cache disks are added to the selected Flash Pool aggregate.

## Adding capacity disks

You can increase the size of an existing non-root aggregate or a root aggregate containing disks by adding capacity disks. You can use System Manager to add HDDs or SSDs of the selected ONTAP disk type and to modify the RAID group options.

### Before you begin

- The aggregate must be online.
- There must be sufficient compatible spare disks.

### About this task

- It is a best practice to add disks that are of the same size as the other disks in the aggregate.

  If you add disks that are smaller in size than the other disks in the aggregate, the aggregate becomes suboptimal in configuration, which in turn might cause performance issues.

  If you add disks that are larger in size than the disks that are available in a pre-existing RAID group within the aggregate, then the disks are downsized, and their space is reduced to that of the other disks in that RAID group. If a new RAID group is created in the aggregate and similar sized disks remain in the new RAID group, the disks will not be downsized.

  If you add disks that are not of the same size as the other disks in the aggregate, the selected disks might not be added; instead, other disks with a usable size between 90 percent and 105 percent of the specified size are automatically added. For example, for a 744 GB disk, all of the disks in the range of 669 GB through 781 GB are eligible for selection. For all of the spare disks in this range, ONTAP first selects only partitioned disks, then selects only unpartitioned disks, and finally selects both partitioned disks and unpartitioned disks.
- You cannot use System Manager to add HDDs to the following configurations:
  - Aggregates containing only SSDs
  - Root aggregates containing partitioned disks

  You must use the command-line interface to add HDDs to these configurations.
- For shared disks, RAID-DP is the only supported RAID type.
- You cannot use SSDs with storage pool.
- If the RAID group type is RAID-DP, and if you are adding FSAS or MSATA type of disks that are equal to or larger than 10 TB in size, then you can add them only to **`Specific RAID group`**, and not to **`New RAID group`** or **`All RAID groups`**.

  The disks are added after downsizing the disk size to the size of the disks in the pre-existing RAID group of the existing aggregate.
- If the RAID group type is RAID-TEC, and if you are adding FSAS or MSATA type of disks that are equal to or larger than 10 TB in size, then you can add them to **`All RAID groups`**, **`New RAID group`**, and **`Specific RAID group`**.

  The disks are added after downsizing the disk size to the size of the disks in the pre-existing RAID group of the existing aggregate.

### Steps

1. Choose one of the following methods:
   - Click **Applications & Tiers** > **Storage Tiers**.
   - Click **Storage** > **Aggregates & Disks** > **Aggregates**.
2. In the **Storage Tiers** window, select the aggregate to which you want to add capacity disks, and then click **More Actions > Add Capacity**.
3. Specify the following in the **Add Capacity** dialog box:
   a. The disk type for the capacity disks by using the **Disk Type to Add** option.
   b. The number of capacity disks by using the **Number of Disks or Partitions** option.

4. Specify the RAID group to which the capacity disks are to be added by using the **Add Disks To** option.

   By default, System Manager adds the capacity disks to **`All RAID groups`**.

   a. Click **Change**.
   b. In the **RAID Group Selection** dialog box, specify the RAID group as **`New RAID group`** or **`Specific RAID group`** by using the **Add Disks To** option.

      Shared disks can be added only to the **`New RAID group`** option.
5. Click **Add**.
   For mirrored aggregates, an Add Capacity dialog box is displayed with the information that twice the number of selected disks will be added.
6. In the **Add Capacity** dialog box, click **Yes** to add the capacity disks.

### Result

The capacity disks are added to the selected aggregate, and the aggregate size is increased.

### Related concepts

*What compatible spare disks are* on page 125

In System Manager, compatible spare disks are disks that match the properties of other disks in the aggregate. When you want to increase the size of an existing aggregate by adding HDDs (capacity disks) or change the RAID type of an aggregate from RAID4 to RAID-DP, the aggregate must contain sufficient compatible spare disks.

## Changing the RAID group when adding capacity disks

While adding capacity disks (HDDs) to an aggregate, you can change the RAID group to which you want to add the disks by using System Manager.

### About this task

- If the RAID type is RAID-DP, and if you are adding FSAS or MSATA type of disks that are equal to or larger than 10 TB in size, then you can add them only to **`Specific RAID group`**, and not to **`New RAID group`** or **`All RAID groups`**.

  The disks are added after downsizing the disk size to the size of the existing aggregates.
- If the RAID group is RAID-TEC, and if you are adding FSAS or MSATA type of disks that are equal to or larger than 10 TB in size, then you can add them to **`All RAID groups`**, **`New RAID group`**, and **`Specific RAID group`**.

  The disks are added after downsizing the disk size to the size of the existing aggregates.

### Steps

1. Choose one of the following methods:
   - Click **Applications & Tiers** > **Storage Tiers**.
   - Click **Storage** > **Aggregates & Disks** > **Aggregates**.
2. In the **Storage Tiers** window, select the aggregate to which you want to add capacity disks, and then click **More Actions > Add Capacity**.
3. In the **Add Capacity** dialog box, perform the following steps:
   a. Click **Change**.
   b. In the **Change RAID Configuration** dialog box, specify the RAID group to which you want to add the capacity disks.

      You can change the default value **`All RAID groups`** to either **`Specific RAID group`** or **`New RAID group`**.
   c. Click **Save**.

## Moving FlexVol volumes

You can nondisruptively move a FlexVol volume to a different aggregate or a different node for capacity utilization and improved performance by using System Manager.

### Before you begin

If you are moving a data protection volume, data protection mirror relationships must be initialized before you move the volume.

### About this task

- When you move a volume that is hosted on a Flash Pool aggregate, only the data that is stored in the HDD tier is moved to the destination aggregate.

  The cache data that is associated with the volume is not moved to the destination aggregate. Therefore, some performance degradation might occur after the volume move.
- You cannot move volumes from a SnapLock aggregate.
- You cannot move volumes from an SVM that is configured for disaster recovery to a FabricPool-enabled aggregate.

### Steps

1. Choose one of the following methods:
   - Click **Applications & Tiers** > **Storage Tiers**.
   - Click **Storage** > **Aggregates & Disks** > **Aggregates**.
2. Select the aggregate that contains the volume, and then click **More Actions > Volume Move**.
3. Type or select information as prompted by the wizard.
4. Confirm the details, and then click **Finish** to complete the wizard.

## Mirroring aggregates

You can use System Manager to protect data and to provide increased resiliency by mirroring data in real-time, within a single aggregate. Mirroring aggregates removes single points of failure in connecting to disks and array LUNs.

### Before you begin

There must be sufficient free disks in the other pool to mirror the aggregate.

### About this task

You cannot mirror a Flash Pool aggregate when the cache source is storage pool.

### Steps

1. Choose one of the following methods:
   - Click **Applications & Tiers** > **Storage Tiers**.
   - Click **Storage** > **Aggregates & Disks** > **Aggregates**.
2. Select the aggregate that you want to mirror, and then click **More Actions > Mirror**.

   **Note:** SyncMirror is not supported on FabricPool-enabled aggregates.
3. In the **Mirror this aggregate** dialog box, click **Mirror** to initiate the mirroring.

## Viewing aggregate information

You can use the Aggregates window in System Manager to view the name, status, and space information about an aggregate.

### Steps

1. Choose one of the following methods:
   - Click **Applications & Tiers** > **Storage Tiers**.
   - Click **Storage** > **Aggregates & Disks** > **Aggregates**.
2. Click on the aggregate name to view the details of the selected aggregate.

## Installing a CA certificate if you use StorageGRID Webscale

For ONTAP to authenticate with StorageGRID as the object store for a FabricPool-enabled aggregate, you can install a StorageGRID CA certificate on the cluster.

### Steps

1. Follow the StorageGRID system documentation to copy the CA certificate of the StorageGRID system by using the Grid Management Interface.

   *StorageGRID Webscale 11.1 Administrator Guide*

   While adding StorageGRID as an external capacity tier, a message is displayed if the CA certificate is not installed.
2. Add the StorageGRID CA certificate.

   **Note:** The fully qualified domain name (FQDN) that you specify must match the custom common name on the StorageGRID CA certificate.

### Related tasks

*Adding an external capacity tier* on page 131

You can use System Manager to add an external capacity tier to an SSD aggregate or a VMDISK aggregate. External capacity tiers provide storage for infrequently used data.

## How moving a FlexVol volume works

Knowing how moving a FlexVol volume works helps you to determine whether the volume move satisfies service-level agreements and to understand where a volume move is in the volume move process.

FlexVol volumes are moved from one aggregate or node to another within the same storage virtual machine (SVM). A volume move does not disrupt client access during the move.

Moving a volume occurs in multiple phases:

- A new volume is made on the destination aggregate.
- The data from the original volume is copied to the new volume.
  During this time, the original volume is intact and available for clients to access.
- At the end of the move process, client access is temporarily blocked.
  During this time the system performs a final replication from the source volume to the destination volume, swaps the identities of the source and destination volumes, and changes the destination volume to the source volume.
- After completing the move, the system routes client traffic to the new source volume and resumes client access.

The move is not disruptive to client access because the time in which client access is blocked ends before clients notice a disruption and time out. Client access is blocked for 35 seconds by default. If the volume move operation cannot finish in the time that access is denied, the system aborts this final phase of the volume move operation and allows client access. The system attempts the final

phase three times by default. After the third attempt, the system waits an hour before attempting the final phase sequence again. The system runs the final phase of the volume move operation until the volume move is complete.

## How you can use effective ONTAP disk type for mixing HDDs

Starting with Data ONTAP 8.1, certain ONTAP disk types are considered equivalent for the purposes of creating and adding to aggregates, and managing spares. ONTAP assigns an effective disk type for each disk type. You can mix HDDs that have the same effective disk type.

When the `raid.disktype.enable` option is set to **`off`**, you can mix certain types of HDDs within the same aggregate. When the `raid.disktype.enable` option is set to **`on`**, the effective disk type is the same as the ONTAP disk type. Aggregates can be created using only one disk type. The default value for the `raid.disktype.enable` option is **`off`**.

Starting with Data ONTAP 8.2, the option `raid.mix.hdd.disktype.capacity` must be set to **`on`** to mix disks of type BSAS, FSAS, and ATA. The option `raid.mix.hdd.disktype.performance` must be set to **`on`** to mix disks of type FCAL and SAS.

The following table shows how the disk types map to the effective disk type:

| ONTAP disk type | Effective disk type |
|---|---|
| FCAL | SAS |
| SAS | SAS |
| ATA | FSAS |
| BSAS | FSAS |
| FCAL and SAS | SAS |
| MSATA | MSATA |
| FSAS | FSAS |

## What compatible spare disks are

In System Manager, compatible spare disks are disks that match the properties of other disks in the aggregate. When you want to increase the size of an existing aggregate by adding HDDs (capacity disks) or change the RAID type of an aggregate from RAID4 to RAID-DP, the aggregate must contain sufficient compatible spare disks.

Disk properties that must match are the disk type, disk size (can be a higher size disk in case the same disk size is not available), disk RPM, checksum, node owner, pool, and shared disk properties. If you use higher sized disks, you must be aware that disk downsizing occurs and the size of all disks are reduced to the lowest disk size. Existing shared disks are matched with higher size non-shared disks, and the non-shared disks are converted to shared disks and added as spares.

If RAID mixing options, such as disk type mixing and disk RPM mixing, are enabled for the RAID group, the disk type and disk RPM of the existing disks of the aggregate are matched with the effective disk type and effective disk RPM of the spare disks to obtain compatible spares.

### Related tasks

*Adding capacity disks* on page 121
You can increase the size of an existing non-root aggregate or a root aggregate containing disks by adding capacity disks. You can use System Manager to add HDDs or SSDs of the selected ONTAP disk type and to modify the RAID group options.

*Editing aggregates* on page 117

You can use System Manager to change the aggregate name, RAID type, and RAID group size of an existing aggregate when required.

## How System Manager works with hot spares

A hot spare is a disk that is assigned to a storage system but not used by any RAID group. Hot spares do not contain any data and are assigned to a RAID group when a disk failure occurs in the RAID group. System Manager uses the largest disk as the hot spare.

When there are different disk types in the RAID group, the largest-sized disk of each disk type is left as the hot spare. For example, if there are 10 SATA disks and 10 SAS disks in the RAID group, the largest-sized SATA disk and the largest-sized SAS disk are served as hot spares.

If the largest-sized disk is partitioned, then the hot spares are provided separately for partitioned and non-partitioned RAID groups. If the largest-sized disk is unpartitioned, then a single spare disk is provided.

The largest-sized non-partitioned disk is left as a hot spare if there are root partitions in the disk group. When a non-partitioned disk of the same size is not available, then spare root partitions are left as hot spares for the root partitioned group.

A single spare disk can serve as a hot spare for multiple RAID groups. System Manager calculates the hot spares based on the value set in the option `raid.min_spare_count` at the node level. For example, if there are 10 SSDs in an SSD RAID group and the option `raid.min_spare_count` is set to **`1`** at the node level, System Manager leaves 1 SSD as the hot spare and uses the other 9 SSDs for SSD-related operations. Similarly, if there are 10 HDDs in an HDD RAID group and the option `raid.min_spare_count` is set to **`2`** at the node level, System Manager leaves 2 HDDs as hot spares and uses the other 8 HDDs for HDD-related operations.

System Manager enforces the hot spare rule for RAID groups when you create an aggregate, edit an aggregate, and when you add HDDs or SSDs to an aggregate. The hot spare rule is also used when you create a storage pool or add disks to an existing storage pool.

There are exceptions to the hot spare rule in System Manager:

- For MSATA or disks in a multi-disk carrier, the number of hot spares is twice the value set at the node level and the number must not be less than 2 at any time.
- Hot spares are not used if the disks are part of array LUNs or virtual storage appliances.

## Rules for displaying disk types and disk RPM

When you are creating an aggregate and adding capacity disks to an aggregate, you should understand the rules that apply when disk types and disk RPM are displayed.

When the disk type mixing and the disk RPM mixing options are not enabled, the actual disk type and actual disk RPM are displayed.

When these mixing options are enabled, the effective disk type and effective disk RPM are displayed instead of the actual disk type and actual disk RPM. For example, when the disk mixing option is enabled, System Manager displays BSAS disks as FSAS. Similarly, when the disk RPM mixing option is enabled, if the RPM of the disks is 10K and 15K, System Manager displays the effective RPM as 10K.

## How mirrored aggregates work

Mirrored aggregates have two *plexes* (copies of their data), which use the SyncMirror functionality to duplicate the data to provide redundancy.

When a mirrored aggregate is created (or when a second plex is added to an existing unmirrored aggregate), ONTAP copies the data in the original plex (plex0) to the new plex (plex1). The plexes are physically separated (each plex has its own RAID groups and its own pool), and the plexes are

updated simultaneously. This provides added protection against data loss if more disks fail than the RAID level of the aggregate protects against or there is a loss of connectivity, because the unaffected plex continues to serve data while you fix the cause of the failure. After the plex that had a problem is fixed, the two plexes resynchronize and reestablish the mirror relationship.

The disks and array LUNs on the system are divided into two pools: pool0 and pool1. Plex0 gets its storage from pool0 and plex1 gets its storage from pool1.

The following diagram shows an aggregate composed of disks with the SyncMirror functionality enabled and implemented. A second plex has been created for the aggregate, plex1. The data in plex1 is a copy of the data in plex0, and the RAID groups are also identical. The 32 spare disks are allocated to pool0 or pool1, 16 disks for each pool.

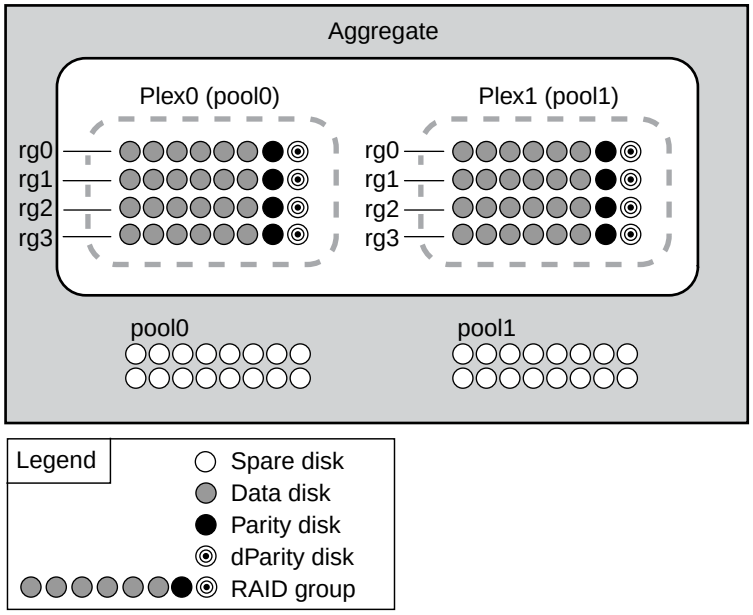


The following diagram shows an aggregate composed of array LUNs with the SyncMirror functionality enabled and implemented. A second plex has been created for the aggregate, plex1. Plex1 is a copy of plex0, and the RAID groups are also identical.

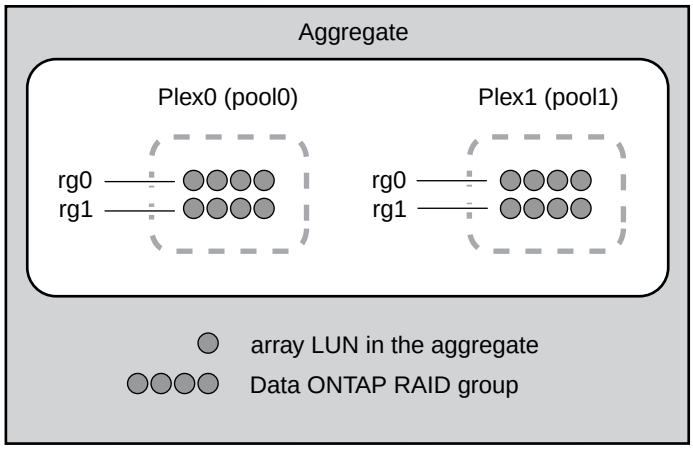


## What a FabricPool is

FabricPool is a hybrid storage solution that uses an all flash (all SSD) aggregate as the performance tier and an object store as the external capacity tier. Data in a FabricPool is stored in a

tier based on whether it is frequently accessed or not. Using a FabricPool helps you reduce storage cost without compromising performance, efficiency, or protection.

### Related tasks

*Adding an external capacity tier* on page 131
You can use System Manager to add an external capacity tier to an SSD aggregate or a VMDISK aggregate. External capacity tiers provide storage for infrequently used data.

*Attaching an aggregate to an external capacity tier* on page 132
You can use System Manager to attach an All Flash aggregate to an external capacity tier. You can store infrequently used data in external capacity tiers.

## Storage recommendations for creating aggregates

Starting with System Manager 9.4, you can create aggregates based on storage recommendations. However, you must determine whether create aggregates based on storage recommendations is supported in your environment. If it is not, you must decide the RAID policy and disk configuration, and then create the aggregates manually.

System Manager analyzes the available spare disks in the cluster and generates a recommendation about how the spare disks should be used to create aggregates according to best practices. System Manager displays the summary of recommended aggregates including their names and usable size.

In many cases, the storage recommendation will be optimal for your environment. However, if your cluster is running ONTAP 9.3 or earlier, or if your environment includes the following configurations, you must create aggregates manually:

- Aggregates using third-party array LUNs
- Virtual disks with Cloud Volumes ONTAP or ONTAP Select
- MetroCluster configurations
- SyncMirror functionality
- MSATA disks
- Flash Pool aggregates
- Multiple disk types or sizes are connected to the node

In addition, if any of the following disk conditions exist in your environment, you must rectify the disk conditions before you use the storage recommendation to create aggregates:

- Missing disks
- Fluctuation in spare disk numbers
- Unassigned disks
- Non-zeroed spares
- Disks that are undergoing maintenance testing

### Related tasks

*Zeroing spare disks* on page 38
You can use System Manager to erase all the data and to format the spare disks by writing zeros to the disk. These disks can then be used in new aggregates.

### Related information

*Disk and aggregate management*

## Storage Tiers window

You can use the Storage Tiers window to view cluster-wide space details and to add and view aggregate details.

The Internal Tier panel, or the Performance Tier panel if the cluster has all flash (all SSD) aggregates, displays cluster-wide space details such as the sum of the total sizes of all of the aggregates, the space used by the aggregates in the cluster, and the available space in the cluster.

The External Capacity Tier panel displays the total licensed external capacity tiers in the cluster, the licensed space that is used in the cluster, and the licensed space that is available in the cluster. The External Capacity Tier panel also displays the unlicensed external capacity that is used.

Aggregates are grouped by type, and the aggregate panel displays details about the total aggregate space, space used, and the available space. You can select the aggregate and perform any of the aggregate-related actions.

### Command buttons

**Add Aggregate**

Enables you to create an aggregate.

**Actions**

Provides the following options:

**Change status to**

Changes the status of the selected aggregate to one of the following statuses:

- Online
  Read and write access to the volumes that are contained in this aggregate is allowed.
- Offline
  Read and write access is not allowed.
- Restrict
  Some operations such as parity reconstruction are allowed, but data access is not allowed.

**Add Capacity**

Enables you to add capacity (HDDs or SSDs) to existing aggregates.

**Add Cache**

Enables you to add cache disks (SSDs) to existing HDD aggregates or Flash Pool aggregates.

You cannot add cache disks to FabricPool-enabled aggregates.

This option is not available for a cluster containing nodes with All Flash Optimized personality.

**Mirror**

Enables you to mirror the aggregates.

**Volume Move**

Enables you to move a FlexVol volume.

### Details area

You can click the aggregate name to view detailed information about the aggregate.

**Overview tab**

Displays detailed information about the selected aggregate, and displays a pictorial representation of the space allocation of the aggregate, the space savings of the aggregate, and the performance of the aggregate .

**Disk Information tab**

Displays the disk layout information for the selected aggregate.

**Volumes tab**

Displays details about the total number of volumes on the aggregate, the total aggregate space, and the space committed to the aggregate.

**Performance tab**

Displays graphs that show the performance metrics of the aggregates, including total transfers and IOPS. Performance metrics data for read, write, and total transfers is displayed, and the data for SSDs and HDDs is recorded separately.

Changing the client time zone or the cluster time zone impacts the performance metrics graphs. You should refresh your browser to view the updated graphs.

### Related tasks

*Adding an external capacity tier* on page 131
You can use System Manager to add an external capacity tier to an SSD aggregate or a VMDISK aggregate. External capacity tiers provide storage for infrequently used data.

*Attaching an aggregate to an external capacity tier* on page 132
You can use System Manager to attach an All Flash aggregate to an external capacity tier. You can store infrequently used data in external capacity tiers.

*Deleting an external capacity tier* on page 134
You can use System Manager to delete an external capacity tier that you no longer require.

*Editing an external capacity tier* on page 134
You can use System Manager to modify the configuration information of an external capacity tier. The configuration details that you can edit include the name, fully qualified domain name (FQDN), port, access key ID, secret key, and object store certificate.

*Provisioning storage through aggregates* on page 38
You can create an aggregate based on storage recommendations or manually depending on your requirement. You can create Flash Pool aggregates, SnapLock aggregates, and a FabricPool-enabled aggregates to provide storage for one or more volumes by using System Manager.

*Deleting aggregates* on page 118
You can use System Manager to delete aggregates when you no longer require the data in the aggregates. However, you cannot delete the root aggregate because it contains the root volume, which contains the system configuration information.

*Editing aggregates* on page 117
You can use System Manager to change the aggregate name, RAID type, and RAID group size of an existing aggregate when required.

## Configuring and managing external capacity tiers

Storing data in tiers can enhance the efficiency of your storage system. You manage storage tiers by using FabricPool-enabled aggregates. External capacity tiers stores data in a tier based on whether the data is frequently accessed.

### Before you begin

- You must be running ONTAP 9.2 or later.

- You must have all flash (all SSD) aggregates

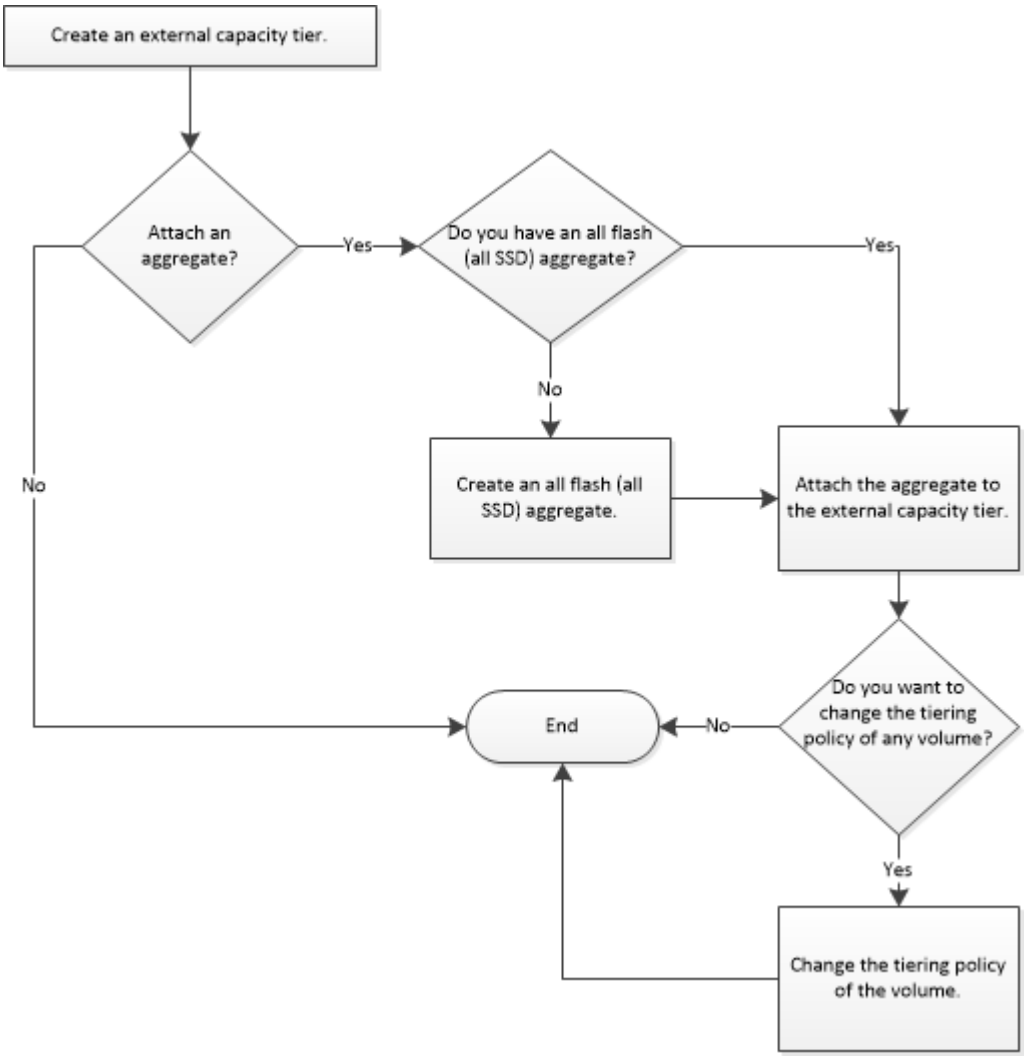


## Adding an external capacity tier

You can use System Manager to add an external capacity tier to an SSD aggregate or a VMDISK aggregate. External capacity tiers provide storage for infrequently used data.

### Before you begin

- You must have the access key ID and secret key to connect to the object store.
- You must have created a bucket inside the object store.
- Network connectivity must exist between the cluster and the external capacity tier.
- If communication between the external capacity tier and cluster is encrypted using SSL or TLS, the required certificates must be installed.

### About this task

The supported object stores that can be used as external capacity tiers are StorageGRID, Amazon AWS S3, and Microsoft Azure Blob storage.

**Note:**

- Azure Stack, which is an on-premises Azure services, is not supported.

- If you want to use Amazon AWS S3 or Microsoft Azure Blob storage as an external capacity tier, you must have the FabricPool capacity license.

### Steps

1. Click **Storage** > **Aggregates & Disks** > **External Capacity Tier**.
2. Click **Add External Capacity Tier**.
   The Add External Capacity Tier window is displayed.
3. Specify the server name that hosts the external capacity tier, the port to access the external capacity tier, the access key ID of the external capacity tier, the secret key of the external tier, and the container name.
4. Enable the **SSL** option if you want to transfer the data securely to the external capacity tier.
5. If you want to add an external capacity tier for StorageGRID, enable the **Object Store Certificate** option, copy the contents of the certificate, and then paste the certificate contents in the signed certification.
6. From the **IPspace** list, select the IPspace that is used to connect to the external capacity tier.
7. Click **Save** to save the external capacity tier.
8. Click **Save and Attach Aggregates** to save the external capacity tier and to attach aggregates to the external capacity tier.

### Related concepts

*What external capacity tiers and tiering policies are* on page 135
External capacity tiers provide storage for infrequently accessed data. You can attach an all flash (all SSD) aggregate to an external capacity tier to store infrequently used data. You can use tiering policies to decide whether data should be moved to an external capacity tier.

*What a FabricPool is* on page 127
FabricPool is a hybrid storage solution that uses an all flash (all SSD) aggregate as the performance tier and an object store as the external capacity tier. Data in a FabricPool is stored in a tier based on whether it is frequently accessed or not. Using a FabricPool helps you reduce storage cost without compromising performance, efficiency, or protection.

### Related tasks

*Installing a CA certificate if you use StorageGRID Webscale* on page 124
For ONTAP to authenticate with StorageGRID as the object store for a FabricPool-enabled aggregate, you can install a StorageGRID CA certificate on the cluster.

### Related reference

*Storage Tiers window* on page 129
You can use the Storage Tiers window to view cluster-wide space details and to add and view aggregate details.

## Attaching an aggregate to an external capacity tier

You can use System Manager to attach an All Flash aggregate to an external capacity tier. You can store infrequently used data in external capacity tiers.

### Before you begin

You must have added an external capacity tier to the cluster.

### Steps

1. Click **Storage** > **Aggregates & Disks** > **External Capacity Tier**.
2. In the **Used in Aggregates** column, click **Attach Aggregates**.
   The Attach Aggregates window is displayed.

3. Select the aggregate that you want to attach to the external capacity tier.
4. If you want to update the tiering policy of the volumes that are on the selected aggregate, click **View and Update Tiering Policy for Volumes**.
5. Click the **Change Tiering Policy** list, and then update the tiering policy.
6. Click **Save**.

### Related concepts

*What external capacity tiers and tiering policies are* on page 135
External capacity tiers provide storage for infrequently accessed data. You can attach an all flash (all SSD) aggregate to an external capacity tier to store infrequently used data. You can use tiering policies to decide whether data should be moved to an external capacity tier.

*What a FabricPool is* on page 127
FabricPool is a hybrid storage solution that uses an all flash (all SSD) aggregate as the performance tier and an object store as the external capacity tier. Data in a FabricPool is stored in a tier based on whether it is frequently accessed or not. Using a FabricPool helps you reduce storage cost without compromising performance, efficiency, or protection.

### Related reference

*Storage Tiers window* on page 129
You can use the Storage Tiers window to view cluster-wide space details and to add and view aggregate details.

## Provisioning storage by creating a FabricPool-enabled aggregate manually

You can use System Manager to create a FabricPool-enabled aggregate to attach an external capacity tier to the SSD aggregate.

### Before you begin

- You must have created an external capacity tier and attached the external capacity tier to the cluster in which the SSD aggregate resides.
- An on-premises external capacity tier must have been created.
- A dedicated network connection must exist between the external capacity tier and the aggregate.

### About this task

- The supported external capacity tiers are StorageGRID, Amazon AWS S3, and Microsoft Azure Blob storage.

  **Note:**

  - Azure Stack, which is an on-premises Azure services, is not supported.
  - If you want to use Amazon AWS S3 or Microsoft Azure Blob storage as an external capacity tier, you must have the FabricPool capacity license.

- FabricPool-enabled aggregates are not supported on MetroCluster configurations.

### Steps

1. Create a FabricPool-enabled aggregate by using one of the following methods:
   - Click **Applications & Tiers** > **Storage Tiers** > **Add Aggregate**.
   - Click **Storage** > **Aggregate & Disks** > **Aggregates** > **Create**.
2. Enable the **Manually Create Aggregate** option to create an aggregate.
3. To create a FabricPool-enabled aggregate:
   a. Specify the name of the aggregate, the disk type, and the number of disks or partitions to include in the aggregate.

**Note:** Only all flash (all SSD) aggregates support FabricPool-enabled aggregates.

The minimum hot spare rule is applied to the disk group that has the largest disk size.

b. Optional: Modify the RAID configuration of the aggregate:

   **i.** Click **Change**.

   **ii.** In the Change RAID Configuration dialog box, specify the RAID type and the RAID group size.
   RAID-DP is the only supported RAID type for shared disks.

   **iii.** Click **Save**.

4. Select the **FabricPool** checkbox, and then select an external capacity tier from the list.
5. Click **Create**.

## Changing the tiering policy of a volume

You can use System Manager to change the default tiering policy of a volume to control whether the data of the volume is moved to the capacity tier when the data becomes inactive.

### Steps

1. Click **Storage** > **Volumes**.
2. From the **Volumes on SVM** list, select **All SVMs**.
3. Select the volume for which you want to change the tiering policy, and then click **More Actions** > **Change Tiering Policy**.
4. Select the required tiering policy from the **Tiering Policy** list, and then click **Save**.

## Editing an external capacity tier

You can use System Manager to modify the configuration information of an external capacity tier. The configuration details that you can edit include the name, fully qualified domain name (FQDN), port, access key ID, secret key, and object store certificate.

### Steps

1. Click **Storage** > **Aggregates & Disks** > **External Capacity Tier**.
2. Select the external capacity tier that you want to edit, and then click **Edit**.
3. In the **Edit External Capacity Tier** window, modify the external capacity tier name, FQDN, port, access key ID, secret key, and object store certificate, as required.
4. Click **Save**.

### Related reference

*Storage Tiers window* on page 129
You can use the Storage Tiers window to view cluster-wide space details and to add and view aggregate details.

## Deleting an external capacity tier

You can use System Manager to delete an external capacity tier that you no longer require.

### Before you begin

You must have deleted the FabricPool-enabled aggregate that is associated with the external capacity tier.

### Steps

1. Click **Storage** > **Aggregates & Disks** > **External Capacity Tier**.
2. Select the external capacity tier that you want to delete, and then click **Delete**.

### Related reference

*Storage Tiers window* on page 129
You can use the Storage Tiers window to view cluster-wide space details and to add and view aggregate details.

## What external capacity tiers and tiering policies are

External capacity tiers provide storage for infrequently accessed data. You can attach an all flash (all SSD) aggregate to an external capacity tier to store infrequently used data. You can use tiering policies to decide whether data should be moved to an external capacity tier.

You can set one of the following tiering policies on a volume:

**Snapshot-only**

Moves the Snapshot copies of only those volumes that are currently not being referenced by the active file system. Snapshot-only policy is the default tiering policy.

**None**

Prevents the data on the volume from being moved to an external capacity tier.

**Backup**

Moves the newly transferred data of a data protection (DP) volume to the external capacity tier.

**Auto**

Moves the inactive (cold) data and the Snapshot copies from the active file system to the external capacity tier.

### Related tasks

*Adding an external capacity tier* on page 131
You can use System Manager to add an external capacity tier to an SSD aggregate or a VMDISK aggregate. External capacity tiers provide storage for infrequently used data.

*Attaching an aggregate to an external capacity tier* on page 132
You can use System Manager to attach an All Flash aggregate to an external capacity tier. You can store infrequently used data in external capacity tiers.

## What inactive (cold) data is

Infrequently accessed data in a performance tier is known as inactive (cold) data. By default, data that is not accessed for a period of 31 days becomes inactive.

Inactive data is displayed at the aggregate level, cluster level, and volume level. The inactive data for an aggregate or a cluster is displayed only if inactive scanning is complete on that aggregate or cluster. By default, inactive data is displayed for FabricPool-enabled aggregates. If you want to view inactive data for SSD aggregates, you must manually enable inactive data.

## External Capacity Tier window

You can use System Manager to add, edit, and delete external capacity tiers and to view external capacity tier details.

The External Capacity Tier panel displays the total number of licensed external capacity tiers in the cluster, the licensed space that is used in the cluster, and the licensed space that is available in the cluster. The External Capacity Tier panel also displays the unlicensed external capacity that is used.

### Command buttons

**Add**

Enables you to add an external capacity tier.

**Attach Aggregates**

Enables you to attach aggregates to an external capacity tier.

**Delete**

Enables you to delete a selected external capacity tier.

**Edit**

Enables you to modify the properties of a selected external capacity tier.

### Details area

You can view detailed information about external capacity tiers such as the list of external capacity tiers, the details of the object stores, the aggregates used, and the used capacity.

If you create an external capacity tier other than StorageGRID, Amazon AWS S3, and Microsoft Azure Blob storage by using the command-line interface (CLI), this external capacity tier is displayed as Others in System Manager. You can then attach aggregates to this external capacity tier.

# Aggregates

You can use System Manager to create aggregates to support the differing security, backup, performance, and data sharing requirements of your users.

## Aggregates window

You can use the Aggregates window to create, display, and manage information about aggregates.

- *Aggregates window* on page 136
- *Aggregate list* on page 137
- *Details area* on page 138
- *Command buttons* on page 136

### Command buttons

**Create**

Opens the Create Aggregate dialog box, which enables you to create an aggregate.

**Edit**

Opens the Edit Aggregate dialog box, which enables you to change the name of an aggregate or the level of RAID protection that you want to provide for the aggregate.

**Delete**

Deletes the selected aggregate.

**Note:** This button is disabled for the root aggregate.

**More Actions**

Provides the following options:

**Change status to**

Changes the status of the selected aggregate to one of the following statuses:

- Online
  Read and write access to the volumes that are contained in this aggregate is allowed.

- Offline
  Read and write access is not allowed.
- Restrict
  Some operations—such as parity reconstruction—are allowed, but data access is not allowed.

**Add Capacity**

Enables you to add capacity (HDDs or SSDs) to existing aggregates.

**Add Cache**

Enables you to add cache disks (SSDs) to existing HDD aggregates or Flash Pool aggregates.

This button is not available for a cluster containing nodes with All Flash Optimized personality.

**Mirror**

Enables you to mirror the aggregates.

**Volume Move**

Enables you to move a FlexVol volume.

**Attach External Capacity Tier**

Enables you to attach an external capacity tier to the aggregate.

**Refresh**

Updates the information in the window.

### Aggregate list

Displays the name and the space usage information for each aggregate.

**Status**

Displays the status of the aggregate.

**Name**

Displays the name of the aggregate.

**Node**

Displays the name of the node to which the disks of the aggregate are assigned.

This field is available only at the cluster level.

**Type**

Displays the type of the aggregate.

This field is not displayed for a cluster containing nodes with All Flash Optimized personality.

**Used (%)**

Displays the percentage of space that is used in the aggregate.

**Available Space**

Displays the available space in the aggregate.

**Used Space**

Displays the amount of space that is used for data in the aggregate.

**Total Space**

Displays the total space of the aggregate.

**FabricPool**

Displays whether the selected aggregate is attached to an external capacity tier.

**External Capacity Tier**

If the selected aggregate is attached to an external capacity tier, displays the name of the external capacity tier.

**Volume Count**

Displays the number of volumes that are associated with the aggregate.

**Disk Count**

Displays the number of disks that are used to create the aggregate.

**Flash Pool**

Displays the total cache size of the Flash Pool aggregate. A value of NA indicates that the aggregate is not a Flash Pool aggregate.

This field is not displayed for a cluster containing nodes with All Flash Optimized personality.

**Mirrored**

Displays whether the aggregate is mirrored.

**SnapLock Type**

Displays the SnapLock type of the aggregate.

### Details area

Select an aggregate to view information about the selected aggregate. You can click Show More Details to view detailed information about the selected aggregate.

**Overview tab**

Displays detailed information about the selected aggregate, and displays a pictorial representation of the space allocation of the aggregate, the space savings of the aggregate, which includes the total logical space used, total physical space used, overall savings from storage efficiency, data reduction ratio, FlexClone volume ratio, and Snapshot copies ratio, and the performance of the aggregate in IOPS and total data transfers.

**Disk Information tab**

Displays disk layout information such as the name of the disk, disk type, physical size, usable size, disk position, disk status, plex name, plex status, RAID group, RAID type, and storage pool (if any) for the selected aggregate. The disk port that is associated with the disk primary path and the disk name with the disk secondary path for a multipath configuration are also displayed.

**Volumes tab**

Displays details about the total number of volumes on the aggregate, total aggregate space, and the space committed to the aggregate.

**Performance tab**

Displays graphs that show the performance metrics of the aggregates, including total transfers and IOPS. Performance metrics data for read, write, and total transfers is displayed, and the data for SSDs and HDDs is recorded separately.

Changing the client time zone or the cluster time zone impacts the performance metrics graphs. You should refresh your browser to view the updated graphs.

### Related tasks

*Provisioning storage through aggregates* on page 38

You can create an aggregate based on storage recommendations or manually depending on your requirement. You can create Flash Pool aggregates, SnapLock aggregates, and a FabricPool-enabled aggregates to provide storage for one or more volumes by using System Manager.

*Deleting aggregates* on page 118
You can use System Manager to delete aggregates when you no longer require the data in the aggregates. However, you cannot delete the root aggregate because it contains the root volume, which contains the system configuration information.

*Editing aggregates* on page 117
You can use System Manager to change the aggregate name, RAID type, and RAID group size of an existing aggregate when required.

# Storage pools

You can use System Manager to create storage pools to enable SSDs to be shared by multiple Flash Pool aggregates.

### Related information

*Disk and aggregate management*

## Creating a storage pool

A storage pool is a collection of SSDs (cache disks). You can use System Manager to combine SSDs to create a storage pool, which enables you to share the SSDs and SSD spares between an HA pair for allocation to two or more Flash Pool aggregates at the same time.

### Before you begin

- Both nodes of the HA pair must be up and running in order to allocate SSDs and SSD spares through a storage pool.
- Storage pools must have a minimum of 3 SSDs.
- All SSDs in a storage pool must be owned by the same HA pair.

### About this task

System Manager enforces the hot spare rule for SSD RAID groups when you use SSDs for adding disks to a storage pool. For example, if there are 10 SSDs in the SSD RAID group and the option `raid.min_spare_count` is set to **`1`** at the node level, System Manager leaves 1 SSD as the hot spare and uses the other 9 SSDs for SSD-related operations.

You cannot use partitioned SSDs when creating a storage pool by using System Manager.

### Steps

1. Click **Storage** > **Aggregates & Disks** > **Storage Pools**.
2. In the **Storage Pools** window, click **Create**.
3. In the **Create Storage Pool** dialog box, specify the name for the storage pool, disk size, and the number of disks.
4. Click **Create**.

### Related reference

*Storage Pools window* on page 144

You can use the Storage Pools window to create, display, and manage a dedicated cache of SSDs, also known as *storage pools*. These storage pools can be associated with a non-root aggregate to provide SSD cache and with a Flash Pool aggregate to increase its size.

## Adding disks to a storage pool

You can add SSDs to an existing storage pool and increase its cache size by using System Manager.

### Before you begin

Both nodes of the HA pair must be up and running in order to allocate SSDs and SSD spares through a storage pool.

### About this task

- The SSDs that you add to a storage pool are distributed proportionally among the aggregates using the storage pool cache and to the free space of the storage pool.
- System Manager enforces the hot spare rule for SSD RAID groups when you use SSDs for adding disks to a storage pool.
  For example, if there are 10 SSDs in the SSD RAID group and the option `raid.min_spare_count` is set to **`1`** at the node level, System Manager leaves 1 SSD as the hot spare and uses the other 9 SSDs for SSD-related operations.
- You cannot use partitioned SSDs when adding disks to a storage pool by using System Manager.

### Steps

1. Click **Storage** > **Aggregates & Disks** > **Storage Pools**.
2. In the **Storage Pools** window, select the storage pool, and then click **Add Disks**.
3. In the **Add Disks** dialog box, specify the number of disks that you want to add.
4. Click **Next**.
5. In the **Summary** dialog box, review how the cache is distributed among various aggregates and the free space of the storage pool.
6. Click **Add**.

### Related reference

*Storage Pools window* on page 144

You can use the Storage Pools window to create, display, and manage a dedicated cache of SSDs, also known as *storage pools*. These storage pools can be associated with a non-root aggregate to provide SSD cache and with a Flash Pool aggregate to increase its size.

## Deleting storage pools

You might want to delete a storage pool when the cache of the storage pool is not optimal or when it is no longer used by any aggregate or Flash Pool aggregate. You can delete a storage pool by using the Delete Storage Pool dialog box in System Manager.

### Before you begin

The storage pool must not be used by any aggregate.

### Steps

1. Click **Storage** > **Aggregates & Disks** > **Storage Pools**.
2. In the **Storage Pools** window, select the storage pool that you want to delete, and then click **Delete**.
3. In the **Delete Storage Pool** dialog box, click **Delete**.

### Related reference

*Storage Pools window* on page 144

You can use the Storage Pools window to create, display, and manage a dedicated cache of SSDs, also known as *storage pools*. These storage pools can be associated with a non-root aggregate to provide SSD cache and with a Flash Pool aggregate to increase its size.

## How to use SSD storage pools

To enable SSDs to be shared by multiple Flash Pool aggregates, you can add the SSDs to a *storage pool*. After you add an SSD to a storage pool, you can no longer manage the SSD as a stand-alone entity. You must use the storage pool to assign or allocate the storage that is provided by the SSD.

You can create storage pools for a specific high-availability (HA) pair. Then, you can add allocation units from that storage pool to one or more Flash Pool aggregates that are owned by the same HA pair. Just as disks must be owned by the same node that owns an aggregate before the disks can be allocated to it, storage pools can provide storage only to the Flash Pool aggregates that are owned by one of the nodes that owns the storage pool.

If you have to increase the amount of Flash Pool cache on your system, you can add more SSDs to a storage pool, up to the maximum RAID group size for the RAID type of the Flash Pool caches that are using the storage pool. When you add an SSD to an existing storage pool, you increase the size of the storage pool's allocation units, including any allocation units that are already allocated to a Flash Pool aggregate.

You can use only one spare SSD for a storage pool, so that if an SSD in that storage pool becomes unavailable, ONTAP can use the spare SSD to reconstruct the partitions of the malfunctioning SSD. You do not have to reserve any allocation units as spare capacity; ONTAP can use only a full, unpartitioned SSD as a spare for the SSDs in a storage pool.

After you add an SSD to a storage pool, you cannot remove the SSD, just as you cannot remove disks from an aggregate. If you want to use the SSDs in a storage pool as discrete drives again, you must destroy all of the Flash Pool aggregates to which the storage pool's allocation units have been allocated, and then destroy the storage pool.

## Requirements and best practices for using SSD storage pools

Some technologies cannot be combined with Flash Pool aggregates that use SSD storage pools.

You cannot use the following technologies with Flash Pool aggregates that use SSD storage pools for their cache storage:

- MetroCluster
- SyncMirror functionality
  Mirrored aggregates can coexist with Flash Pool aggregates that use storage pools; however, Flash Pool aggregates cannot be mirrored.
- Physical SSDs
  Flash Pool aggregates can use SSD storage pools or physical SSDs, but not both.

SSD storage pools must conform to the following rules:

- SSD storage pools can contain only SSDs; HDDs cannot be added to an SSD storage pool.
- SSD storage pools can contain between 3 and 28 SSDs.
  If an SSD storage pool contains more SSDs than the maximum RAID4 RAID group size for SSDs, then the SSD storage pool cannot be used for a Flash Pool aggregate whose cache has a RAID type of RAID4.
- All of the SSDs in an SSD storage pool must be owned by the same high-availability (HA) pair.
- You cannot use SSDs that have been partitioned for root-data partitioning in a storage pool.

If you provide storage from a single storage pool to two caches with different RAID types, and you expand the size of the storage pool beyond the maximum RAID group size for RAID4, the extra partitions in the RAID4 allocation units go unused. Therefore, it is a best practice to keep your cache RAID types homogenous for a storage pool.

You cannot change the RAID type of cache RAID groups that are allocated from a storage pool. You set the RAID type for the cache before adding the first allocation units, and you cannot change the RAID type later.

When you create a storage pool or add SSDs to an existing storage pool, you must use the same size SSDs. If a failure occurs and no spare SSD of the correct size exists, ONTAP can use a larger SSD to replace the failed SSD. However, the larger SSD is right-sized to match the size of the other SSDs in the storage pool, resulting in lost SSD capacity.

You can use only one spare SSD for a storage pool. If the storage pool provides allocation units to the Flash Pool aggregates that are owned by both nodes in the HA pair, then the spare SSD can be owned by either node. However, if the storage pool provides allocation units only to the Flash Pool aggregates that are owned by one of the nodes in the HA pair, then the SSD spare must be owned by that same node.

## Considerations for when to use SSD storage pools

SSD storage pools provide many benefits, but they also introduce some restrictions that you should be aware of when deciding whether to use SSD storage pools or dedicated SSDs.

SSD storage pools make sense only when they are providing cache to two or more Flash Pool aggregates. SSD storage pools provide the following benefits:

- Increased storage utilization for SSDs used in Flash Pool aggregates
  SSD storage pools reduce the overall percentage of SSDs needed for parity by enabling you to share parity SSDs between two or more Flash Pool aggregates.
- Ability to share spares between HA partners
  Because the storage pool is effectively owned by the HA pair, one spare, owned by one of the HA partners, can function as a spare for the entire SSD storage pool if needed.
- Better utilization of SSD performance
  The high performance provided by SSDs can support access by both controllers in an HA pair.

These advantages must be weighed against the costs of using SSD storage pools, which include the following items:

- Reduced fault isolation
  The loss of a single SSD affects all RAID groups that include one of its partitions. In this situation, every Flash Pool aggregate that has cache allocated from the SSD storage pool that contains the affected SSD has one or more RAID groups in reconstruction.
- Reduced performance isolation
  If the Flash Pool cache is not properly sized, there can be contention for the cache between the Flash Pool aggregates that are sharing it. This risk can be mitigated with proper cache sizing and QoS controls.
- Decreased management flexibility
  When you add storage to a storage pool, you increase the size of all Flash Pool caches that include one or more allocation units from that storage pool; you cannot determine how the extra capacity is distributed.

## Considerations for adding SSDs to an existing storage pool versus creating a new one

You can increase the size of your SSD cache in two ways—by adding SSDs to an existing SSD storage pool or by creating a new SSD storage pool. The best method for you depends on your configuration and plans for the storage.

The choice between creating a new storage pool or adding storage capacity to an existing one is similar to deciding whether to create a new RAID group or add storage to an existing one:

- If you are adding a large number of SSDs, creating a new storage pool provides more flexibility because you can allocate the new storage pool differently from the existing one.
- If you are adding only a few SSDs, and increasing the RAID group size of your existing Flash Pool caches is not an issue, then adding SSDs to the existing storage pool keeps your spare and parity costs lower, and automatically allocates the new storage.

If your storage pool is providing allocation units to Flash Pool aggregates whose caches have different RAID types, and you expand the size of the storage pool beyond the maximum RAID4 RAID group size, the newly added partitions in the RAID4 allocation units are unused.

## Why you add disks to storage pools

You can add SSDs to an existing storage pool and increase its cache size. When you add SSDs to a storage pool that has allocation units already allocated to Flash Pool aggregates, you increase the cache size of each of those aggregates and the total cache of the storage pool.

If the allocation units of the storage pool are not yet allocated, adding SSDs to that storage pool does not affect the SSD cache size.

When you add SSDs to an existing storage pool, the SSDs must be owned by one node or the other of the same HA pair that already owned the existing SSDs in the storage pool. You can add SSDs that are owned by either node of the HA pair.

## How storage pool works

A *storage pool* is a collection of SSDs. You can combine SSDs to create a storage pool, which enables you to share the SSDs and SSD spares across multiple Flash Pool aggregates, at the same time.

Storage pools consist of allocation units, which you can use to provide SSDs and SSD spares to aggregates or to increase the existing SSD size.

After you add an SSD to a storage pool, you can no longer use the SSD as an individual disk. You must use the storage pool to assign or allocate the storage provided by the SSD.

### Related tasks

*Provisioning storage by creating a Flash Pool aggregate manually* on page 40
You can use System Manager to create a Flash Pool aggregate manually, or to convert an existing HDD aggregate to a Flash Pool aggregate by adding SSDs. When you create a new HDD aggregate, you can provision an SSD cache to it and create a Flash Pool aggregate.

*Provisioning cache by adding SSDs* on page 119
You can use System Manager to add SSDs as either storage pools or dedicated SSDs to provision cache. By adding SSDs, you can convert a non-root aggregate or a root aggregate that does not

contain partitioned disks to a Flash Pool aggregate, or increase the cache size of an existing Flash Pool aggregate.

## Storage Pools window

You can use the Storage Pools window to create, display, and manage a dedicated cache of SSDs, also known as *storage pools*. These storage pools can be associated with a non-root aggregate to provide SSD cache and with a Flash Pool aggregate to increase its size.

This page is not available for a cluster containing nodes with All Flash Optimized personality.

- *Command buttons* on page 144
- *Storage pools list* on page 144
- *Details tab* on page 145

### Command buttons

**Create**

Opens the Create Storage Pool dialog box, which enables you to create a storage pool.

**Add Disks**

Opens the Add Disks dialog box, which enables you to add cache disks to a storage pool.

**Delete**

Deletes the selected storage pool.

**Refresh**

Updates the information in the window.

### Storage pools list

**Name**

Displays the name of the storage pool.

**Total Cache**

Displays the total cache size of the storage pool.

**Spare Cache**

Displays the available spare cache size of the storage pool.

**Used Cache (%)**

Displays the percentage of used cache size of the storage pool.

**Allocation Unit**

Displays the minimum allocation unit of the total cache size that you can use to increase the size of your storage pool.

**Owner**

Displays the name of the HA pair or the node with which the storage pool is associated.

**State**

Displays the state of the storage pool, which can be Normal, Degraded, Creating, Deleting, Reassigning, or Growing.

**Is Healthy**

Displays whether storage pool is healthy or not.

### Details tab

Displays detailed information about the selected storage pool, such as the name, health, storage type, disk count, total cache, spare cache, used cache size (in percent), and allocation unit. The tab also displays the names of the aggregates that are provisioned by the storage pool.

### Disks tab

Displays detailed information about the disks in the selected storage pool, such as the names, disk types, useable size, and total size.

#### Related tasks

*Adding disks to a storage pool* on page 140
You can add SSDs to an existing storage pool and increase its cache size by using System Manager.

*Creating a storage pool* on page 139
A storage pool is a collection of SSDs (cache disks). You can use System Manager to combine SSDs to create a storage pool, which enables you to share the SSDs and SSD spares between an HA pair for allocation to two or more Flash Pool aggregates at the same time.

*Deleting storage pools* on page 140
You might want to delete a storage pool when the cache of the storage pool is not optimal or when it is no longer used by any aggregate or Flash Pool aggregate. You can delete a storage pool by using the Delete Storage Pool dialog box in System Manager.

# Disks

You can use System Manager to manage disks.

#### Related information

*Disk and aggregate management*

*FlexArray virtualization installation requirements and reference*

*ONTAP concepts*

## Reassigning disks to nodes

You can use System Manager to reassign the ownership of spare disks from one node to another node to increase the capacity of an aggregate or storage pool.

### About this task

- You can reassign disks if the following conditions are true:
  - The container type of the selected disks must be "spare" or "shared".
  - The disks must be connected to nodes in an HA configuration.
  - The disks must be visible to the node.
- You *cannot* reassign a disk if the following conditions are true:
  - The container type of the selected disk is "shared", and the data partition is not spare.
  - The disk is associated with a storage pool.
- You cannot reassign the data partition of shared disks if storage failover is not enabled on the nodes that are associated with the shared disks.
- For partition disks, you can reassign only the data partition of the disks.
- For MetroCluster configurations, you cannot use System Manager to reassign disks.
  You must use the command-line interface to reassign disks for MetroCluster configurations.

### Steps

1. Click **Storage** > **Aggregates & Disks** > **Disks**.

2. In the **Disks** window, select the **Inventory** tab.
3. Select the disks that you want to reassign, and then click **Assign**.
4. In the **Warning** dialog box, click **Continue**.
5. In the **Assign Disks** dialog box, select the node to which you want to reassign the disks.
6. Click **Assign**.

## Viewing disk information

You can use the Disks window in System Manager to view the name, size, and container details of disks along with graphical information about capacity disks and cache disks.

### Steps

1. Click **Storage** > **Aggregates & Disks** > **Disks**.
2. Select the disk that you want to view information about from the displayed list of disks.
3. Review the disk details.

### Related reference

*Disks window* on page 149

You can use the Disks window to view all the disks in your storage system.

## How ONTAP reports disk types

ONTAP associates a type with every disk. ONTAP reports some disk types differently than the industry standards; you should understand how ONTAP disk types map to industry standards to avoid confusion.

When ONTAP documentation refers to a disk type, it is the type used by ONTAP unless otherwise specified. *RAID disk types* denote the role that a specific disk plays for RAID. RAID disk types are not related to ONTAP disk types.

For a specific configuration, the disk types that are supported depend on the storage system model, the shelf type, and the I/O modules that are installed in the system.

The following tables show how ONTAP disk types map to industry standard disk types for the SAS and FC storage connection types, and for storage arrays.

### SAS-connected storage

| ONTAP disk type | Disk class | Industry standard disk type | Description |
|---|---|---|---|
| BSAS | Capacity | SATA | Bridged SAS-SATA disks with added hardware to enable them to be plugged into a SAS-connected storage shelf |
| FSAS | Capacity | NL-SAS | Near Line SAS |
| MSATA | Capacity | SATA | SATA disk in multi-disk carrier storage shelf |
| SAS | Performance | SAS | Serial-Attached SCSI |
| SSD | Ultra-performance | SSD | Solid-state drives |

### FC-connected storage

| ONTAP disk type | Disk class | Industry standard disk type |
|---|---|---|
| ATA | Capacity | SATA |
| FCAL | Performance | FC |

### Storage arrays

| ONTAP disk type | Disk class | Industry standard disk type | Description |
|---|---|---|---|
| LUN | N/A | LUN | Logical storage device that is backed by storage arrays and used by ONTAP as a disk<br>These LUNs are referred to as *array LUNs* to distinguish them from the LUNs that ONTAP serves to clients. |

### Related information

*NetApp Hardware Universe*

*NetApp Technical Report 3437: Storage Subsystem Resiliency Guide*

## Minimum number of hot spares required for disks

Having insufficient spares increases the risk of a disk failure with no available spare, resulting in a degraded RAID group. A spare disk is also required to provide important information (a *core file*) to technical support in case of a controller disruption.

MSATA disks, or disks in a multi-disk carrier, should have four hot spares during steady state operation, and you should never allow the number of MSATA hot spares to dip below two.

For RAID groups composed of SSDs, you should have at least one spare disk.

For all other ONTAP disk types, you should have at least one matching or appropriate hot spare available for each kind of disk installed in your storage system. However, having two available hot spares for all disks provides the best protection against disk failure. Having at least two available hot spares provides the following benefits:

- When you have two or more hot spares for a data disk, ONTAP can put that disk into the maintenance center if required.

  ONTAP uses the maintenance center to test suspect disks and to take offline any disk that shows problems.
- Having two hot spares means that when a disk fails, you still have a spare disk available if another disk fails before you replace the first failed disk.

A single spare disk can serve as a hot spare for multiple RAID groups. However, if any disk in those RAID groups fails, then no spare disk is available for any future disk failures or for a core file until the spare disk is replaced. Therefore, it is a best practice to have more than one spare.

## Spare requirements for multi-disk carrier disks

Maintaining the proper number of spares for disks in multi-disk carriers is critical for optimizing storage redundancy and minimizing the amount of time that ONTAP must spend copying disks to achieve an optimal disk layout.

You must maintain a minimum of two hot spares for multi-disk carrier disks at all times. To support the use of the Maintenance Center and to avoid issues caused by multiple concurrent disk failures, you should maintain at least four hot spares for steady state operation, and replace failed disks promptly.

If two disks fail at the same time with only two available hot spares, ONTAP might not be able to swap the contents of both the failed disk and its carrier mate to the spare disks. This scenario is called a *stalemate*. If this happens, you are notified through EMS messages and AutoSupport messages. When the replacement carriers become available, you must follow the instructions that are provided by the EMS messages or you must contact technical support to recover from the stalemate.

## Shelf configuration requirements for multi-disk carrier storage shelves

You can combine multi-disk carrier disk shelves with single-disk carrier disk shelves (standard disk shelves) on the same storage system and within in the same stack.

## How to determine when it is safe to remove a multi-disk carrier

Removing a multi-disk carrier before it is safe to do so can result in one or more RAID groups becoming degraded, or possibly even a storage disruption. System Manager enables you to determine when it is safe to remove a multi-disk carrier.

When a multi-disk carrier has to be replaced, the following events must have occurred before you can remove the carrier safely:

- An AutoSupport message must have been logged indicating that the carrier is ready to be removed.
- An EMS message must have been logged indicating that the carrier is ready to be removed.
- The state of both disks in the carrier must be displayed as **`broken`** in the Disks window.

  You must remove the disks only after the carrier mate of a failed disk is evacuated. You can click Details to view the disk evacuation status in the Properties tab of the Disks window.
- The fault LED (amber) on the carrier must be lit continuously indicating that it is ready for removal.
- The activity LED (green) must be turned off indicating there is no disk activity.
- The shelf digital display only shows the shelf ID number.

**Attention:** You cannot reuse the carrier mate of a failed disk. When you remove a multi-disk carrier that contains a failed disk, you must replace it with a new carrier.

## Considerations for sizing RAID groups

Configuring an optimum RAID group size requires a trade-off of factors. You must decide which factors—speed of RAID rebuild, assurance against risk of data loss due to drive failure, optimizing I/O performance, and maximizing data storage space—are most important for the aggregate that you are configuring.

When you create larger RAID groups, you maximize the space available for data storage for the same amount of storage used for parity (also known as the "parity tax"). On the other hand, when a disk fails in a larger RAID group, reconstruction time is increased, impacting performance for a longer period of time. In addition, having more disks in a RAID group increases the probability of a multiple disk failure within the same RAID group.

### HDD or array LUN RAID groups

You should follow these guidelines when sizing your RAID groups composed of HDDs or array LUNs:

- All RAID groups in an aggregate should have a similar number of disks.
  The RAID groups do not have to be exactly the same size, but you should avoid having any RAID group that is less than one half the size of other RAID groups in the same aggregate when possible.
- The recommended range of RAID group size is between 12 and 20.
  The reliability of performance disks can support a RAID group size of up to 28, if needed.
- If you can satisfy the first two guidelines with multiple RAID group sizes, you should choose the larger size.

### SSD RAID groups in Flash Pool aggregates

The SSD RAID group size can be different from the RAID group size for the HDD RAID groups in a Flash Pool aggregate. Usually, you should ensure that you have only one SSD RAID group for a Flash Pool aggregate, to minimize the number of SSDs required for parity.

### SSD RAID groups in SSD aggregates

You should follow these guidelines when sizing your RAID groups composed of SSDs:

- All RAID groups in an aggregate should have a similar number of drives.
  The RAID groups do not have to be exactly the same size, but you should avoid having any RAID group that is less than one half the size of other RAID groups in the same aggregate when possible.
- For RAID-DP, the recommended range of RAID group size is between 20 and 28.

## Disks window

You can use the Disks window to view all the disks in your storage system.

- *Command buttons* on page 149
- *Summary* on page 149
- *Inventory* on page 150
- *Inventory details area* on page 151

### Command buttons

**Assign**

Assigns or reassigns the ownership of the disks to a node.

This button is enabled only if the container type of the selected disks is unassigned, spare, or shared.

**Zero Spares**

Erases all the data, and formats the spare disks and array LUNs.

**Refresh**

Updates the information in the window.

### Tabs

### Summary

Displays detailed information about the disks in the cluster, including the size of the spare disks and assigned disks. The tab also graphically displays information about spare disks, aggregates, and root aggregates for HDDs and information about spare disks, disks in a storage pool, aggregates, Flash Pool aggregates, and root aggregates for cache disks (SSDs).

The HDD panel is not displayed for systems with All Flash Optimized personality.

The details panel provides additional information about partitioned and unpartitioned spare disks (disk type, node, disk size, RPM, checksum, number of available disks, and spare capacity), in tabular format.

### Inventory

**Name**

Displays the name of the disk.

**Container Type**

Displays the purpose for which the disk is used. The possible values are Aggregate, Broken, Foreign, Label Maintenance, Maintenance, Shared, Spare, Unassigned, Volume, Unknown, and Unsupported.

**Partition Type**

Displays the partition type of the disk.

**Node Name**

Displays the name of the node that contains the aggregate.

This field is available only at the cluster level.

**Home owner**

Displays the name of the home node to which this disk is assigned.

**Current owner**

Displays the name of the node that currently owns this disk.

**Root owner**

Displays the name of the node that currently owns the root partition of this disk.

**Data Owner**

Displays the name of the node that currently owns the data partition of this disk.

**Data1 Owner**

Displays the name of the node that currently owns the data1 partition of the disk.

**Data2 Owner**

Displays the name of the node that currently owns the data2 partition of the disk.

**Storage Pool**

Displays the name of the storage pool with which the disk is associated.

**Type**

Displays the type of the disk.

**Firmware Version**

Displays the firmware version of the disk.

**Model**

Displays the model of the disk.

**RPM**

Displays the effective speed of the disk drive when the option `raid.mix.hdd.rpm.capacity` is enabled, and displays the actual speed of the disk drive when the option `raid.mix.hdd.rpm.capacity` is disabled.

This field is not applicable to SSDs.

**Effective Size**

Displays the usable space available on the disk.

**Physical Space**

Displays the total physical space of the disk.

**Shelf**

Displays the shelf on which the physical disks are located.

This field is hidden by default.

**Bay**

Displays the bay within the shelf for the physical disk.

This field is hidden by default.

**Pool**

Displays the name of the pool to which the selected disk is assigned.

This field is hidden by default.

**Checksum**

Displays the type of the checksum.

This field is hidden by default.

**Carrier ID**

Specifies information about disks that are located within the specified multi-disk carrier. The ID is a 64-bit value.

This field is hidden by default.

### Inventory details area

The area below the inventory tab displays detailed information about the selected disk, including information about the aggregate or volume (if applicable), vendor ID, zeroing state (in percent), serial number of the disk, and error details in case of a broken disk. For shared disks, the Inventory details area displays the names of all the aggregates, including the root and the non-root aggregates.

### Related tasks

*Viewing disk information* on page 146

You can use the Disks window in System Manager to view the name, size, and container details of disks along with graphical information about capacity disks and cache disks.

# Array LUNs

You can use System Manager to assign array LUNs to an existing aggregate and manage array LUNs.

### Related information

*FlexArray virtualization installation requirements and reference*

## Assigning array LUNs

You can use System Manager to assign unassigned array LUNs to an existing aggregate to increase the size of the aggregate.

### About this task

- You can assign array LUNs if the following conditions are true:
  - The container type of the selected array LUNs must be "unassigned".

  - The disks must be connected to nodes in an HA pair.
  - The disks must be visible to the node.
- For MetroCluster configurations, you cannot use System Manager to assign array LUNs as spares.

  You must use the command-line interface instead.

### Steps

1. Click **Storage** > **Aggregates & Disks** > **Array LUNs**.
2. Select the array LUNs, and then click **Assign**.
3. In the **Assign Array LUNs** dialog box, select the node to which you want to assign the array LUNs.
4. Click **Assign**.

## Reassigning spare array LUNs to nodes

You can use System Manager to reassign the ownership of spare array LUNs from one node to another to increase the capacity of an aggregate.

### About this task

- You can reassign array LUNs if the following conditions are true:
  - The container type of the selected array LUNs must be "spare".
  - The disks must be connected to nodes in an HA pair.
  - The disks must be visible to the node.
- For MetroCluster configurations, you cannot use System Manager to reassign array LUNs as spares.

  You must use the command-line interface instead.

### Steps

1. Click **Storage** > **Aggregates & Disks** > **Array LUNs**.
2. Select the spare array LUNs that you want to reassign, and then click **Assign**.
3. In the **Warning** dialog box, click **Continue**.
4. In the **Assign Array LUNs** dialog box, select the node to which you want to reassign the spare array LUNs.
5. Click **Assign**.

## Zeroing spare array LUNs

You can use System Manager to erase all the data and to format the spare array LUNs by writing zeros to the array LUNs. These array LUNs can then be used in new aggregates.

### About this task

When you zero the spare array LUNs, all the spares in the cluster, including disks, are zeroed. You can zero the spare array LUNs for a specific node or for the entire cluster.

### Steps

1. Click **Storage** > **Aggregates & Disks** > **Array LUNs**.
2. Click **Zero Spares**.
3. In the **Zero Spares** dialog box, select a node or "All nodes" from which you want to zero the array LUNs.
4. Select the **Zero all non-zeroed spares** check box to confirm the zeroing operation.
5. Click **Zero Spares**.

## About disks and array LUNs

A disk is the basic unit of storage for storage systems that use ONTAP software to access native disk shelves. An array LUN is the basic unit of storage that third-party storage arrays provide to storage systems that run ONTAP software.

ONTAP software enables you to assign ownership to your disks and array LUNs, and to add them to an aggregate. ONTAP software also provides a number of ways to manage your disks, including removing them, replacing them, and sanitizing them. Because array LUNs are provided by the third-party storage array, you use the third-party storage array for all other management tasks for array LUNs.

You can create an aggregate using either disks or array LUNs. After you have created the aggregate, you manage it using ONTAP software in exactly the same way, whether it was created from disks or array LUNs.

## Array LUNs window

The Array LUNs window enables you to assign ownership to your array LUNs and to add them to an aggregate.

The Array LUNs link in the left navigation pane is displayed only if there are any spare array LUNs, or if the V_StorageAttach license is installed.

- *Command buttons* on page 153
- *Array LUN list* on page 153
- *Details area* on page 154

### Command buttons

**Assign**

Enables you to assign or reassign the ownership of array LUNs to a node.

**Zero Spares**

Erases all the data, and formats the spare array LUNs and disks.

**Refresh**

Updates the information in the window.

### Array LUN list

Displays information such as the name, state, and vendor for each array LUN.

**Name**

Specifies the name of the array LUN.

**State**

Specifies the state of the array LUN.

**Vendor**

Specifies the name of the vendor.

**Used Space**

Specifies the space used by the array LUN.

**Total Size**

Specifies the size of the array LUN.

**Container**

Specifies the aggregate to which the array LUN belongs.

**Node name**

Specifies the name of the node to which the array LUN belongs.

**Home owner**

Displays the name of the home node to which the array LUN is assigned.

**Current owner**

Displays the name of the node that currently owns the array LUN.

**Array name**

Specifies the name of the array.

**Pool**

Displays the name of the pool to which the selected array LUN is assigned.

### Details area

The area below the Array LUNs list displays detailed information about the selected array LUN.

# Nodes

You can use System Manager to view the details of the nodes in the cluster.

## Initializing the ComplianceClock time

You can use System Manager to initialize the ComplianceClock time to the current cluster time. You must initialize the ComplianceClock time in order to create SnapLock aggregates.

### Before you begin

The SnapLock license must be installed.

### About this task

You cannot modify or stop the ComplianceClock time after it is initialized.

### Steps

1. Click **Storage** > **Nodes**.
2. Select the node, and then click **Initialize ComplianceClock**.
3. In the **Initialize ComplianceClock** dialog box, click **Yes** to initialize the ComplianceClock time to the current cluster time.

## Nodes window

You can use the Nodes window to view the details of the nodes in a cluster.

- *Command buttons*
- *Nodes list*

### Command buttons

**Initialize ComplianceClock**

Initializes the ComplianceClock of the selected node to the current value of the system clock.

**Refresh**

Updates the information in the window.

### Nodes list

**Name**

Displays the name of the node.

**State**

Displays the state of the node (whether the node is up or down).

**Up Time**

Displays the duration for which the node is up.

**ONTAP Version**

Displays the ONTAP version that is installed on the node.

**Model**

Displays the platform model number of the node.

**System ID**

Displays the ID of the node.

**Serial No**

Displays the serial number of the node.

**All Flash Optimized**

Displays whether the node has an All Flash Optimized personality.

### Details area

Displays detailed information about the selected node.

**Details tab**

Displays information related to the selected node such as the name of the node, the state of the node, and the duration for which the node is up.

**Performance tab**

Displays the throughput, IOPS, and latency of the selected node.

Changing the client time zone or the cluster time zone impacts the performance metrics graphs. You should refresh your browser to view the updated graphs.

# Flash Cache

You can use System Manager to manage Flash Cache.

## Enabling or disabling Flash Cache

You can enable or disable the WAFL external cache functionality for a storage system that has a PAM II card or Flash Cache module installed by using System Manager. You can enable Flash Cache based on the workload requirements of your storage system.

### Steps

1. Click **Configuration** > **Flash Cache Module**.
2. Select the node.
3. Click **Enable** or **Disable**, as required.

## How Flash Cache works

Using Flash Caches improves the performance of a storage system.

You can configure Flash Cache and disks based on the workload requirements of a storage system. By determining the read workload (number of read operations) served by Flash Cache and disks, you can analyze the performance of the storage system.

Flash Cache does not contain any data during storage system boot or when control is returned to the storage system after a takeover event. Therefore, disks serve all the data read requests of the storage system.

The Flash Cache module is slowly populated with data when data read requests are served. Because the data read requests served by Flash Cache are faster than those served by the disks, the performance of the storage system improves.

Data read requests served by the Flash Cache module replace the data read requests served by the disks and, therefore, the performance improvement in the storage system is directly related to the disk reads that are replaced. To understand the impact of Flash Cache on storage system performance, you must view the read workload graph when the Flash Cache contains data.

## Flash Cache window

You can use the Flash Cache window to enable or disable Flash Cache for a storage system that has a Flash Cache module installed. You can also view the read workload statistics.

This Flash Cache window is not available for a cluster containing nodes with All Flash Optimized personality.

### Command buttons

**Enable/Disable**

Enables or disables Flash Cache.

### Flash Cache Read Workload

Displays a graph specifying the rate of the read workload that is served by the disks and the Flash Cache module, which indicates the performance of the storage system.

### Details area

Displays information about the system read latency (in seconds), the caching mode that specifies the caching configuration, the state of Flash Cache (enabled or disabled), and the size of the Flash Cache (in GB). If there are multiple Flash Cache cards, the total cache size from all of the cards is displayed.

**Note:** The Flash Cache size that is displayed differs from the actual Flash cache size for the following reasons:

- System Manager reports only the usable capacity that is provided by ONTAP.
- A portion of the total SSD capacity is reserved for storing metadata.

# Events

You can use System Manager to view the event log and event notifications.

## Events window

You can use the Events window to view the event log and event notifications.

### Command buttons

**Refresh**

Updates the information in the window.

### Events list

**Time**

Displays the time when the event occurred.

**Node**

Displays the node and the cluster on which the event occurred.

**Severity**

Displays the severity of the event. The possible severity levels are:

- Emergency
  Specifies that the event source unexpectedly stopped, and the system experienced unrecoverable data loss. You must take corrective action immediately to avoid extended downtime.
- Alert
  Specifies that the event source has an alert, and action must be taken to avoid downtime.
- Critical
  Specifies that the event source is critical, and might lead to service disruption if corrective action is not taken immediately.
- Error
  Specifies that the event source is still performing, and a corrective action is required to avoid service disruption.
- Warning
  Specifies that the event source experienced an occurrence that you must be aware of. Events of this severity might not cause service disruption; however, corrective action might be required.
- Notice
  Specifies that the event source is normal, but the severity is a significant condition that you must be aware of.
- Informational
  Specifies that the event source has an occurrence that you must be aware of. No corrective action might be required.
- Debug
  Specifies that the event source includes a debugging message.

By default, the alert severity type, emergency severity type, and the error severity type are displayed.

**Source**

Displays the source of the event.

**Event**

Displays the description of the event.

### Details area

Displays the event details, including the event description, message name, sequence number, message description, and corrective action for the selected event.

# System alerts

You can use System Manager to monitor different parts of a cluster.

### Related information

*System administration*

## Acknowledging system health alerts

You can use System Manager to acknowledge and respond to system health alerts for subsystems. You can use the information displayed to take the recommended action and correct the problem reported by the alert.

### Steps

1. Click **Events & Jobs** > **System Alerts**.
2. In the **System Alerts** window, click the arrow icon next to the name of subsystem.
3. Select the alert that you want to acknowledge, and then click **Acknowledge**.
4. Type your name, and then click **Acknowledge**.

### Related reference

*System Alerts window* on page 160

You can use the System Alerts window to learn more about system health alerts. You can also acknowledge, delete, and suppress alerts from the window.

## Suppressing system health alerts

You can use System Manager to suppress system health alerts that do not require any intervention from you.

### Steps

1. Click **Events & Jobs** > **System Alerts**.
2. In the **System Alerts** window, click the arrow icon next to the name of subsystem.
3. Select the alert that you want to suppress, and then click **Suppress**.
4. Type your name, and then click **Suppress**.

### Related reference

*System Alerts window* on page 160

You can use the System Alerts window to learn more about system health alerts. You can also acknowledge, delete, and suppress alerts from the window.

## Deleting system health alerts

You can use System Manager to delete system health alerts to which you have already responded.

### Steps

1. Click **Events & Jobs** > **System Alerts**.
2. In the **System Alerts** window, click the arrow icon next to the name of subsystem.
3. Select the alert that you want to delete, and then click **Delete**.
4. Click **OK**.

### Related reference

*System Alerts window* on page 160

You can use the System Alerts window to learn more about system health alerts. You can also acknowledge, delete, and suppress alerts from the window.

## Available cluster health monitors

There are several health monitors that monitor different parts of a cluster. Health monitors help you to recover from errors within ONTAP systems by detecting events, sending alerts to you, and deleting events as they clear.

| Health monitor name (identifier) | Subsystem name (identifier) | Purpose |
|---|---|---|
| Cluster switch (cluster-switch) | Switch (Switch-Health) | Monitors cluster network switches and management network switches for temperature, utilization, interface configuration, redundancy (cluster network switches only), and fan and power supply operation. The cluster switch health monitor communicates with switches through SNMP. SNMPv2c is the default setting. |
| MetroCluster Fabric | Switch | Monitors the MetroCluster configuration back-end fabric topology and detects misconfigurations such as incorrect cabling and zoning, and ISL failures. |
| MetroCluster Health | Interconnect, RAID, and storage | Monitors FC-VI adapters, FC initiator adapters, left-behind aggregates and disks, and inter-cluster ports |
| Node connectivity (node-connect) | CIFS nondisruptive operations (CIFS-NDO) | Monitors SMB connections for nondisruptive operations to Hyper-V applications. |
| | Storage (SAS-connect) | Monitors shelves, disks, and adapters at the node level for appropriate paths and connections. |
| System | not applicable | Aggregates information from other health monitors. |
| System connectivity (system-connect) | Storage (SAS-connect) | Monitors shelves at the cluster level for appropriate paths to two HA clustered nodes. |

## Ways to respond to system health alerts

When a system health alert occurs, you can acknowledge it, learn more about it, repair the underlying condition, and prevent it from occurring again.

When a health monitor raises an alert, you can respond in any of the following ways:

- Get information about the alert, which includes the affected resource, alert severity, probable cause, possible effect, and corrective actions.
- Get detailed information about the alert, such as the time when the alert was raised and whether anyone else has acknowledged the alert already.
- Get health-related information about the state of the affected resource or subsystem, such as a specific shelf or disk.
- Acknowledge the alert to indicate that someone is working on the problem, and identify yourself as the "Acknowledger."
- Resolve the problem by taking the corrective actions provided in the alert, such as fixing cabling to resolve a connectivity problem.
- Delete the alert, if the system did not automatically clear it.
- Suppress an alert to prevent it from affecting the health status of a subsystem.

Suppressing is useful when you understand a problem. After you suppress an alert, it can still occur, but the subsystem health displays as "ok-with-suppressed." when the suppressed alert occurs.

## System Alerts window

You can use the System Alerts window to learn more about system health alerts. You can also acknowledge, delete, and suppress alerts from the window.

### Command buttons

**Acknowledge**

Enables you to acknowledge the selected alert to indicate that the problem is being addressed and identifies the person who clicks the button as the "Acknowledger."

**Suppress**

Enables you to suppress the selected alert to prevent the system from notifying you about the same alert again and identifies you as the "Suppressor."

**Delete**

Deletes the selected alert.

**Refresh**

Updates the information in the window.

### Alerts list

**SubSystem (No. of Alerts)**

Displays the name of the subsystem, such as the SAS connection, switch health, CIFS NDO, or MetroCluster, for which the alert is generated.

**Alert ID**

Displays the alert ID.

**Node**

Displays the name of the node for which the alert is generated.

**Severity**

Displays the severity of the alert as Unknown, Other, Information, Degraded, Minor, Major, Critical, or Fatal.

**Resource**

Displays the resource that generated the alert, such as a specific shelf or disk.

**Time**

Displays the time when the alert was generated.

### Details area

The details area displays detailed information about the alert, such as the time when the alert was generated and whether the alert has been acknowledged. The area also includes information about the probable cause and possible effect of the condition generated by the alert, and the recommended actions to correct the problem reported by the alert.

### Related tasks

*Acknowledging system health alerts* on page 158

You can use System Manager to acknowledge and respond to system health alerts for subsystems. You can use the information displayed to take the recommended action and correct the problem reported by the alert.

*Suppressing system health alerts* on page 158

You can use System Manager to suppress system health alerts that do not require any intervention from you.

*Deleting system health alerts* on page 158

You can use System Manager to delete system health alerts to which you have already responded.

# AutoSupport notifications

You can use System Manager to configure AutoSupport notifications that help you to monitor your storage system health.

## Setting up AutoSupport notifications

You can use the Edit AutoSupport Settings dialog box in System Manager to set up AutoSupport notifications by specifying an email address from which email notifications are sent and adding multiple email host names.

### Steps

1. Click > **AutoSupport**.
2. Select the node, and then click **Edit**.
3. In the **Email Recipient** tab, type the email address from which email notifications are sent, specify the email recipients and the message content for each email recipient, and add the mail hosts.

   You can add up to five email addresses of the host names.
4. In the **Others** tab, select a transport protocol for delivering the email messages from the drop-down list and specify the HTTP or HTTPS proxy server details.
5. Click **OK**.
6. Verify that configuration you have set for AutoSupport notification is set up correctly in the **AutoSupport** window.

## Enabling or disabling AutoSupport settings

You can enable or disable AutoSupport settings on your storage system by using System Manager. AutoSupport messages enable you to monitor your storage system health or send notifications to technical support and your internal support organization.

### About this task

The AutoSupport option is enabled by default.

### Steps

1. Click > **AutoSupport**.
2. Select the node, and then click **Enable** or **Disable**.
3. Click **OK**.
4. Verify that the AutoSupport status correctly displays the change you made.

## Adding AutoSupport email recipients

You can use the **Email Recipient** tab of the Edit AutoSupport Settings dialog box in System Manager to add email addresses of the recipients of AutoSupport notifications.

### Steps

1. Click > **AutoSupport**.
2. Select the node, and then click **Edit**.

3. In the **Email Recipient** tab, type the address of the email recipient, specify whether the recipient receives a full message or a short message, and then click **Add**.
4. Click **OK**.
5. Verify that the details you specified are displayed in the **AutoSupport** window.

## Testing AutoSupport settings

You can use the AutoSupport Test dialog box in System Manager to test that you have configured the AutoSupport settings correctly.

### Steps

1. Click ⚙ > **AutoSupport**.
2. Select the node, and then click **Test**.
3. In the **AutoSupport Test** dialog box, enter the AutoSupport subject text "Test AutoSupport" or any text that notifies the recipients that you are testing the AutoSupport settings.
4. Click **Test**.
   An email message with the subject "Test AutoSupport" or the text that you typed in the **AutoSupport subject** field is sent to the specified recipients.

## Generating AutoSupport data

You can use System Manager to generate AutoSupport data for a single node or multiple nodes to monitor their health and to send notifications to technical support.

### Steps

1. Click ⚙ > **AutoSupport**.
2. Select the node, and then click **AutoSupport Request** > **Generate AutoSupport**.
   By default, the AutoSupport data is generated for all nodes.
3. In the **Generate AutoSupport** dialog box, perform the following steps:
   a. If you want to generate AutoSupport data for a specific node, clear the **Generate Autosupport data for all nodes** check box, and then select the node.
   b. Type the case number.
4. Click **Generate**.
5. In the **Confirmation** dialog box, click **OK**.

## Viewing AutoSupport summary

System Manager enables you to view the status and details of all the previous AutoSupport data in order to review the data that has been sent to technical support. You can also view the information to understand the health and performance of your storage system.

### Steps

1. Click ⚙ > **AutoSupport**.
2. Select the node, and then click **AutoSupport Request** > **View Previous Summary**.
   The AutoSupport data for all the nodes is displayed.
3. Click **OK**.

## AutoSupport severity types

AutoSupport messages have severity types that help you understand the purpose of each message —for example, to draw immediate attention to an emergency problem, or only to provide information.

Messages have one of the following severities:

- **Alert**: Alert messages indicate that a next-higher level event might occur if you do not take some action.

  You must take an action against alert messages within 24 hours.
- **Emergency**: Emergency messages are displayed when a disruption has occurred.

  You must take an action against emergency messages immediately.
- **Error**: Error conditions indicate what might happen if you ignore.
- **Notice**: Normal but significant condition.
- **Info**: Informational message provides details about the issue, which you can ignore.
- **Debug**: Debug-level messages provide instructions you should perform.

If your internal support organization receives AutoSupport messages through email, the severity appears in the subject line of the email message.

## AutoSupport window

The AutoSupport window enables you to view the current AutoSupport settings for your system. You can also change your system's AutoSupport settings.

### Command buttons

**Enable**

Enables AutoSupport notification.

**Disable**

Disables AutoSupport notification.

**Edit**

Opens the Edit AutoSupport Settings dialog box, which enables you to specify an email address from which email notifications are sent and to add multiple email addresses of the host names.

**Test**

Opens the AutoSupport Test dialog box, which enables you to generate an AutoSupport test message.

**AutoSupport Request**

Provides the following AutoSupport requests:

**Generate AutoSupport**

Generates AutoSupport data for a selected node or all nodes.

**View Previous Summary**

Displays the status and details of all the previous AutoSupport data.

**Refresh**

Updates the information in the window.

### Details area

The details area displays AutoSupport setting information such as the node name, AutoSupport status, transport protocol used, and name of the proxy server.

#### Related tasks

*Setting up a support page* on page 26

Setting up the support page completes the cluster setup, and involves setting up the AutoSupport messages and event notifications, and for single-node clusters, configuring system backup.

# Jobs

You can use System Manager to manage job tasks such as displaying job information and monitoring the progress of a job.

## Jobs

*Jobs* are asynchronous task and typically long-running volume operations, such as copying, moving, or mirroring data. Jobs are placed in a job queue and are run when resources are available. The cluster administrator can perform all the tasks related to job management.

A job can be one of the following categories:

- A *server-affiliated* job is placed in queue by the management framework to be run in a specific node.
- A *cluster-affiliated* job is placed in queue by the management framework to be run in any node in the cluster.
- A *private* job is specific to a node and does not use the replicated database (RDB) or any other cluster mechanism.

  You require the advanced privilege level or higher to run the commands to manage private jobs.

You can manage jobs in the following ways:

- Displaying job information, including the following:
  - Jobs on a per-node basis
  - Cluster-affiliated jobs
  - Completed jobs
  - Job history
- Monitoring a job's progress
- Displaying information about the initialization state for job managers.

You can determine the outcome of a completed job by checking the event log.

## Job window

You can use the Job window to manage job tasks such as displaying job information and monitoring the progress of a job.

### Command button

**Refresh**

Updates the information in the window.

### Tabs

**Current Jobs**

This tab displays information about the job tasks that are in progress.

**Job History**

This tab displays information about all the jobs.

### Job list

**Job ID**

Displays the ID of the job.

**Start Time**

Displays the start time of the job.

**Job Name**

Displays the name of the job.

**Node**

Displays the name of the node.

**State**

Displays the state of the job.

**Job Description**

Displays the description of the job.

**Progress**

Displays the state of the job.

**Schedule Name**

Displays the name of the schedule.

# Flash Pool statistics

You can use System Manager to view the real-time SSD tier read and write workloads for a selected Flash Pool aggregate.

## Flash Pool aggregate Statistics window

You can view the real-time SSD tier read and write workloads for a selected Flash Pool aggregate.

This page is not available for a cluster containing nodes with All Flash Optimized personality.

**Displaying Statistics for Flash Pool aggregate**

From the list of Flash Pool aggregates, you can select the Flash Pool aggregate whose statistics you want to view.

### SSD Cache Read Workload

Displays a graphical view of the total read requests that are sent to the Flash Pool aggregate in comparison with the read operations that are performed by the SSD tier.

### SSD Cache Write Workload

Displays a graphical view of the total write requests that are sent to the Flash Pool aggregate in comparison with the write operations that are performed by the SSD tier.

# Managing logical storage

You can use System Manager to manage the logical storage such as storage virtual machines (SVMs), volumes, Qtrees, protocols, policies and so on.

## Storage Virtual Machines

You can use System Manager to manage the SVMs in your cluster.

**Related information**

*SAN administration*

*ONTAP concepts*

### SVM Dashboard window

The dashboard provides a cumulative at-a-glance information about your storage virtual machine (SVM) and its performance. You can use the Dashboard window to view important information related to your SVM such as the protocols configured, the volumes that are nearing capacity, and the performance.

#### SVM Details

This window displays details about the SVM through various panels such as the Protocol Status panel, Volumes Nearing Capacity panel, Applications panel, and SVM performance panel.

**Protocol Status**

Provides an overview of the protocols that are configured for the SVM. You can click the protocol name to view the configuration.

If a protocol is not configured or if a protocol license is not available for the SVM, you can click the protocol name to configure the protocol or to add the protocol license.

**Volumes Nearing Capacity**

Displays information about the volumes that are nearing capacity utilization of 80 percent or more, and that therefore require immediate attention or corrective action.

**Applications**

Displays information about the top five applications of the SVM. You can view the top five applications based on either IOPS (from low to high or from high to low) or capacity (from low to high or from high to low). You must click the specific bar chart to view more information about the application. For capacity, the total space, used space, and available space are displayed, and for IOPS, the IOPS details are displayed.

**Note:** The used size displayed in the Applications on SVM window does not equal the used size in the CLI.

You can click **View details** to open the Applications window of the specific application. You can click **View all applications** to view all of the applications for the SVM.

The refresh interval for the Applications panel is one minute.

**SVM Performance**

Displays the performance metrics of the protocols in the SVM, including latency and IOPS.

If the information about SVM performance cannot be retrieved from ONTAP, you cannot view the respective graph. In such cases, System Manager displays the specific error message.

The refresh interval for the SVM Performance panel is 15 seconds.

## Monitoring SVMs

The dashboard in System Manager enables you to monitor the health and performance of a storage virtual machine (SVM).

### Steps

1. Click **Storage** > **SVMs**.
2. Click on the SVM that you want to monitor.
3. View the details in the dashboard panels.

## Editing SVM settings

You can use System Manager to edit the properties of storage virtual machines (SVMs), such as the name service switch, name mapping switch, and aggregate list.

### About this task

- You can edit the values of the following SVM properties:
  - Name service switch
  - Protocols that are enabled to serve data

    **Note:** The CIFS protocol that is configured on the SVM continues to serve data even when you disable the protocol on that SVM.
  - The list of aggregates that are available to create volumes

    **Note:** For FlexVol volumes, you can assign aggregates only if you have delegated administration to an SVM administrator.
- System Manager does not display the values of the name service switch and the name mapping switch for an SVM that is created through the command-line interface, or whose services are not configured and are not set to the default values by ONTAP.

  You can use the command-line interface to view the services because the Services tab is disabled.

  System Manager displays the name service switch and the name mapping switch of an SVM only when it is created by using System Manager or when the services of the SVM are set to the default values by ONTAP.

### Steps

1. Click **Storage** > **SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **Details** tab, modify the required data protocols.
4. In the **Resource Allocation** tab, choose one of the following methods to delegate volume creation:

| If you want to provision volume creation... | Then... |
|---|---|
| For all aggregates | Select the **Do not delegate volume creation** option. |
| For specific aggregates | **a.** Select the **Delegate volume creation** option.<br>**b.** Select the required aggregates for delegating volume creation. |

5. In the **Service** tab, specify the name service switch sources for the required database types and the order in which they should be consulted to retrieve name service information.

   The default values for each of the database types are as follows:

- hosts: files, dns
- namemap: files
- group: files
- netgroup: files
- passwd: files

6. Click **Save and Close**.

### Related concepts

*How ONTAP name service switch configuration works* on page 171

ONTAP stores name service configuration information in a table that is the equivalent of the `/etc/nsswitch.conf` file on UNIX systems. You must understand the function of the table and how ONTAP uses it so that you can configure it appropriately for your environment.

## Deleting SVMs

You can use System Manager to delete storage virtual machines (SVMs) that you no longer require from the storage system configuration.

### Before you begin

You must have completed the following tasks:

1. Disabled the Snapshot copies, data protection (DP) mirrors, and load-sharing (LS) mirrors for all the volumes

   **Note:** You must use the command-line interface (CLI) to disable LS mirrors.
2. Deleted all the igroups that belong to the SVM manually if you are deleting SVMs
3. Deleted all the portsets
4. Deleted all the volumes in the SVM, including the root volume
5. Unmapped the LUNs, taken them offline, and deleted them
6. Deleted the CIFS server if you are deleting SVMs
7. Deleted any customized user accounts and roles that are associated with the SVM
8. Deleted any NVMe subsystems associated with the SVM using the CLI.
9. Stopped the SVM

### About this task

When you delete SVMs, the following objects associated with the SVM are also deleted:

- LIFs, LIF failover groups, and LIF routing groups
- Export policies
- Efficiency policies

If you delete SVMs that are configured to use Kerberos, or modify SVMs to use a different Service Principal Name (SPN), the original service principal of the SVM is not automatically deleted or disabled from the Kerberos realm. You must manually delete or disable the principal. You must have the Kerberos realm administrator's user name and password to delete or disable the principal.

If you want to move data from an SVM to another SVM before you delete the first SVM, you can use the SnapMirror technology to do so.

### Steps

1. Click **Storage** > **SVMs**.
2. Select the SVM that you want to delete, and then click **Delete**.
3. Select the confirmation check box, and then click **Delete**.

## Starting SVMs

You can use System Manager to provide data access from a storage virtual machine (SVM) by starting the SVM.

### Steps

1. Click **Storage** > **SVMs**.
2. Select the SVM that you want to start, and then click **Start**.

### Result

The SVM starts serving data to clients.

## Stopping SVMs

You can use System Manager to stop a storage virtual machine (SVM) if you want to troubleshoot any issue with the SVM, delete the SVM, or stop data access from the SVM.

### Before you begin

All the clients connected to the SVM must be disconnected.

**Attention:** If any clients are connected to the SVM when you stop it, data loss might occur.

### About this task

- You cannot stop SVMs during storage failover (SFO).
- When you stop the SVM, an SVM administrator cannot log in to the SVM.

### Steps

1. Click **Storage** > **SVMs**.
2. Select the SVM that you want to stop, and then click **Stop**.

### Result

The SVM stops serving data to clients.

## Managing SVMs

storage virtual machine (SVM) administrators can administer SVMs and its resources, such as volumes, protocols, and services, depending on the capabilities assigned by the cluster administrator. SVM administrators cannot create, modify, or delete SVMs.

**Note:** SVM administrators cannot log in to System Manager.

SVM administrators might have all or some of the following administration capabilities:

- Data access protocol configuration
  SVM administrators can configure data access protocols, such as NFS, CIFS, iSCSI, and Fibre Channel (FC) protocol (Fibre Channel over Ethernet or FCoE included).
- Services configuration
  SVM administrators can configure services such as LDAP, NIS, and DNS.
- Storage management
  SVM administrators can manage volumes, quotas, qtrees, and files.
- LUN management in a SAN environment
- Management of Snapshot copies of the volume
- Monitoring SVM
  SVM administrators can monitor jobs, network connection, network interface, and the SVM health.

**Related information**

*NetApp Documentation: ONTAP 9*

## Types of SVMs

A cluster consists of four types of SVMs, which help in managing the cluster and its resources and data access to the clients and applications.

A cluster contains the following types of SVMs:

- Admin SVM
  The cluster setup process automatically creates the admin SVM for the cluster. The admin SVM represents the cluster.
- Node SVM
  A node SVM is created when the node joins the cluster, and the node SVM represents the individual nodes of the cluster.
- System SVM (advanced)
  A system SVM is automatically created for cluster-level communications in an IPspace.
- Data SVM
  A data SVM represents the data serving SVMs. After the cluster setup, a cluster administrator must create data SVMs and add volumes to these SVMs to facilitate data access from the cluster.
  A cluster must have at least one data SVM to serve data to its clients.

## Why you use SVMs

SVMs provide data access to clients regardless of the physical storage or controller, similar to any storage system. SVMs provide benefits such as nondisruptive operations, scalability, security, and unified storage.

SVMs provide the following benefits:

- Multi-tenancy
  SVM is the fundamental unit of secure multi-tenancy, which enables partitioning of the storage infrastructure so that it appears as multiple independent storage systems. These partitions isolate the data and management.
- Nondisruptive operations
  SVMs can operate continuously and nondisruptively for as long as they are needed. SVMs help clusters to operate continuously during software and hardware upgrades, addition and removal of nodes, and all administrative operations.
- Scalability
  SVMs meet on-demand data throughput and the other storage requirements.
- Security
  Each SVM appears as a single independent server, which enables multiple SVMs to coexist in a cluster while ensuring no data flows among them.
- Unified storage
  SVMs can serve data concurrently through multiple data access protocols. SVMs provide file-level data access through NAS protocols, such as CIFS and NFS, and block-level data access through SAN protocols, such as iSCSI, FC/FCoE, and NVMe. SVMs can serve data to SAN and NAS clients independently at the same time.
- Delegation of management
  SVM administrators have privileges assigned by the cluster administrator.

## How ONTAP name service switch configuration works

ONTAP stores name service configuration information in a table that is the equivalent of the `/etc/nsswitch.conf` file on UNIX systems. You must understand the function of the table and how ONTAP uses it so that you can configure it appropriately for your environment.

The ONTAP name service switch table determines which name service sources ONTAP consults in which order to retrieve information for a certain type of name service information. ONTAP maintains a separate name service switch table for each SVM.

### Database types

The table stores a separate name service list for each of the following database types:

| Database type | Defines name service sources for... | Valid sources are... |
|---|---|---|
| hosts | Converting host names to IP addresses | files, dns |
| group | Looking up user group information | files, nis, ldap |
| passwd | Looking up user information | files, nis, ldap |
| netgroup | Looking up netgroup information | files, nis, ldap |
| namemap | Mapping user names | files, ldap |

### Source types

The sources specify which name service source to use for retrieving the appropriate information.

| Specify source type... | To look up information in... | Managed by the command families... |
|---|---|---|
| files | Local source files | `vserver services name-service unix-user`<br>`vserver services name-service unix-group`<br>`vserver services name-service netgroup`<br>`vserver services name-service dns hosts` |
| nis | External NIS servers as specified in the NIS domain configuration of the SVM | `vserver services name-service nis-domain` |
| ldap | External LDAP servers as specified in the LDAP client configuration of the SVM | `vserver services name-service ldap` |
| dns | External DNS servers as specified in the DNS configuration of the SVM | `vserver services name-service dns` |

Even if you plan to use NIS or LDAP for both data access and SVM administration authentication, you should still include **`files`** and configure local users as a fallback in case NIS or LDAP authentication fails.

### Related tasks

*Editing SVM settings* on page 167

You can use System Manager to edit the properties of storage virtual machines (SVMs), such as the name service switch, name mapping switch, and aggregate list.

## Storage Virtual Machines window

You can use the Storage Virtual Machines window to manage your storage virtual machines (SVMs) and display information about them.

You cannot manage (create, edit, delete, start, or stop) an SVM configured for disaster recovery (DR) by using System Manager. Also, you cannot view the storage objects associated with the SVM configured for disaster recovery in the application interface.

### Command buttons

**Create**

Opens the Storage Virtual Machine (SVM) Setup wizard, which enables you to create a new SVM.

**Edit**

Opens the Edit Storage Virtual Machine dialog box, which enables you to modify the properties, such as the name service switch, name mapping switch, and aggregate list, of a selected SVM.

**Delete**

Deletes the selected SVMs.

**Start**

Starts the selected SVM.

**Stop**

Stops the selected SVM.

**Manage**

Manages the storage, policies, and configuration for the selected SVM.

**Refresh**

Updates the information in the window.

### SVM list

The SVM list displays the name of each SVM and the allowed protocols on it.

You can view only data SVMs by using System Manager.

**Name**

Displays the name of the SVM.

**State**

Displays the SVM state, such as Running, Starting, Stopped, or Stopping.

**Subtype**

Displays the subtype of the SVM, which can be one of the following:

- default
  Specifies that the SVM is a data-serving SVM.
- dp-destination
  Specifies that the SVM is configured for disaster recovery.
- sync-source
  Specifies that the SVM is in the primary site of a MetroCluster configuration.
- sync-destination
  Specifies that the SVM is in the surviving site of a MetroCluster configuration.

**Allowed Protocols**

Displays the allowed protocols, such as CIFS and NFS, on each SVM.

**IPspace**

Displays the IPspace of the associated SVM.

**Volume Type**

Displays the allowed volume type, such as FlexVol volume, on each SVM.

**Configuration State**

Displays whether the configuration state of the SVM is locked or unlocked.

### Details area

The area below the SVM list displays detailed information, such as the type of volumes allowed, language, and Snapshot policy, about the selected SVM.

You can also configure the protocols that are allowed on this SVM. If you have not configured the protocols while creating the SVM, you can click the protocol link to configure the protocol.

You cannot configure protocols for an SVM configured for disaster recovery by using System Manager.

**Note:** If the FCP service is already started for the SVM, clicking the FC/FCoE link opens the Network Interfaces window.

The color indicates the status of the protocol configuration:

| Status | Description |
|---|---|
| Green | LIFs exist and the protocol is configured.<br>You can click the link to view the configuration details.<br>**Note:** Configuration might be partially completed. However, service is running. You can create the LIFs and complete the configuration from the Network Interfaces window. |
| Yellow | Indicates one of the following:<br>• LIFs exist. Service is created but is not running.<br>• LIFs exist. Service is not created.<br>• Service is created. LIFs do not exist. |
| Grey | The protocol is not configured. You can click the protocol link to configure the protocol. |
| Grey border | The protocol license has expired or is missing. You can click the protocol link to add the licenses in the Licenses page. |

### Peer Storage Virtual Machines area

Displays a list of the SVMs that are peered with the selected SVM along with details of the applications that are using the peer relationship.

# Volumes

You can use System Manager to create, edit, and delete volumes.

You can access all the volumes in the cluster by using the Volumes tab or you can access the volumes specific to an SVM by using **SVMs** > **Volumes**.

**Note:** The Volumes tab is displayed only if you have enabled the CIFS and NFS licenses.

### Related information

*ONTAP concepts*
*Logical storage management*

## Editing volume properties

You can modify volume properties such as the volume name, security style, fractional reserve, and space guarantee by using System Manager. You can modify storage efficiency settings (deduplication schedule, deduplication policy, and compression) and space reclamation settings.-

### About this task

- You can set the fractional reserve to either zero percent or 100 percent.
- Data compression is not supported on 32-bit volumes.
- For Data ONTAP 8.3.1 clusters, you can enable both inline compression and background compression for Cloud Volumes ONTAP for AWS (AWS).

  Compression is not supported for Data ONTAP Edge.
- You cannot rename a SnapLock Compliance volume.

### Steps

1. Click **Storage** > **Volumes**.
2. From the **Volumes on SVM** list, select the storage virtual machine (SVM) in which the volume that you want to edit resides.
3. Select the volume that you want to modify, and then click **Edit**.
   The Edit Volume dialog box is displayed.
4. In the **General** tab, modify the following properties as required:
   - Volume name
   - Security style for the volume
   - Thin provisioning
5. Click the **Storage Efficiency** tab, and enable storage efficiency by configuring the following properties:
   - Deduplication
   - Data compression

   You cannot enable background compression for a volume that is contained by an aggregate with All Flash Optimized personality. You can enable only inline compression for these volumes.

   You can enable inline deduplication only on a volume that is contained by an aggregate with All Flash Optimized personality or on a volume in a Flash Pool aggregate.
6. For SnapLock volumes, click the **SnapLock** tab, and perform the following steps:
   a. Specify the autocommit period.

      The autocommit period determines how long a file in the volume must remain unchanged before the file is committed to WORM state.
   b. Specify the minimum retention period and maximum retention period.

The values must be in the range of 1 day through 70 years or Infinite.

c. Select the default retention period.

The default retention period must be within the specified minimum retention period and maximum retention period.

7. Click the **Advanced** tab, and enable the following properties:

- If you want the volume to automatically grow when the used space in the volume is above the grow threshold, select **Grow**.
- If you want the volume to grow or shrink in size in response to the amount of used space, select **Grow or Shrink**.

a. Specify the maximum size to which the volume can grow.

- Enable automatic deletion of older Snapshot copies by choosing one of the following options:
  - Try
    Deletes the Snapshot copies that are not locked by any other subsystems.
  - Destroy
    Deletes the Snapshot copies that are locked by the data-backing functionality.
  - Disrupt
    Deletes the Snapshot copies that can disrupt the data transfer.
- Select the caching policy that you want to assign to the volume.
  This option is available only for FlexVol volumes in a Flash Pool aggregate.
- Select the retention priority for cached data in the volume.
  This option is available only for FlexVol volumes in a Flash Pool aggregate.
- Specify the fractional reserve that you want to set for the volume.
- Update the access time for reading the file.
  This option is disabled for SnapLock volumes.

8. Click **Save and Close**.

### Related tasks

*Setting up CIFS* on page 247
You can use System Manager to enable and configure CIFS servers to allow CIFS clients to access the files on the cluster.

### Related reference

*Volumes window* on page 203
You can use the Volumes window to manage your volumes, and to display information about these volumes.

## Editing data protection volumes

You can use System Manager to modify the volume name for a data protection (DP) volume. If the source volume does not have storage efficiency enabled, you might want to enable storage efficiency only on the destination volume.

### About this task

You cannot modify storage efficiency on a mirror DP volume.

### Steps

1. Click **Storage** > **Volumes**.
2. From the **Volumes on SVM** list, select the storage virtual machine (SVM) in which the DP volume that you want to edit resides.
3. Select the volume that you want to modify, and then click **Edit**.

4. In the **Edit Data Protection Volume** dialog box, modify the volume name.
5. Select **Enable Storage Efficiency**.

   If storage efficiency is already enabled on the volume, then the check box is selected by default.
6. Optional: Click the **Advanced** tab, and perform the following steps:
   a. Select the caching policy that you want to assign to the volume.
   b. Select the retention priority for the cached data in the volume.

   These options are available only for data protection FlexVol volumes in a Flash Pool aggregate.
7. Click **Save**.

## Deleting volumes

You can use System Manager to delete a volume when you no longer require the data that a volume contains, or if you have copied the data that a volume contains to another location. When you delete a volume, all the data in the volume is destroyed, and you cannot recover this data.

### Before you begin

- If the FlexVol volume is cloned, the FlexClone volumes must be either split from the parent volume or destroyed.
- The volume must be unmounted and must be in the offline state.
- If the volume is in one or more SnapMirror relationships, the SnapMirror relationships must be deleted.
- You can delete a complete SnapLock Enterprise volume or a file in a SnapLock Enterprise volume; however, you cannot delete only the data within a file in a SnapLock Enterprise volume.
- You cannot delete a SnapLock Compliance volume if data is committed to the volume.

### Steps

1. Click **Storage** > **Volumes**.
2. From the **Volumes on SVM** list, select the storage virtual machine (SVM) in which the volume that you want to delete resides.
3. Select the volumes that you want to delete, and then click **Delete**.
4. Select the confirmation check box, and then click **Delete**.

### Related reference

*Volumes window* on page 203

You can use the Volumes window to manage your volumes, and to display information about these volumes.

## Creating FlexClone volumes

You can use System Manager to create a FlexClone volume when you require a writable, point-in-time copy of an existing FlexVol volume. You might want to create a copy of a volume for testing or to provide access to the volume for additional users without giving them access to the production data.

### Before you begin

- The FlexClone license must be installed on the storage system.
- The volume that you want to clone must be online and must be a non-root volume.

### About this task

The base Snapshot copy that is used to create a FlexClone volume of a SnapMirror destination is marked as busy and cannot be deleted. If a FlexClone volume is created from a Snapshot copy that

is not the most recent Snapshot copy, and that Snapshot copy no longer exists on the source volume, all SnapMirror updates to the destination volume fail.

### Steps

1. Click **Storage** > **Volumes**.
2. From the **Volumes on SVM** list, select **All SVMs**.
3. Select the FlexVol volume that you want to clone from the list of volumes.
4. Click **More Actions** > **Clone** > **Create** > **Volume**.
5. In the **Create FlexClone Volume** dialog box, type the name of the FlexClone volume that you want to create.
6. Optional: If you want to enable thin provisioning for the new FlexClone volume, select **Thin Provisioning**.

   By default, this setting is the same as that of the parent volume.
7. Create a Snapshot copy or select an existing Snapshot copy that you want to use as the base Snapshot copy for creating the FlexClone volume.
8. Click **Clone**.

### Related reference

*Volumes window* on page 203

You can use the Volumes window to manage your volumes, and to display information about these volumes.

## Creating FlexClone files

You can use System Manager to create a FlexClone file, which is a writable copy of a parent file. You can use these copies to test applications.

### Before you begin

- The file that is cloned must be part of the active file system.
- The FlexClone license must be installed on the storage system.

### About this task

- FlexClone files are supported only for FlexVol volumes.

  You can create a FlexClone file of a parent file that is within a volume by accessing the parent file from the volume in which it resides, not from the parent volume.
- You cannot create a FlexClone file on a SnapLock volume.

### Steps

1. Click **Storage** > **Volumes**.
2. From the **Volumes on SVM** list, select **All SVMs**.
3. Select the volume in which you want to create a FlexClone file from the list of volumes.
4. Click **More Actions** > **Clone** > **Create** > **File**.
5. In the **Create FlexClone File** dialog box, select the file that you want to clone, and then specify a name for the FlexClone file.
6. Click **Clone**.

### Result

The FlexClone file is created in the same volume as the parent file.

### Related reference

*Volumes window* on page 203

You can use the Volumes window to manage your volumes, and to display information about these volumes.

## Splitting a FlexClone volume from its parent volume

If you want a FlexClone volume to have its own disk space instead of using the disk space of its parent volume, you can split the volume from its parent by using System Manager. After the split, the FlexClone volume becomes a normal FlexVol volume.

### Before you begin

The FlexClone volume must be online.

### About this task

The clone-splitting operation deletes all of the existing Snapshot copies of the clone. The Snapshot copies that are required for SnapMirror updates are also deleted. Therefore, any subsequent SnapMirror updates might fail.

You can pause the clone-splitting operation if you have to perform any other operation on the volume. You can resume the clone-splitting process after the other operation is complete.

### Steps

1. Click **Storage** > **Volumes**.
2. From the **Volumes on SVM** list, select **All SVMs**.
3. Select the FlexClone volume that you want to split from its parent volume.
4. Click **More Actions** > **Clone** > **Split**.
5. Confirm the FlexClone volume details for the clone-splitting operation, and then click **Start Split** in the confirmation dialog box.

### Related reference

*Volumes window* on page 203

You can use the Volumes window to manage your volumes, and to display information about these volumes.

## Viewing the FlexClone volume hierarchy

You can use System Manager to view the hierarchy of FlexClone volumes and their parent volumes.

### Steps

1. Click **Storage** > **Volumes**.
2. From the **Volumes on SVM** list, select **All SVMs**.
3. Select the required volume from the list of volumes.
4. Click **More Actions** > **Clone** > **View Hierarchy**.

### Result

Volumes that have at least one child FlexClone volume are displayed. The FlexClone volumes are displayed as children of their respective parent volumes.

### Related reference

*Volumes window* on page 203

You can use the Volumes window to manage your volumes, and to display information about these volumes.

## Changing the status of a volume

You can use System Manager to change the status of a volume when you want to take a volume offline, bring a volume back online, or restrict access to a volume.

### Before you begin

- If you want a volume to be the target of a volume copy operation or a SnapMirror replication operation, the volume must be in the restricted state.
- If you want to take a NAS volume offline, the NAS volume must be unmounted.

### About this task

You can take a volume offline to perform maintenance on the volume, to move the volume, or to destroy the volume. When a volume is offline, the volume is unavailable for read or write access by clients. You cannot take a root volume offline.

### Steps

1. Click **Storage** > **Volumes**.
2. From the **Volumes on SVM** list, select **All SVMs**.
3. Select the volume for which you want to modify the status.
4. From the **More Actions** > **Change status to** menu, select the required volume status.
5. Click **Ok** in the confirmation dialog box to change the volume status.

### Related reference

*Volumes window* on page 203
You can use the Volumes window to manage your volumes, and to display information about these volumes.

## Viewing the list of saved Snapshot copies

You can use System Manager to view the list of all of the saved Snapshot copies for a selected volume from the Snapshot Copies tab in the lower pane of the Volumes window. You can use the list of saved Snapshot copies to rename, restore, or delete a Snapshot copy.

### Before you begin

The volume must be online.

### About this task

You can view Snapshot copies for only one volume at a time.

### Steps

1. Click **Storage** > **Volumes**.
2. From the **Volumes on SVM** list, select **All SVMs**.
3. Expand the required volume.
4. Click the **Show More Details** link to view more information about the volume.
5. Click the **Snapshot Copies** tab.
   The list of available Snapshot copies for the selected volume is displayed.

## Creating Snapshot copies outside a defined schedule

You can use System Manager to create a Snapshot copy of a volume outside a defined schedule to capture the state of the file system at a specific point in time.

### Steps

1. Click **Storage** > **Volumes**.
2. From the **Volumes on SVM** list, select **All SVMs**.
3. Select the volume from the list of volumes.
4. Click **More Actions** > **Manage Snapshots** > **Create**.
5. In the **Create Snapshot Copy** dialog box, if you want to change the default name, specify a new name for the Snapshot copy.

   Valid characters are ASCII characters, numerals, hyphens (-), underscores (_), periods (.), and the plus (+) symbol.

   The default name of a Snapshot copy consists of the volume name and the timestamp.
6. Click **Create**.
7. Verify that the Snapshot copy that you created is included in the list of Snapshot copies in the **Snapshot Copies** tab.

### Related reference

*Volumes window* on page 203

You can use the Volumes window to manage your volumes, and to display information about these volumes.

## Setting the Snapshot copy reserve

You can use System Manager to reserve space (in percentage) for the Snapshot copies in a volume. By setting the Snapshot copy reserve, you can allocate enough disk space for the Snapshot copies so that they do not consume the active file system space.

### About this task

- The default space that is reserved for Snapshot copies is 5 percent for SAN and VMware volumes.

### Steps

1. Click **Storage** > **Volumes**.
2. From the **Volumes on SVM** list, select **All SVMs**.
3. Select the volume for which you want to set the Snapshot copy reserve.
4. Click **More Actions** > **Manage Snapshots** > **Configuration Settings**.
5. Type or select the percentage of volume space that you want to reserve for the Snapshot copies, and then click **OK**.

### Related reference

*Volumes window* on page 203

You can use the Volumes window to manage your volumes, and to display information about these volumes.

## Hiding the Snapshot copy directory

You can use System Manager to hide the Snapshot copy directory (`.snapshot`) so that the Snapshot copy directory is not visible when you view your volume directories. By default, the `.snapshot` directory is visible.

### Steps

1. Click **Storage** > **Volumes**.
2. From the **Volumes on SVM** list, select **All SVMs**.
3. Select the volume for which you want to hide the Snapshot copy directory.
4. Click **More Actions** > **Manage Snapshots** > **Configuration Settings**.
5. Ensure that the **Make snapshot directory (.snapshot) visible** option is not selected, and then click **OK**.

### Related reference

*Volumes window* on page 203

You can use the Volumes window to manage your volumes, and to display information about these volumes.

## Scheduling automatic creation of Snapshot copies

You can use System Manager to set up a schedule for the automatic creating automatic Snapshot copies of a volume. You can specify the time and frequency of creating the copies. You can also specify the number of Snapshot copies that are saved.

### Steps

1. Click **Storage** > **Volumes**.
2. From the **Volumes on SVM** list, select **All SVMs**.
3. Select the required volume from the list of volumes.
4. Click **More Actions** > **Manage Snapshots** > **Configuration Settings**.
5. In the **Configure Volume Snapshot Copies** dialog box, select **Enable scheduled Snapshot Copies**.
6. Select a Snapshot policy.

   You can schedule the creation of only policy-based Snapshot copies.
7. Click **OK** to save your changes and start your Snapshot copy schedule.

### Related reference

*Volumes window* on page 203

You can use the Volumes window to manage your volumes, and to display information about these volumes.

## Restoring a volume from a Snapshot copy

You can use System Manager to restore a volume to a state that is recorded in a previously created Snapshot copy to retrieve lost information. When you restore a volume from a Snapshot copy, the restore operation overwrites the existing volume configuration. Any changes that were made to the data in the volume after the Snapshot copy was created are lost.

### Before you begin

- The SnapRestore license must be installed on your system.

- If the FlexVol volume that you want to restore contains a LUN, the LUN must be unmounted or unmapped.
- There must be enough space available for the restored volume.
- Users accessing the volume must be notified that you are going to revert a volume, and that the data from the selected Snapshot copy replaces the current data in the volume.

### About this task

- If the volume that you restore contains junction points to other volumes, the volumes that are mounted on these junction points will not be restored.
- You cannot restore Snapshot copies for SnapLock Compliance volumes.

### Steps

1. Click **Storage** > **Volumes**.
2. From the **Volumes on SVM** list, select **All SVMs**.
3. Select the volume that you want to restore from a Snapshot copy.
4. Click **More Actions** > **Manage Snapshots** > **Restore**.
5. Select the appropriate Snapshot copy, and then click **Restore**.
6. Select the confirmation check box, and then click **Restore**.

### Related reference

*Volumes window* on page 203

You can use the Volumes window to manage your volumes, and to display information about these volumes.

## Extending the expiry date of Snapshot copies

You can use System Manager to extend the expiry date of the Snapshot copies in a volume.

### Before you begin

The SnapLock license must be installed on your system.

### About this task

You can extend the expiry date only for Snapshot copies in a data protection (DP) volume that is the destination in a SnapLock for SnapVault relationship.

### Steps

1. Click **Storage** > **Volumes**.
2. From the **Volumes on SVM** list, select **All SVMs**.
3. Expand the required volume.
4. Click the **Show More Details** link to view more information about the volume.
5. Click the **Snapshot Copies** tab.
   The list of available Snapshot copies for the selected volume is displayed.
6. Select the Snapshot copy that you want to modify, and then click **Extend Expiry Date**.
7. In the **Extend Expiry Date** dialog box, specify the expiry date.

   The values must be in the range of 1 day through 70 years or Infinite.
8. Click **OK**.

## Renaming Snapshot copies

You can use System Manager to rename a Snapshot copy to help you organize and manage your Snapshot copies.

### About this task

- You cannot rename the Snapshot copies (which are committed to WORM state) of a SnapLock DP volume that is in a SnapVault relationship.

### Steps

1. Click **Storage** > **Volumes**.
2. From the **Volumes on SVM** list, select **All SVMs**.
3. Click the required volume.
4. Click the **Show More Details** link to view more information about the volume.
5. Click the **Snapshot Copies** tab.
   The list of available Snapshot copies for the selected volume is displayed.
6. Select the Snapshot copy that you want to rename, and then click **More Actions** > **Rename**.
7. Specify a new name, and then click **Rename**.

   Valid characters are ASCII characters, numerals, hyphens (-), underscores (_), periods (.), and the plus (+) symbol.
8. Verify the Snapshot copy name in the **Snapshot Copies** tab of the **Volumes** window.

### Related reference

*Volumes window* on page 203
You can use the Volumes window to manage your volumes, and to display information about these volumes.

## Deleting Snapshot copies

You can delete a Snapshot copy to conserve disk space or to free disk space by using System Manager. You can also delete a Snapshot copy if the Snapshot copy is no longer required.

### Before you begin

If you want to delete a Snapshot copy that is busy or locked, you must have released the Snapshot copy from the application that was using the Snapshot copy.

### About this task

- You cannot delete the base Snapshot copy in a parent volume if a FlexClone volume is using that Snapshot copy.

  The base Snapshot copy is the Snapshot copy that is used to create a FlexClone volume. The base Snapshot copy always displays the status `busy` and Application Dependency as `busy,vclone` in the parent volume.
- You cannot delete a locked Snapshot copy that is used in a SnapMirror relationship.

  The Snapshot copy is locked and is required for the next update.
- You cannot delete a Snapshot copy from a SnapLock DP volume that is used in a SnapVault relationship before the expiry time of the Snapshot copy.
- You cannot delete the unexpired Snapshot copies (which are committed to WORM state) of a SnapLock DP volume that is in a SnapVault relationship.

### Steps

1. Click **Storage** > **Volumes**.
2. From the **Volumes on SVM** list, select **All SVMs**.
3. Expand the required volume.

4. Click the **Show More Details** link to view more information about the volume.
5. Click the **Snapshot Copies** tab.
   The list of available Snapshot copies for the selected volume is displayed.
6. Select the Snapshot copy that you want to delete.
7. Click **Delete**.
8. Select the confirmation check box, and then click **Delete**.

### Related reference

*Volumes window* on page 203
You can use the Volumes window to manage your volumes, and to display information about these volumes.

### Related information

*NetApp Documentation: ONTAP 9*

## Resizing volumes

When your volume reaches nearly full capacity, you can increase the size of the volume, delete some Snapshot copies, or adjust the Snapshot reserve. You can use the Volume Resize wizard in System Manager to provide more free space.

### About this task

- For a volume that is configured to grow automatically, you can modify the limit to which the volume can grow automatically based on the increased size of the volume.
- You cannot resize a data protection volume if its mirror relationship is broken or if a reverse resynchronization operation has been performed on the volume.
  Instead, you must use the command-line interface (CLI).

### Steps

1. Click **Storage** > **Volumes**.
2. From the **Volumes on SVM** list, select **All SVMs**.
3. Select the volume that you want to resize.
4. Click **More Actions** > **Resize**.
5. Type or select information as prompted by the wizard.
6. Confirm the details, and then click **Finish** to complete the wizard.
7. Verify the changes that you made to the available space and the total space of the volume in the **Volumes** window.

### Related reference

*Volumes window* on page 203
You can use the Volumes window to manage your volumes, and to display information about these volumes.

## Enabling storage efficiency on a volume

You can use System Manager to enable storage efficiency and to configure both deduplication and data compression or only deduplication on a volume to save storage space. If you have not enabled storage efficiency when you created the volume, you can do so later by editing the volume.

### Before you begin

- The volume must be online.
- If you want to use a policy-based deduplication schedule, you must have created an efficiency policy.

### About this task

- You can enable background compression only if you have enabled background deduplication.
- You can enable inline compression and inline deduplication with or without enabling background compression and background deduplication, respectively.
- You can enable inline deduplication only on volumes that are contained by an aggregate with All Flash Optimized personality and on volumes that are contained by a Flash Pool aggregate.

### Steps

1. Click **Storage** > **Volumes**.
2. From the **Volumes on SVM** list, select **All SVMs**.
3. Select the volume for which you want to enable storage efficiency, and then click **Edit**.
4. In the **Edit Volume** dialog box, click **Storage Efficiency**.
5. Select the **Background Deduplication** check box.
6. Select one of the following methods to run deduplication:

| If you want to run deduplication... | Then... |
|---|---|
| Based on a storage efficiency policy | **a.** Ensure that the **Policy based** option is selected.<br>**b.** Click **Choose**, and then select a storage efficiency policy.<br>**c.** Click **OK**. |
| When required | Select the **On-demand** option. |

7. Optional: Select the **Background Compression** check box to enable background compression.

   You cannot enable background compression for a volume that is contained by an aggregate with All Flash Optimized personality.
8. Optional: Select the **Inline Compression** check box to compress data while it is being written to the volume.

   By default, inline compression is enabled on volumes that are contained by an aggregate with All Flash Optimized personality.
9. Optional: Select the **Inline Deduplication** check box to run deduplication before data is written to the disk.

   By default, inline deduplication is enabled on volumes that are contained by an aggregate with All Flash Optimized personality.
10. Click **Save and Close**.

### Related reference

*Volumes window* on page 203
You can use the Volumes window to manage your volumes, and to display information about these volumes.

## Changing the deduplication schedule

You can use System Manager to change the deduplication schedule by choosing to run deduplication manually, automatically, or on a schedule that you specify.

### Steps

1. Click **Storage** > **Volumes**.
2. From the **Volumes on SVM** list, select **All SVMs**.

3. Select the read/write volume for which you want to modify the deduplication schedule.
4. Click **Edit**, and then click the **Storage Efficiency** tab.
5. Change the deduplication schedule as required.
6. Click **Save and Close**.

### Related reference

*Volumes window* on page 203
You can use the Volumes window to manage your volumes, and to display information about these volumes.

## Running deduplication operations

You can use System Manager to run deduplication immediately after creating a volume, or to schedule deduplication to run at a specified time.

### Before you begin

- Deduplication must be enabled on the volume.
- The volume must be online and mounted.

### About this task

Deduplication is a background process that consumes system resources during the operation; therefore, it might affect other operations that are in progress. You must cancel deduplication before you can perform any other operation.

### Steps

1. Click **Storage** > **Volumes**.
2. From the **Volumes on SVM** list, select **All SVMs**.
3. Select the volume for which you want to run deduplication.
4. Click **More Actions** > **Storage Efficiency**.
5. If you are running deduplication on the volume for the first time, run deduplication on the entire volume data by selecting **Scan Entire Volume** in the **Storage Efficiency** dialog box.
6. Click **Start**.
7. View the last-run details of the deduplication operation in the **Storage Efficiency** tab of the **Volumes** window.

### Related reference

*Volumes window* on page 203
You can use the Volumes window to manage your volumes, and to display information about these volumes.

## Moving FlexVol volumes between aggregates or nodes

You can nondisruptively move a FlexVol volume to a different aggregate or a different node for capacity utilization and improved performance by using System Manager.

### Before you begin

If you are moving a data protection (DP) volume, the data protection mirror relationships must be initialized before you move the volume.

### About this task

You cannot move SnapLock volumes between aggregates and nodes.

### Steps

1. Click **Storage** > **Volumes**.

2. From the **Volumes on SVM** list, select **All SVMs**.
3. Select the volume that you want to move.
4. Click **More Actions** > **Move**.
5. In the **Move Volume dialog box**, select the destination aggregate or node for the volume, and then change the tiering policy.

   **Note:**

   - You cannot change the tiering policy of a root volume.
   - You cannot move the root volume to FabricPool.
   - For read/write volumes, you can set the tiering policy as "back up" during the volume move.
     The tiering policy changes to "snapshot-only" after the move.
   - Capacity tier values that are displayed in the "Used After Move" in both the source aggregate and destination aggregate are estimated values.
     For the exact values, you must navigate to the Aggregate window and view the details of a specific aggregate.

6. Click **Move**.

### Manually triggering the cutover for volume move

For a volume move operation, you can use System Manager to manually trigger the cutover when the volume enters the cutover deferred phase. You can set the duration of the cutover and the cutover action to be performed by the system if the operation fails within that duration.

#### Steps

1. Click **Storage** > **SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. Click the **Volumes** tab.
4. Expand the volume for which the volume move operation has been initiated.
5. Click the **Show More Details** link to view more information about the volume.
6. In the **Overview** tab, click **Cutover**.
7. In the **Cutover** dialog box, click **Advanced Options**.
8. Optional: Specify the cutover action and the cutover window period.
9. Click **OK**.

## Assigning volumes to Storage QoS

You can limit the throughput of FlexVol volumes and FlexGroup volumes by assigning them to storage Quality of Service (QoS) policy groups. You can assign storage QoS for new volumes, or you can modify the storage QoS details of the volumes that are already assigned to a policy group by using System Manager.

#### About this task

- You can assign storage QoS only to read/write (rw) volumes that are online.
- You cannot assign storage QoS to a volume if the following storage objects are assigned to a policy group:
  - Parent storage virtual machine (SVM) of the volume
  - Child LUNs of the volume
  - Child files of the volume
- You can assign storage QoS or modify the QoS details for a maximum of 10 volumes simultaneously.

### Steps

1. Click **Storage** > **Volumes**.
2. From the **Volumes on SVM** list, select **All SVMs**.
3. Select one or more volumes for which you want to assign storage QoS.
4. Click **More Actions** > **Storage QoS**.
5. In the **Quality of Service Details** dialog box, select the **Manage Storage Quality of Service** check box if you want to manage the workload performance of the FlexVol volume.

   If some of the volumes that you selected are already assigned to a policy group, the changes that you make might affect the performance of these volumes.
6. Create a new storage QoS policy group or select an existing policy group to control the input/output (I/O) performance of the FlexVol volume:

| If you want to... | Do this... |
|---|---|
| Create a new policy group | **a.** Select **New Policy Group**.<br>**b.** Specify the policy group name.<br>**c.** Specify the minimum throughput limit.<br>• You can set the minimum throughput limit only on an AFF platform.<br>• You cannot set the minimum throughput limit for volumes on a FabricPool-enabled aggregate.<br>• If you do not specify the minimum throughput value or if the minimum throughput value is set to 0, the system automatically displays "None" as the value.<br>This value is case-sensitive.<br>**d.** Specify the maximum throughput limit to prevent the workload of the objects in the policy group from exceeding the specified throughput limit.<br>• The minimum throughput limit and the maximum throughput limit must be of the same unit type.<br>• If you do not specify the minimum throughput limit, then you can set the maximum throughput limit in IOPS, B/s, KB/s, MB/s, and so on.<br>• If you do not specify the maximum throughput value, the system automatically displays "Unlimited" as the value.<br>This value is case-sensitive. The unit that you specify does not affect the maximum throughput. |

| If you want to... | Do this... |
|---|---|
| Select an existing policy group | **a.** Select **Existing Policy Group**, and then click **Choose** to select an existing policy group from the Select Policy Group dialog box.<br>**b.** Specify the minimum throughput limit.<br>• You can set the minimum throughput limit only on an AFF platform.<br>• You cannot set the minimum throughput limit for volumes on a FabricPool-enabled aggregate.<br>• If you do not specify the minimum throughput value or if the minimum throughput value is set to 0, the system automatically displays "None" as the value.<br>This value is case-sensitive.<br>**c.** Specify the maximum throughput limit to prevent the workload of the objects in the policy group from exceeding the specified throughput limit.<br>• The minimum throughput limit and the maximum throughput limit must be of the same unit type.<br>• If you do not specify the minimum throughput limit, then you can set the maximum throughput limit in IOPS, B/s, KB/s, MB/s, and so on.<br>• If you do not specify the maximum throughput value, the system automatically displays "Unlimited" as the value.<br>This value is case-sensitive. The unit that you specify does not affect the maximum throughput.<br>If the policy group is assigned to more than one object, the maximum throughput that you specify is shared among the objects. |

7. Optional: Click the link that specifies the number of volumes if you want to review the list of selected volumes, and then click **Discard** if you want to remove any volumes from the list.

   The link is displayed only when multiple volumes are selected.
8. Click **OK**.

## Creating a mirror relationship from a source SVM

You can use System Manager to create a mirror relationship from the source storage virtual machine (SVM), and to assign a mirror policy and schedule to the mirror relationship. The mirror copy enables quick availability of data if the data on the source volume is corrupted or lost.

### Before you begin

- The SnapMirror license must be enabled on the source cluster and destination cluster.

  **Note:**

  - For some platforms, it is not mandatory for the source cluster to have the SnapMirror license enabled if the destination cluster has the SnapMirror license and Data Protection Optimization (DPO) license enabled.
  - After the DPO license is enabled on the destination cluster, you must refresh the browser of the source cluster to enable the `Protect` option.

- While mirroring a SnapLock volume, the SnapMirror license must be installed on both the source cluster and destination cluster, and the SnapLock license must be installed on the destination cluster.
- The source cluster and destination cluster, and the source SVM and destination SVM must be in a healthy peer relationship.
- The destination aggregate must have space available.

- FlexVol volumes must be online and read/write.
- The SnapLock aggregate type must be the same.
- A maximum of 25 volumes can be protected in one selection.
- If you are connecting from a cluster running ONTAP 9.2 or earlier to a remote cluster on which Security Assertion Markup Language (SAML) authentication is enabled, password-based authentication must be enabled on the remote cluster.

### About this task

- System Manager does not support a cascade relationship.

  For example, a destination volume in a relationship cannot be the source volume in another relationship.
- You can create a mirror relationship between SnapLock volumes of the same type only.

  For example, if the source volume is a SnapLock Enterprise volume, then the destination volume must also be a SnapLock Enterprise volume.
- If encryption is enabled on the source volume and the destination cluster is running a version of ONTAP software earlier than ONTAP 9.3, then encryption is disabled on the destination volume by default.

### Steps

1. Click **Storage** > **Volumes**.
2. From the **Volumes on SVM** list, select **All SVMs**.
3. Select the volumes for which you want to create mirror relationships, and then click **More Actions** > **Protect**.

   The **Protect** option is available only for a read/write volume.
4. Select **Mirror** from the **Volume Relationship Type** list.
5. Specify the cluster and the SVM, and then enter a name suffix for the destination volume.

   If the specified cluster is running a version of ONTAP software earlier than ONTAP 9.3, then only peered SVMs are listed. If the specified cluster is running ONTAP 9.3 or later, peered SVMs and permitted SVMs are listed.
6. Optional: Click ⚙, update the protection policy and protection schedule, select **FabricPool-enabled aggregate**, and then initialize the protection relationship.
7. Click **Save**.

### Result

A new destination volume of type *dp* is created with the following default settings:

- Autogrow is enabled.
- Compression is disabled.
- The language attribute is set to match the language attribute of the source volume.

If the destination FlexVol volume is on a different SVM than the source FlexVol volume, then a peer relationship is created between the two SVMs if the relationship does not already exist.

A mirror relationship is created between the source volume and the destination volume. The base Snapshot copy is transferred to the destination volume if you have opted to initialize the relationship.

#### Related reference

*Protection window* on page 336

You can use the Protection window to create and manage mirror relationships, vault relationships, and mirror vault relationships and to display details about these relationships. The Protection window does not display load-sharing (LS) relationships and transition (TDP) relationships.

## Creating a vault relationship from a source SVM

You can use System Manager to create a vault relationship from the source storage virtual machine (SVM), and to assign a vault policy to the vault relationship to create a backup vault. In the event of data loss or corruption on a system, backed-up data can be restored from the backup vault destination.

### Before you begin

- The SnapVault license or SnapMirror license must be enabled on both the source cluster and the destination cluster.

  **Note:**

  - For some platforms, it is not mandatory for the source cluster to have the SnapVault license or the SnapMirror license enabled if the destination cluster has the SnapVault license or the SnapMirror license, and DPO license enabled.
  - After the DPO license is enabled on the destination cluster, you must refresh the browser of the source cluster to enable the `Protect` option.

- The source cluster and destination cluster, and the source SVM and destination SVM must be in a healthy peer relationship.
- The destination aggregate must have space available.
- The source aggregate and the destination aggregate must be 64-bit aggregates.
- A vault (XDP) policy must exist.

  If a vault policy does not exist, you must create a vault policy or accept the default vault policy (named XDPDefault) that is automatically assigned.
- FlexVol volumes must be online and read/write.
- The SnapLock aggregate type must be the same.
- A maximum of 25 volumes can be protected in one selection.
- If you are connecting from a cluster running ONTAP 9.2 or earlier to a remote cluster on which Security Assertion Markup Language (SAML) authentication is enabled, password-based authentication must be enabled on the remote cluster.

### About this task

- System Manager does not support a cascade relationship.

  For example, a destination volume in a relationship cannot be the source volume in another relationship.
- You can create a lock-vault relationship only between a non-SnapLock (primary) volume and a Snaplock destination (secondary) volume.
- If encryption is enabled on the source volume and the destination cluster is running a version of ONTAP software earlier than ONTAP 9.3, then encryption is disabled on the destination volume by default.

### Steps

1. Click **Storage** > **Volumes**.
2. From the **Volumes on SVM** list, select **All SVMs**.
3. Select the volumes for which you want to create vault relationships, and then click **More Actions** > **Protect**.

   The **Protect** option is available only for a read/write volume.

4. Select **Vault** from the **Volume relationship Type** list.
5. Specify the cluster and the SVM, and then enter a name suffix for the destination volume.

   If the specified cluster is running a version of ONTAP software earlier than ONTAP 9.3, then only peered SVMs are listed. If the specified cluster is running ONTAP 9.3 or later, peered SVMs and permitted SVMs are listed.
6. Optional: Click ⚙, update the protection policy and protection schedule, enable SnapLock properties on the destination volume, select a FabricPool-enabled aggregate, and then initialize the protection relationship.
7. Click **Save**.

### Related reference

*Protection window* on page 336

You can use the Protection window to create and manage mirror relationships, vault relationships, and mirror vault relationships and to display details about these relationships. The Protection window does not display load-sharing (LS) relationships and transition (TDP) relationships.

## Creating a mirror and vault relationship from a source SVM

You can use System Manager to create a mirror and vault relationship from the source storage virtual machine (SVM). Creating this relationship enables you to better protect your data by periodically transferring data from the source volume to the destination volume. It also enables you to retain data for long periods by creating backups of the source volume.

### Before you begin

- The source cluster must be running ONTAP 8.3.2 or later.
- The SnapMirror license must be enabled on the source cluster and destination cluster.

  **Note:**

  - For some platforms, it is not mandatory for the source cluster to have the SnapMirror license enabled if the destination cluster has the SnapMirror license and Data Protection Optimization (DPO) license enabled.
  - After the DPO license is enabled on the destination cluster, you must refresh the browser of the source cluster to enable the `Protect` option.

- The source cluster and destination cluster must be in a healthy peer relationship.
- The source SVM and destination SVM must be either in a healthy peer relationship or the destination SVM must have permission to peer.
- The destination aggregate must have space available.
- The source aggregate and the destination aggregate must be 64-bit aggregates.
- FlexVol volumes must be online and read/write.
- The SnapLock aggregate type must be the same.
- A maximum of 25 volumes can be protected in one selection.
- If you are connecting from a cluster running ONTAP 9.2 or earlier to a remote cluster on which Security Assertion Markup Language (SAML) authentication is enabled, password-based authentication must be enabled on the remote cluster.

### About this task

- System Manager does not support a cascade relationship.

  For example, a destination volume in a relationship cannot be the source volume in another relationship.
- If encryption is enabled on the source volume and the destination cluster is running a version of ONTAP software earlier than ONTAP 9.3, then encryption is disabled on the destination volume by default.

### Steps

1. Click **Storage** > **Volumes**.
2. From the **Volumes on SVM** list, select **All SVMs**.
3. Select the volumes for which you want to create mirror and vault relationships, and then click **More Actions** > **Protect**.

   The **Protect** option is available only for a read/write volume.
4. Select **Mirror and Vault** from the **Volume relationship Type** list.
5. Specify the cluster and the SVM, and then enter a name suffix for the destination volume.

   If the specified cluster is running a version of ONTAP software earlier than ONTAP 9.3, then only peered SVMs are listed. If the specified cluster is running ONTAP 9.3 or later, peered SVMs and permitted SVMs are listed.
6. Optional: Click ⚙, update the protection policy and protection schedule, select **FabricPool-enabled aggregate**, and then initialize the protection relationship.
7. Click **Save**.

## Creating an NFS datastore for VMware

You can use the Create NFS Datastore for VMware wizard in System Manager to create an NFS datastore for VMware. You can create a volume for the NFS datastore and specify the ESX servers that can access the NFS datastore.

### Before you begin

The NFS service must be licensed.

### Steps

1. Click **Storage** > **Volumes**.
2. From the **Volumes on SVM** list, select **All SVMs**.
3. Select the volume, and then click **More Actions** > **Provision Storage for VMware**.
4. In the **Create NFS Datastore for VMware** wizard, type or select information as required.
5. Confirm the details, and then click **Finish** to complete the wizard.

## Changing the tiering policy of a volume

You can use System Manager to change the default tiering policy of a volume to control whether the data of the volume is moved to the capacity tier when the data becomes inactive.

### Steps

1. Click **Storage** > **Volumes**.
2. From the **Volumes on SVM** list, select **All SVMs**.
3. Select the volume for which you want to change the tiering policy, and then click **More Actions** > **Change Tiering Policy**.
4. Select the required tiering policy from the **Tiering Policy** list, and then click **Save**.

## Creating FlexGroup volumes

You can use System Manager to create a FlexGroup volume by selecting specific aggregates or by selecting system-recommended aggregates.

### About this task

- You can create only read/write (rw) FlexGroup volumes.
- You cannot create FlexGroup volumes in a MetroCluster configuration.

### Steps

1. Click **Storage** > **Volumes**.
2. Click **Create** > **Create FlexGroup**.
3. In the **Create FlexGroup Volume** window, specify a name for the FlexGroup volume.

   By default, the aggregates are selected as per best practices.

   **Note:** If you want to select aggregates as per your requirement, click (advanced options).
4. Specify a size for the FlexGroup volume.
5. Enable the **Encrypted** button to enable encryption for the volume.

   This option is available only if you have enabled the Volume Encryption license and if the corresponding platform supports encryption.
6. Optional: Click to specify the advanced options.
   a. In the **General Details** section, select the space reserve and security style, and then set the UNIX permission for the volume.

      When thin provisioning is enabled, space is allocated to the volume from the aggregate only when data is written to the volume.

      **Note:** For AFF (AFF) storage systems, thin provisioning is enabled by default and for other storage systems, thick provisioning is enabled by default.
   b. In **Optimize space** section, enable fractional reserve, and then grow or shrink the volume, as required.
   c. In the **Storage Efficiency** section, enable deduplication on the volume.

      System Manager uses the default deduplication schedule. If the specified volume size exceeds the limit that is required for running deduplication, the volume is created and deduplication is not enabled.

      For systems with All Flash Optimized personality, inline compression and the `auto` deduplication schedule is enabled by default.
   d. In the **Quality of Service** section, specify the policy group to control the input/output (I/O) performance of the FlexGroup volume.
   e. Click **Apply** to update the changes.
7. In the **Protection** section, enable the **Volume Protection** option and select the **Volume Relationship Type**.
8. Click **Create** to create the FlexGroup volume.

### Related reference

*Volumes window* on page 203

You can use the Volumes window to manage your volumes, and to display information about these volumes.

## Editing FlexGroup volumes

You can use System Manager to edit the properties of an existing FlexGroup volume.

### Before you begin

The FlexGroup volume must be online.

### Steps

1. Click **Storage** > **Volumes**.
2. From the **Volumes on SVM** list, select **All SVMs**.
3. Select the FlexGroup volume that you want to modify, and click **Edit**.

4. Specify the size to which you want to resize the FlexGroup volume.

   By default, existing aggregates are used to resize the FlexGroup volume.

   **Note:** If you want to expand the FlexGroup volume by adding new resources, you should click (advanced options).

5. Specify the percentage of the Snapshot copy reserve.
6. Optional: Click to modify the FlexGroup volume settings.
   a. In the **General Details** section, select the space reserve and security style, and then set the UNIX permission for the volume.

      When thin provisioning is enabled, space is allocated to the volume from the aggregate only when data is written to the volume.

      **Note:** For AFF storage systems, thin provisioning is enabled by default and for other storage systems, thick provisioning is enabled by default.

   b. In **Optimize space** section, enable fractional reserve, and then grow or shrink the volume, as required.
   c. In the **Storage Efficiency** section, enable deduplication on the volume.
   d. Click **Apply** to update the changes.
7. Click **Save** to save the changes.

### Related reference

*Volumes window* on page 203

You can use the Volumes window to manage your volumes, and to display information about these volumes.

## Resizing FlexGroup volumes

You can use System Manager to resize a FlexGroup volume by resizing existing resources or by adding new resources.

### Before you begin

- To resize a FlexGroup volume, there must be enough free space on the existing aggregates.
- To expand a FlexGroup volume, there must be enough free space on the aggregate that you are using for expansion.

### Steps

1. Click **Storage** > **Volumes**.
2. From the **Volumes on SVM** list, select **All SVMs**.
3. Select the FlexGroup volume that you want to resize, and then click **More Actions** > **Resize**.
4. In the **Resize FlexGroup Volume** window, specify the size to which you want to resize the FlexGroup volume.

   By default, existing aggregates are used to resize the FlexGroup volume.

   **Note:** If you want to expand the FlexGroup volume by adding new resources, click (advanced options).

5. Specify the percentage of the Snapshot copy reserve.
6. Click **Resize** to resize the FlexGroup volume.

### Related reference

*Volumes window* on page 203

You can use the Volumes window to manage your volumes, and to display information about these volumes.

## Changing the status of a FlexGroup volume

You can use System Manager to change the status of a FlexGroup volume when you want to take a FlexGroup volume offline, bring a FlexGroup volume back online, or restrict access to a FlexGroup volume.

### About this task

System Manager does not support constituent-level management for FlexGroup volumes.

### Steps

1. Click **Storage** > **Volumes**.
2. From the **Volumes on SVM** list, select **All SVMs**.
3. Select the FlexGroup volume for which you want to modify the status.
4. Click **More Actions** > **Change status to**, and then update the FlexGroup volume status by selecting the required status.

### Related reference

*Volumes window* on page 203

You can use the Volumes window to manage your volumes, and to display information about these volumes.

## Deleting FlexGroup volumes

You can use System Manager to delete a FlexGroup volume when you no longer require the FlexGroup volume.

### Before you begin

- The junction path of the FlexGroup volume must be unmounted.
- The FlexGroup volume must be offline.

### About this task

System Manager does not support constituent level of management for FlexGroup volumes.

### Steps

1. Click **Storage** > **Volumes**.
2. From the **Volumes on SVM** list, select **All SVMs**.
3. Select the FlexGroup volume that you want to delete, and then click **Delete**.
4. Select the confirmation check box, and then click **OK**.

### Related reference

*Volumes window* on page 203

You can use the Volumes window to manage your volumes, and to display information about these volumes.

## Viewing FlexGroup volume information

You can use System Manager to view information about a FlexGroup volume. You can view a graphical representation of the space allocated, the protection status, and the performance of a FlexGroup volume.

### About this task

You can also view the Snapshot copies that are available for the FlexGroup volume, the data protection relationships for the FlexGroup volume, and the average performance metrics, read performance metrics, and write performance metrics of the FlexGroup volume based on latency, IOPS, and throughput.

### Steps

1. Click **Storage** > **Volumes**.
2. From the **Volumes on SVM** list, select **All SVMs**.
3. From the displayed list of FlexGroup volumes, select the FlexGroup volume about which you want to view information.
   The information about the FlexGroup volume, the space allocated to the FlexGroup volume, the protection status of the FlexGroup volume, and the performance information about the FlexGroup volume are displayed.
4. Click the **Show More Details** link to view more information about the FlexGroup volume.
5. Click the **Snapshot Copies** tab to view the Snapshot copies of the FlexGroup volume.
6. Click the **Data Protection** tab to view the data protection relationships for the FlexGroup volume.
7. Click the **Storage Efficiency** tab to view the storage efficiency settings.
8. Click the **Performance** tab to view the average performance metrics, read performance metrics, and write performance metrics of the FlexGroup volume based on latency, IOPS, and throughput.

### Related reference

*Volumes window* on page 203
You can use the Volumes window to manage your volumes, and to display information about these volumes.

## What NetApp Volume Encryption is

NetApp Volume Encryption is the process of protecting the user data, including the metadata, by encrypting the data before storing it on the disk. The data is decrypted and provided to the user only after proper authentication is provided.

To encrypt data, an encryption key is required. Each volume is assigned an encryption key to encrypt/decrypt operations of its data.

## Snapshot configuration

You can configure Snapshot copies by setting a schedule for an existing Snapshot policy. Starting with ONTAP 9.4, you can have less than 1024 Snapshot copies of a FlexVol volume.

## How volume guarantees work for FlexVol volumes

Volume guarantees (sometimes called *space guarantees*) determine how space for a volume is allocated from its containing aggregate—whether or not the space is preallocated for the volume.

The guarantee is an attribute of the volume.

You set the guarantee when you create a new volume; you can also change the guarantee for an existing volume, provided that sufficient free space exists to honor the new guarantee.

Volume guarantee types can be **`volume`** (the default type) or **`none`**.

- A guarantee type of **`volume`** allocates space in the aggregate for the entire volume when you create the volume, regardless of whether that space is used for data yet.
  The allocated space cannot be provided to or allocated for any other volume in that aggregate.
- A guarantee of **`none`** allocates space from the aggregate only as it is needed by the volume.
  The amount of space consumed by volumes with this guarantee type grows as data is added instead of being determined by the initial volume size, which might leave space unused if the volume data does not grow to that size. The maximum size of a volume with a guarantee of **`none`** is not limited by the amount of free space in its aggregate. It is possible for the total size of all volumes associated with an aggregate to exceed the amount of free space for the aggregate, although the amount of space that can actually be used is limited by the size of aggregate.
  Writes to LUNs or files (including space-reserved LUNs and files) contained by that volume could fail if the containing aggregate does not have enough available space to accommodate the write.

When space in the aggregate is allocated for a **`volume`** guarantee for an existing volume, that space is no longer considered free in the aggregate, even if the volume is not yet using the space. Operations that consume free space in the aggregate, such as creation of aggregate Snapshot copies or creation of new volumes in the containing aggregate, can occur only if there is enough available free space in that aggregate; these operations are prevented from using space already allocated to another volume.

When the free space in an aggregate is exhausted, only writes to volumes or files in that aggregate with preallocated space are guaranteed to succeed.

Guarantees are honored only for online volumes. If you take a volume offline, any allocated but unused space for that volume becomes available for other volumes in that aggregate. When you try to bring that volume back online, if there is insufficient available space in the aggregate to fulfill its guarantee, it will remain offline. You must force the volume online, at which point the volume's guarantee will be disabled.

### Related information

*NetApp Technical Report 3965: NetApp Thin Provisioning Deployment and Implementation Guide Data ONTAP 8.1 (7-Mode)*

## FlexClone volumes and space guarantees

A FlexClone volume inherits its initial space guarantee from its parent volume. For example, if you create a FlexClone volume from a parent volume with a space guarantee of **`volume`**, then the

FlexClone volume's initial space guarantee will be **`volume`** also. You can change the FlexClone volume's space guarantee.

For example, suppose that you have a 100-MB FlexVol volume with a space guarantee of **`volume`**, with 70 MB used and 30 MB free, and you use that FlexVol volume as a parent volume for a new FlexClone volume. The new FlexClone volume has an initial space guarantee of **`volume`**, but it does not require a full 100 MB of space from the aggregate, as it would if you had copied the volume. Instead, the aggregate needs to allocate only 30 MB (100 MB minus 70 MB) of free space to the clone.

If you have multiple clones with the same parent volume and a space guarantee of **`volume`**, they all share the same shared parent space with each other, so the space savings are even greater.

**Note:** The shared space depends on the existence of the shared Snapshot copy (the base Snapshot copy that was used to create the FlexClone volume). If you delete this shared Snapshot copy, you lose the space savings provided by the FlexClone volume.

## Thin provisioning for greater efficiencies using FlexVol volumes

With thin provisioning, when you create volumes and LUNs in a given aggregate, you do not actually allocate any space for those in advance. The space is allocated as data is written to the volumes or LUNs.

The unused aggregate space is available to other volumes and LUNs. By allowing as-needed provisioning and space reclamation, thin provisioning can improve storage utilization and decrease storage costs.

A FlexVol volume can share its containing aggregate with other FlexVol volumes. Therefore, a single aggregate is the shared source of all the storage used by the FlexVol volumes it contains. Flexible volumes are no longer bound by the limitations of the disks on which they reside. A FlexVol volume can be sized based on how much data you want to store in it, rather than on the size of your disk. This flexibility enables you to maximize the performance and capacity utilization of the storage systems. Because FlexVol volumes can access all available physical storage in the system, improvements in storage utilization are possible.

### Example

A 500-GB volume is allocated with only 100 GB of actual data; the remaining 400 GB allocated has no data stored in it. This unused capacity is assigned to a business application, even though the application might not need all 400 GB until later. The allocated but unused 400 GB of excess capacity is temporarily wasted.

With thin provisioning, the storage administrator provisions 500 GB to the business application but uses only 100 GB for the data. The difference is that with thin provisioning, the unused 400 GB is still available to other applications. This approach allows the application to grow transparently, and the physical storage is fully allocated only when the application needs it. The rest of the storage remains in the free pool to be used as needed.

## Using space reservations with FlexVol volumes

Using space reservation, you can provision FlexVol volumes. Thin provisioning appears to provide more storage than is actually available from a given aggregate, as long as not all of that storage is currently being used.

Thick provisioning sets aside enough storage from the aggregate to ensure that any block in the volume can be written to at any time.

Aggregates can provide storage to volumes contained by more than one storage virtual machine (SVM). If you are using thin provisioning, and you need to maintain strict separation between

your SVMs (for example, if you are providing storage in a multi-tenancy environment), you should either use fully allocated volumes (thick provisioning) or ensure that your aggregates are not shared between tenants.

When the space reserve is set to "Default", the ONTAP space reservation settings apply to the volumes.

**Related information**

*NetApp Technical Report 3563: NetApp Thin Provisioning Increases Storage Utilization With On Demand Allocation*

*NetApp Technical Report 3483: Thin Provisioning in a NetApp SAN or IP SAN Enterprise Environment*

## Benefits of storage efficiency

Storage efficiency enables you to store the maximum amount of data for the lowest cost and accommodate rapid data growth while consuming less space. You can use technologies such as RAID-DP, FlexVol, Snapshot copies, deduplication, data compression, SnapMirror, and FlexClone to increase storage utilization and decrease storage costs. When used together, these technologies help to achieve increased performance.

- High-density disk drives, such as serial advanced technology attachment (SATA) drives mitigated with RAID-DP technology, provide increased efficiency and read performance.
- RAID-DP is a double-parity RAID6 implementation that protects against dual disk drive failures.
- Thin provisioning enables you to maintain a common unallocated storage space that is readily available to other applications as required.

  It is based on FlexVol technology.
- Snapshot copies are a point-in-time, read-only view of a data volume, which consume minimal storage space.

  Two Snapshot copies created in sequence differ only by the blocks added or changed in the time interval between the two. This block incremental behavior limits the associated consumption of storage capacity.
- Deduplication saves storage space by eliminating redundant data blocks within a FlexVol volume.
- Data compression stores more data in less space and reduces the time and bandwidth required to replicate data during volume SnapMirror transfers.

  You have to choose the type of compression (inline or background) based on your requirement and the configurations of your storage system. Inline compression checks if data can be compressed, compresses data, and then writes data to the volume. Background compression runs on all the files, irrespective of whether the file is compressible or not, after all the data is written to the volume.
- SnapMirror technology is a flexible solution for replicating data over local area, wide area, and Fibre Channel networks.

  It can serve as a critical component in implementing enterprise data protection strategies. You can replicate your data to one or more storage systems to minimize downtime costs in case of a production site failure. You can also use SnapMirror technology to centralize the backup of data to disks from multiple data centers.
- FlexClone technology copies data volumes, files, and LUNs as instant virtual copies.

  A FlexClone volume, file, or LUN is a writable point-in-time image of the FlexVol volume or another FlexClone volume, file, or LUN. This technology enables you to use space efficiently, storing only data that changes between the parent and the clone.
- The unified architecture integrates multiprotocol support to enable both file-based and block-based storage on a single platform.

With FlexArray Virtualization, you can virtualize your entire storage infrastructure under one interface, and you can apply all the preceding efficiencies to your non-NetApp systems.

## Data compression and deduplication

Beginning with Data ONTAP 8.0.1, data compression is supported with deduplication.

When both data compression and deduplication are enabled on a FlexVol volume, the data is first compressed and then deduplicated. Therefore, deduplication can further increase the space savings by removing duplicate blocks in the FlexVol volume.

Though data compression and deduplication can be enabled on a FlexVol volume, the savings might not be the sum of the savings when each is run individually on a data set. The combined savings can yield higher savings than running deduplication or data compression individually.

You can achieve better savings when you run the data compression scanner before deduplication. This is because data compression scanner cannot run on data that is locked by deduplication, but deduplication can run on compressed data.

The following illustration shows how data is first compressed and then deduplicated:

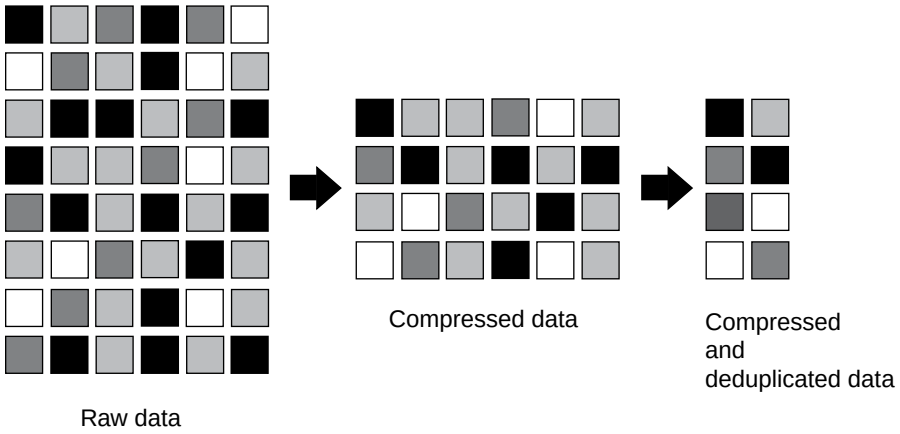


When you run deduplication on a FlexVol volume that contains uncompressed data, it scans all the uncompressed blocks in the FlexVol volume and creates a digital fingerprint for each of the blocks.

**Note:** If a FlexVol volume has compressed data, but the compression option is disabled on that volume, then you might lose the space savings when you run the `sis undo` command.

## Guidelines for using deduplication

You must remember certain guidelines about system resources and free space when using deduplication.

The guidelines are as follows:

- If you have a performance-sensitive solution, you must carefully consider the performance impact of deduplication and measure the impact in a test setup before using deduplication.
- Deduplication is a background process that consumes system resources while it is running.

  If the data does not change very often in a FlexVol volume, it is best to run deduplication less frequently. Multiple concurrent deduplication operations running on a storage system lead to a higher consumption of system resources.
- You must ensure that sufficient free space exists for deduplication metadata in the volumes and aggregates.
- If deduplication is used on the source volume, you must use deduplication on the destination volume.

- You must use automatic mode when possible so that deduplication runs only when significant additional data has been written to each FlexVol volume.
- You must run deduplication before creating a Snapshot copy to obtain maximum savings.
- You must set the Snapshot reserve to greater than 0 if Snapshot copies are used.

## Options for resizing volumes

You can use the Volume Resize wizard to change your volume size, adjust the Snapshot reserve, delete Snapshot copies, and dynamically view the results of your changes.

The Volume Resize wizard displays a bar graph that displays the current space allocations within the volume, including the amount of used and free space. When you make changes to the size or Snapshot reserve of the volume, this graph is updated dynamically to reflect the changes.

You can also use the **Calculate space** button to determine the amount of space that is freed by deleting selected Snapshot copies.

You can use the Volume Resize wizard to make the following changes to your volume:

**Change the volume size**

You can change the total volume size to increase or decrease storage space.

**Adjust Snapshot reserve**

You can adjust the amount of space reserved for Snapshot copies to increase or decrease storage space.

**Delete Snapshot copies**

You can delete Snapshot copies to reclaim volume space.

**Note:** Snapshot copies that are in use cannot be deleted.

**Autogrow**

You can specify the limit to which the volume can be grown automatically, if required.

## Considerations when moving volumes

Moving a volume has many considerations and recommendations that are influenced by the volume you are moving or by the system configuration. You should understand the considerations associated with moving volumes.

- If you move a volume that has inline deduplication enabled from an aggregate with All Flash Optimized personality or a Flash Pool aggregate to an HDD aggregate, inline deduplication is disabled on the volume.
- If you move a volume that has background deduplication and inline compression enabled from an aggregate with All Flash Optimized personality to an HDD aggregate, then background compression, background deduplication, and inline compression are automatically enabled on the volume.
- If you move a volume that has background compression enabled from an HDD aggregate to an aggregate with All Flash Optimized personality, background compression is disabled on the volume.
- If you move a volume from a Flash Pool aggregate to a non-Flash Pool aggregate, the caching policies and retention priority are disabled.
- If you move a volume from a non-Flash Pool aggregate to a Flash Pool aggregate, the **`default`** caching policy and the **`default`** retention priority are automatically assigned to the volume.

## Volumes window

You can use the Volumes window to manage your volumes, and to display information about these volumes.

You cannot view or manage volumes that are in storage virtual machines (SVMs) that are configured for disaster recovery (DR) by using System Manager. You must use the command-line interface (CLI) instead.

**Note:** The command buttons and list of columns will differ based on the type of volume that is selected. You will be able to view only those command buttons and columns that are applicable for the selected volume.

### Command buttons

**Create**

Provides the following options:

**Create FlexVol**

Opens the Create Volume dialog box, which enables you to add FlexVol volumes.

**Create FlexGroup**

Opens the Create FlexGroup window, which enables you to create FlexGroup volumes.

**Edit**

Enables you to edit the properties of the selected volume.

**Delete**

Deletes the selected volume or volumes.

**More Actions**

Provides the following options:

**Change status to**

Changes the status of the selected volume to one of the following statuses:

- Online
- Offline
- Restrict

**Resize**

Enables you to change the size of the volume.

For FlexGroup volumes, you can use existing resources to resize the volumes or you can add new resources to expand the volumes.

**Protect**

Opens the Create Protection Relationship window for the volumes that are selected as source.

**Manage Snapshots**

Provides a list of Snapshot options, including the following:

- Create
  Displays the Create Snapshot dialog box, which you can use to create a Snapshot copy of the selected volume.
- Configuration Settings
  Configures the Snapshot settings.
- Restore

Restores a Snapshot copy of the selected volume.

**Clone**

Provides a list of clone options, including the following:

- Create
  Creates a clone of the selected volume or a clone of a file from the selected volume.
- Split
  Splits the clone from the parent volume.
- View Hierarchy
  Displays information about the clone hierarchy.

**Storage Efficiency**

Opens the Storage Efficiency dialog box, which you can use to manually start deduplication or to abort a running deduplication operation. This button is displayed only if deduplication is enabled on the storage system.

**Move**

Opens the Move Volume dialog box, which you can use to move volumes from one aggregate or node to another aggregate or node within the same SVM.

**Storage QoS**

Opens the Quality of Service details dialog box, which you can use to assign one or more volumes to a new or existing policy group.

**Provision Storage for VMware**

Enables you to create a volume for the NFS datastore and to specify the ESX servers that can access the NFS datastore.

**View Missing Protection Relationship**

Displays read/write volumes that are online and that are not protected, and volumes that have protection relationships and that are not initialized.

**Reset Filters**

Enables you to reset the filters that were set to view missing protection relationships.

**Refresh**

Updates the information in the window.

### Volume list

**Status**

Displays the status of the volume.

**Name**

Displays the name of the volume.

**Style**

Displays the type of the volume.

**SVM**

Displays the storage virtual machine (SVM) that contains the volume.

**Aggregates**

Displays the name of the aggregates belonging to the volume.

**Thin Provisioned**

Displays whether space guarantee is set for the selected volume. Valid values for online volumes are **`Yes`** and **`No`**.

**Root volume**

Displays whether the volume is a root volume.

**Available Space**

Displays the available space in the volume.

**Total Space**

Displays the total space in the volume, which includes the space that is reserved for Snapshot copies.

**% Used**

Displays the amount of space (in percentage) that is used in the volume.

**Logical Used %**

Displays the amount of logical space (in percentage), including space reserves that is used in the volume irrespective of whether logical space reporting is enabled or not.

> **Note:** This field is displayed only if you have enabled logical space reporting using the CLI.

**Logical Space Reporting**

Displays whether logical space reporting is enabled on the volume.

> **Note:** This field is displayed only if you have enabled logical space reporting using the CLI.

**Type**

Displays the type of volume: **`rw`** for read/write, **`ls`** for load sharing, or **`dp`** for data protection.

**Protection Relationship**

Display whether the volume has a protection relationship initiated.

If the relationship is between an ONTAP system and a non-ONTAP system, the value is displayed as **`No`** by default.

**Storage Efficiency**

Displays whether deduplication is enabled or disabled for the selected volume.

**Encrypted**

Displays whether the volume is encrypted or not.

**QoS Policy Group**

Displays the name of the Storage QoS policy group to which the volume is assigned. By default, this column is hidden.

**SnapLock Type**

Displays the SnapLock type of the volume.

**Clone**

Displays whether the volume is a FlexClone volume.

**Is Volume Moving**

Displays whether a volume is being moved from one aggregate to another aggregate, or from one node to another node.

**Tiering Policy**

Displays the tiering policy of a FabricPool-enabled aggregate. The default tiering policy is "snapshot-only".

**Application**

Displays the name of the application that is assigned to the volume.

### Details area

You can select a volume to view information about the selected volume. You can click Show More Details to view detailed information about the selected volume.

**Overview tab**

Displays general information about the selected volume, and displays a pictorial representation of the space allocation of the volume, the protection status of the volume, and the performance of the volume. The Overview tab also displays information about a volume that is being moved, such as the state and phase of the volume move, the destination node and aggregate to which the volume is being moved, the percentage of volume move that is complete, the estimated time to complete the volume move operation, and details of the volume move operation.

The refresh interval for the performance data is 15 seconds.

This tab contains the following command button:

**Cutover**

Opens the Cutover dialog box, which enables you to manually trigger the cutover.

The **Cutover** command button is displayed only if the volume move operation is in the "replication" or "hard deferred" state.

**Snapshot Copies tab**

Displays the Snapshot copies of the selected volume. This tab contains the following command buttons:

**Create**

Opens the Create Snapshot Copy dialog box, which enables you to create a Snapshot copy of the selected volume.

**Configuration Settings**

Configures the Snapshot settings.

**Delete**

Deletes the selected Snapshot copy.

**More Actions > Rename**

Opens the Rename Snapshot Copy dialog box, which enables you to rename a selected Snapshot copy.

**More Actions > Restore**

Restores a Snapshot copy.

**Extend Expiry Date**

Extends the expiry date of a Snapshot copy.

**Refresh**

Updates the information in the window.

**Data Protection tab**

Displays data protection information about the selected volume.

If the source volume (read/write volume) is selected, the tab displays all of the mirror relationships, vault relationships, and mirror and vault relationships that are related to the destination volume (DP volume). If the destination volume is selected, the tab displays the relationship with the source volume.

If some or all of the cluster peer relationships of the local cluster are in an unhealthy state, the Data Protection tab might take some time to display the protection relationships relating to a healthy cluster peer relationship. Relationships relating to unhealthy cluster peer relationships are not displayed.

**Storage Efficiency tab**

Displays information in the following panes:

- Bar graph
  Displays (in graphical format) the volume space that is used by data and Snapshot copies. You can view details about the space used before and after applying settings for storage efficiency savings.
- Details
  Displays information about deduplication properties, including whether deduplication is enabled on the volume, the deduplication mode, the deduplication status, type, and whether inline or background compression is enabled on the volume.
- Last run details
  Provides details about the last-run deduplication operation on the volume. Space savings resulting from compression and deduplication operations that are applied on the data on the volume are also displayed.

**Performance tab**

Displays information about the average performance metrics, read performance metrics, and write performance metrics of the selected volume, including throughput, IOPS, and latency.

Changing the client time zone or the cluster time zone impacts the performance metrics graphs. You must refresh your browser to view the updated graphs.

### Related tasks

*Creating FlexVol volumes* on page 51
You can create a FlexVol volume for your data by using the Create Volume dialog box in System Manager. You should always create a separate volume for your data rather than storing data in the root volume.

*Creating FlexClone volumes* on page 176
You can use System Manager to create a FlexClone volume when you require a writable, point-in-time copy of an existing FlexVol volume. You might want to create a copy of a volume for testing or to provide access to the volume for additional users without giving them access to the production data.

*Creating FlexClone files* on page 177
You can use System Manager to create a FlexClone file, which is a writable copy of a parent file. You can use these copies to test applications.

*Deleting volumes* on page 176
You can use System Manager to delete a volume when you no longer require the data that a volume contains, or if you have copied the data that a volume contains to another location. When you delete a volume, all the data in the volume is destroyed, and you cannot recover this data.

*Setting the Snapshot copy reserve* on page 180

You can use System Manager to reserve space (in percentage) for the Snapshot copies in a volume. By setting the Snapshot copy reserve, you can allocate enough disk space for the Snapshot copies so that they do not consume the active file system space.

*Deleting Snapshot copies* on page 183
You can delete a Snapshot copy to conserve disk space or to free disk space by using System Manager. You can also delete a Snapshot copy if the Snapshot copy is no longer required.

*Creating Snapshot copies outside a defined schedule* on page 180
You can use System Manager to create a Snapshot copy of a volume outside a defined schedule to capture the state of the file system at a specific point in time.

*Editing volume properties* on page 174
You can modify volume properties such as the volume name, security style, fractional reserve, and space guarantee by using System Manager. You can modify storage efficiency settings (deduplication schedule, deduplication policy, and compression) and space reclamation settings.-

*Changing the status of a volume* on page 179
You can use System Manager to change the status of a volume when you want to take a volume offline, bring a volume back online, or restrict access to a volume.

*Enabling storage efficiency on a volume* on page 184
You can use System Manager to enable storage efficiency and to configure both deduplication and data compression or only deduplication on a volume to save storage space. If you have not enabled storage efficiency when you created the volume, you can do so later by editing the volume.

*Changing the deduplication schedule* on page 185
You can use System Manager to change the deduplication schedule by choosing to run deduplication manually, automatically, or on a schedule that you specify.

*Running deduplication operations* on page 186
You can use System Manager to run deduplication immediately after creating a volume, or to schedule deduplication to run at a specified time.

*Splitting a FlexClone volume from its parent volume* on page 178
If you want a FlexClone volume to have its own disk space instead of using the disk space of its parent volume, you can split the volume from its parent by using System Manager. After the split, the FlexClone volume becomes a normal FlexVol volume.

*Resizing volumes* on page 184
When your volume reaches nearly full capacity, you can increase the size of the volume, delete some Snapshot copies, or adjust the Snapshot reserve. You can use the Volume Resize wizard in System Manager to provide more free space.

*Restoring a volume from a Snapshot copy* on page 181
You can use System Manager to restore a volume to a state that is recorded in a previously created Snapshot copy to retrieve lost information. When you restore a volume from a Snapshot copy, the restore operation overwrites the existing volume configuration. Any changes that were made to the data in the volume after the Snapshot copy was created are lost.

*Scheduling automatic creation of Snapshot copies* on page 181
You can use System Manager to set up a schedule for the automatic creating automatic Snapshot copies of a volume. You can specify the time and frequency of creating the copies. You can also specify the number of Snapshot copies that are saved.

*Renaming Snapshot copies* on page 183
You can use System Manager to rename a Snapshot copy to help you organize and manage your Snapshot copies.

*Hiding the Snapshot copy directory* on page 181

You can use System Manager to hide the Snapshot copy directory (`.snapshot`) so that the Snapshot copy directory is not visible when you view your volume directories. By default, the `.snapshot` directory is visible.

*Viewing the FlexClone volume hierarchy* on page 178
You can use System Manager to view the hierarchy of FlexClone volumes and their parent volumes.

*Creating FlexGroup volumes* on page 193
You can use System Manager to create a FlexGroup volume by selecting specific aggregates or by selecting system-recommended aggregates.

*Editing FlexGroup volumes* on page 194
You can use System Manager to edit the properties of an existing FlexGroup volume.

*Resizing FlexGroup volumes* on page 195
You can use System Manager to resize a FlexGroup volume by resizing existing resources or by adding new resources.

*Changing the status of a FlexGroup volume* on page 196
You can use System Manager to change the status of a FlexGroup volume when you want to take a FlexGroup volume offline, bring a FlexGroup volume back online, or restrict access to a FlexGroup volume.

*Deleting FlexGroup volumes* on page 196
You can use System Manager to delete a FlexGroup volume when you no longer require the FlexGroup volume.

*Viewing FlexGroup volume information* on page 197
You can use System Manager to view information about a FlexGroup volume. You can view a graphical representation of the space allocated, the protection status, and the performance of a FlexGroup volume.

# Junction Path

You can use the Junction Path window in System Manager to mount or unmount FlexVol volumes to a junction in the SVM namespace.

## Mounting volumes

You can use System Manager to mount volumes to a junction in the storage virtual machine (SVM) namespace.

### About this task

- If you mount a volume to a junction path with a language setting that is different from that of the immediate parent volume in the path, NFSv3 clients cannot access some of the files because some characters might not be decoded correctly.

  This issue does not occur if the immediate parent directory is the root volume.
- You can mount a SnapLock volume only under the root of the SVM.
- You cannot mount a regular volume under a SnapLock volume.

### Steps

1. Click **Storage** > **Junction Path**.
2. From the **Junction Path on SVM** list, select the SVM on which you want to mount a volume.
3. Click **Mount**, and then select the volume that is to be mounted.
4. Optional: If you want to change the default junction name, specify a new name.
5. Click **Browse**, and then select the junction path to which you want to mount the volume.

6. Click **OK**, and then click **Mount**.
7. Verify the new junction path in the **Details** tab.

## Unmounting FlexVol volumes

You can use the Junction Path option of Storage pane in System Manager to unmount FlexVol volumes from a junction in the storage virtual machine (SVM) namespace.

### Steps

1. Click **Storage** > **Junction Path**.
2. From the **Junction Path on SVM** list, select the SVM from which you want to unmount a volume.
3. Select the volumes that have to be unmounted, and then click **Unmount**.
4. Select the confirmation check box, and then click **Unmount**.

## Changing export policies

When a volume is created, the volume automatically inherits the default export policy of the root volume of the storage virtual machine (SVM). You can use System Manager to change the default export policy that is associated with the volume to redefine the client access to data.

### Steps

1. Click **Storage** > **Junction Path**.
2. From the **Junction Path on SVM** list, select the SVM in which the volume that you want to modify resides.
3. Select the volume, and then click **Change Export Policy**.
4. Select the export policy, and then click **Change**.
5. Verify that the **Export Policy** column in the **Junction Path** window displays the export policy that you applied to the volume.

### Result

The default export policy is replaced with the export policy that you selected.

## Junction Path window

You can use the Junction Path menu to manage the NAS namespace of storage virtual machines (SVMs).

### Command buttons

**Mount**

Opens the Mount Volume dialog box, which enables you to mount a volume to the junction in an SVM namespace.

**Unmount**

Opens the Unmount Volume dialog box, which enables you to unmount a volume from its parent volume.

**Change Export Policy**

Opens the Change Export Policy dialog box, which enables you to change the existing export policy associated with the volume.

**Refresh**

Updates the information in the window.

### Junction Path list

**Path**

Specifies the junction path of the mounted volume. You can click the junction path to view the related volumes and qtrees.

**Storage Object**

Specifies the name of the volume mounted on the junction path. You can also view the qtrees that the volume contains.

**Export Policy**

Specifies the export policy of the mounted volume.

**Security Style**

Specifies the security style for the volume. Possible values include UNIX (for UNIX mode bits), NTFS (for CIFS ACLs), and Mixed (for mixed NFS and CIFS permissions).

### Details tab

Displays general information about the selected volume or qtree, such as the name, type of storage object, junction path of the mounted object, and export policy. If the selected object is a qtree, details about the space hard limit, space soft limit, and space usage are displayed.

# Shares

You can use System Manager to create, edit, and manage shares.

## Creating a CIFS share

You can use System Manager to create a CIFS share that enables you to specify the folder, qtree, or volume that CIFS users can access.

### Before you begin

You must have installed the CIFS license before you set up and start CIFS.

### Steps

1. Click **Storage** > **Shares**.
2. From the **Shares on SVM** list, select the SVM on which you want to create a CIFS share.
3. Click **Create Share**.
4. In the **Create Share** window, click **Browse**, and then select the folder, qtree, or volume that should be shared.
5. Specify a name for the new CIFS share.
6. Optional: Select the **Enable continuous availability for Hyper-V and SQL** check box to permit SMB 3.0 and later clients that support SMB 3.0 to open files persistently during nondisruptive operations.

   Files that are opened by using this option are protected from disruptive events such as failover, giveback, and LIF migration.
7. Select the **Encrypt data while accessing this share** check box to enable SMB 3.0 encryption.
8. Provide a description or comment for the share, and then click **Create**.

### Result

The CIFS share is created with the access permissions set to "Full Control for Everyone" in the group.

### Related tasks

*Setting up CIFS* on page 247

You can use System Manager to enable and configure CIFS servers to allow CIFS clients to access the files on the cluster.

### Related reference

*Shares window* on page 214
You can use the Shares window to manage your shares and to view information about the shares.

## Stopping share access

You can use System Manager to stop a share when you want to remove the shared network access to a folder, qtree, or volume.

### Before you begin

You must have installed the CIFS license.

### Steps

1. Click **Storage** > **Shares**.
2. From the **Shares on SVM** list, select the SVM on which the CIFS share that you want to stop resides.
3. From the list of shares, select the share that you want to stop sharing, and then click **Stop Sharing**.
4. Select the confirmation check box, and then click **Stop**.
5. Verify that the share is no longer listed in the **Shares** window.

### Related reference

*Shares window* on page 214
You can use the Shares window to manage your shares and to view information about the shares.

## Creating home directory shares

You can use System Manager to create a home directory share and to manage home directory search paths.

### Before you begin

CIFS must be set up and started.

### Steps

1. Click **Storage** > **Shares**.
2. Click **Create Home Directory**, and then provide the pattern information that determines how a user is mapped to a directory.
3. Click **Create**.
4. Verify that the home directory that you created is listed in the **Shares** window.

## Editing share settings

You can use System Manager to modify the settings of a share such as the symbolic link settings, share access permissions of users or groups, and the type of access to the share. You can also enable or disable continuous availability of a share over Hyper-V, and enable or disable access-based enumeration (ABE).

### Steps

1. Click **Storage** > **Shares**.
2. Select the share that you want to modify from the list of shares, and then click **Edit**.
3. In the **Edit Share Settings** dialog box, modify the share settings as required:

a. In the **General** tab, enable continuous availability of a share over Hyper-V.

   Enabling continuous availability permits SMB 3.0 and clients that support SMB 3.0 to open files persistently during nondisruptive operations. Files that are opened persistently are protected from disruptive events such as failover, giveback, and LIF migration.

b. In the **Permissions** tab, add users or groups, and then assign permissions to specify the type of access.

c. In the **Options** tab, select the required options.

4. Click **Save and Close**.
5. Verify the changes that you made to the selected share in the **Shares** window.

### Related reference

*Shares window* on page 214

You can use the Shares window to manage your shares and to view information about the shares.

## How ONTAP enables dynamic home directories

ONTAP home directories enable you to configure an SMB share that maps to different directories based on the user that connects to it and a set of variables. Instead of creating separate shares for each user, you can configure one share with a few home directory parameters to define a user's relationship between an entry point (the share) and the home directory (a directory on the SVM).

A user that is logged in as a guest user does not have a home directory and cannot access other users' home directories. There are four variables that determine how a user is mapped to a directory:

**Share name**

This is the name of the share that you create to which the user connects. You must set the home directory property for this share.

The share name can use the following dynamic names:

- *%w* (the user's Windows user name)
- *%d* (the user's Windows domain name)
- *%u* (the user's mapped UNIX user name)

To make the share name unique across all home directories, the share name must contain either the *%w* or the *%u* variable. The share name can contain both the *%d* and the *%w* variable (for example, *%d*/*%w*), or the share name can contain a static portion and a variable portion (for example, home_*%w*).

**Share path**

This is the relative path, which is defined by the share and is therefore associated with one of the share names, that is appended to each search path to generate the user's entire home directory path from the root of the SVM. It can be static (for example, `home`), dynamic (for example, `%w`), or a combination of the two (for example, `eng/%w`).

**Search paths**

This is the set of absolute paths from the root of the SVM that you specify that directs the ONTAP search for home directories. You can specify one or more search paths by using the `vserver cifs home-directory search-path add` command. If you specify multiple search paths, ONTAP tries them in the order specified until it finds a valid path.

**Directory**

This is the user's home directory that you create for the user. The directory name is usually the user's name. You must create the home directory in one of the directories that are defined by the search paths.

As an example, consider the following setup:

- User: John Smith
- User domain: acme
- User name: jsmith
- SVM name: vs1
- Home directory share name #1: home_`%w` - share path: `%w`
- Home directory share name #2: `%w` - share path: `%d/%w`
- Search path #1: `/aggr0home/home`
- Search path #2: `/aggr1home/home`
- Search path #3: `/aggr2home/home`
- Home directory: `/aggr1home/home/jsmith`

Scenario 1: The user connects to `\\vs1\home_jsmith`. This matches the first home directory share name and generates the relative path `jsmith`. ONTAP now searches for a directory named `jsmith` by checking each search path in order:

- `/aggr0home/home/jsmith` does not exist; moving on to search path #2.
- `/aggr1home/home/jsmith` does exist; therefore, search path #3 is not checked; the user is now connected to his home directory.

Scenario 2: The user connects to `\\vs1\jsmith`. This matches the second home directory share name and generates the relative path `acme/jsmith`. ONTAP now searches for a directory named `acme/jsmith` by checking each search path in order:

- `/aggr0home/home/acme/jsmith` does not exist; moving on to search path #2.
- `/aggr1home/home/acme/jsmith` does not exist; moving on to search path #3.
- `/aggr2home/home/acme/jsmith` does not exist; the home directory does not exist; therefore, the connection fails.

## Shares window

You can use the Shares window to manage your shares and to view information about the shares.

- *Command buttons*
- *Shares list*
- *Details area*

### Command buttons

**Create Share**

Opens the Create Share dialog box, which enables you to create a share.

**Create Home Directory**

Opens the Create Home Directory Share dialog box, which enables you to create a new home directory share.

**Edit**

Opens the Edit Settings dialog box, which enables you to modify the properties of a selected share.

**Stop Sharing**

Stops the selected object from being shared.

**Refresh**

Updates the information in the window.

### Shares list

The shares list displays the name and path of each share.

**Share Name**

Displays the name of the share.

**Path**

Displays the complete path name of an existing folder, qtree, or volume that is shared. Path separators can be backward slashes or forward slashes, although ONTAP displays all path separators as forward slashes.

**Home Directory**

Displays the name of the home directory share.

**Comment**

Displays additional descriptions of the share, if any.

**Continuously Available Share**

Displays whether the share is enabled for continuous availability.

### Details area

The area below the shares list displays the share properties and the access rights for each share.

**Properties**

- Name
  Displays the name of the share.
- Oplocks status
  Specifies whether the share uses opportunistic locks (oplocks).
- Browsable
  Specifies whether the share can be browsed by Windows clients.
- Show Snapshot
  Specifies whether Snapshot copies can be viewed by clients.
- Continuously Available Share
  Specifies whether the share is enabled or disabled for continuous availability.
- Access-Based Enumeration
  Specifies whether access-based enumeration (ABE) is enabled or disabled on the share.
- BranchCache
  Specifies whether BranchCache is enabled or disabled on the share.
- SMB Encryption
  Specifies whether data encryption using SMB 3.0 is enabled at the storage virtual machine (SVM) level or at the share level. If SMB encryption is enabled at the SVM level, SMB encryption applies for all of the shares and the value is shown as Enabled (at the SVM level).
- Previous Versions
  Specifies whether the previous versions can be viewed and restored from the client.

**Share access control**

Displays the access rights of the domain users, domain groups, local users, and local groups for the share.

### Related tasks

*Creating a CIFS share* on page 211

You can use System Manager to create a CIFS share that enables you to specify the folder, qtree, or volume that CIFS users can access.

*Stopping share access* on page 212
You can use System Manager to stop a share when you want to remove the shared network access to a folder, qtree, or volume.

*Editing share settings* on page 212
You can use System Manager to modify the settings of a share such as the symbolic link settings, share access permissions of users or groups, and the type of access to the share. You can also enable or disable continuous availability of a share over Hyper-V, and enable or disable access-based enumeration (ABE).

# LUNs

You can use System Manager to manage LUNs.

You can access all the LUNs in the cluster by using the LUNs tab or you can access the LUNs specific to the SVM by using **SVMs** > **LUNs**.

**Note:** The LUNs tab is displayed only if you have enabled the FC/FCoE and iSCSI licenses.

### Related information

*SAN administration*

## Creating FC SAN optimized LUNs

You can use System Manager to create one or more FC SAN optimized LUNs during the initial setup of a cluster on an AFF platform.

### Before you begin

- You must ensure that only one storage virtual machine (SVM) has been created with the name AFF_SAN_DEFAULT_SVM, and that this SVM does not contain any LUNs.
- You must have verified that the hardware setup has been completed successfully.
  *NetApp Documentation: ONTAP 9*

### About this task

- This method is available only during the initial setup of a cluster with two or more nodes.
  System Manager uses only the first two nodes to create LUNs.
- Each LUN is created on a separate volume.
- Volumes are thin provisioned.
- Space reservation is disabled on the created LUNs.
- Most of the cluster configurations are already completed at the factory and are optimized for optimum storage efficiency and performance.
  You must not modify these configurations.

### Steps

1. Log in to System Manager by using your cluster administrator credentials.
   After you create LUNs using this method, you cannot use this method again.

   If you close the dialog box without creating LUNs, you must navigate to the LUNs tab and click **Create** to access the dialog box again.
2. In the **LUN details** area of the **Create LUNs** dialog box, specify the application type:

| If the application type is... | Then... |
|---|---|
| Oracle | **a.** Specify the database name and size.<br>**b.** If you have deployed Oracle Real Application Clusters (RAC), then select the **Oracle RAC** check box.<br>Only two RAC nodes are supported. You must ensure that Oracle RAC has a minimum of two initiators added to the initiator group. |
| SQL | Specify the number of databases and the size of each database. |
| Other | **a.** Specify the name and size of each LUN.<br>**b.** If you want to create more LUNs, click **Add more LUNs**, and then specify the name and size for each LUN. |

Data, log, binary, and temporary LUNs are created based on the selected application type.

3. In the **Map to these Initiators** area, perform these steps:
   a. Specify the initiator group name and the type of operating system.
   b. Add the host initiator WWPN by selecting it from the drop-down list or by typing the initiator in the text box.

   Only one initiator group is created.
4. Click **Create**.
   A summary table is displayed with the LUNs that are created.
5. Click **Close**.

### Related information

*NetApp Documentation: ONTAP 9*

## Application-specific LUN settings

System Manager supports Oracle, SQL, and other application types while creating FC SAN optimized LUNs on an AFF cluster. LUN settings such as the LUN size are determined by rules specific to the application type. For SQL and Oracle, LUN settings are automatically created.

If your cluster contains two or more nodes, System Manager uses only the first two nodes selected by the API to create LUNs. Data aggregates are already created in each of the two nodes. The size of each volume created is equal to the available capacity of the aggregate. The volumes are thin-provisioned and space reservation is disabled on the LUNs.

Storage efficiency policy is enabled by default with the schedule set to "daily" and quality of service (QoS) set to "best_effort". By default, access time (atime) update is enabled on the cluster. However, access time updates are disabled by System Manager while creating volumes and therefore every time a file is read or written, the access time field in the directory is not updated.

**Note:** Enabling the access time update causes performance degradation to the data-serving capability of the cluster.

### LUN settings for SQL

By default, LUNs and volumes are provisioned for a single instance of the SQL server with 2 databases of 1 TB each and 24 physical cores. Space is provisioned for LUNs and volumes according to specific rules for the SQL server. Load balancing is performed for LUNs across the HA pair. You can modify the number of databases. For each database, eight data LUNs and one log LUN is created. One temporary LUN is created for each SQL instance.

The following table provides information about how space is provisioned for the default values of SQL:

| Node | Aggregate | LUN type | Volume name | LUN name | Formula for LUN size | LUN size (GB) |
|---|---|---|---|---|---|---|
| node1 | node1_aggr1 | data | db01_data01 | db01_data01 | Database size ÷ 8 | 125 |
| | | data | db01_data02 | db01_data02 | Database size ÷ 8 | 125 |
| | | data | db01_data03 | db01_data03 | Database size ÷ 8 | 125 |
| | | data | db01_data04 | db01_data04 | Database size ÷ 8 | 125 |
| | | data | db02_data01 | db02_data01 | Database size ÷ 8 | 125 |
| | | data | db02_data02 | db02_data02 | Database size ÷ 8 | 125 |
| | | data | db02_data03 | db02_data03 | Database size ÷ 8 | 125 |
| | | data | db02_data04 | db02_data04 | Database size ÷ 8 | 125 |
| | | log | db01_log | db01_log | Database size ÷ 20 | 50 |
| | | temp | sql_temp | sql_temp | Database size ÷ 3 | 330 |
| node2 | node2_aggr1 | data | db01_data05 | db01_data05 | Database size ÷ 8 | 125 |
| | | data | db01_data06 | db01_data06 | Database size ÷ 8 | 125 |
| | | data | db01_data07 | db01_data07 | Database size ÷ 8 | 125 |
| | | data | db01_data08 | db01_data08 | Database size ÷ 8 | 125 |
| | | data | db02_data05 | db02_data05 | Database size ÷ 8 | 125 |
| | | data | db02_data06 | db02_data06 | Database size ÷ 8 | 125 |
| | | data | db02_data07 | db02_data07 | Database size ÷ 8 | 125 |
| | | data | db02_data08 | db02_data08 | Database size ÷ 8 | 125 |
| | | log | db02_log | db02_log | Database size ÷ 20 | 50 |

### LUN settings for Oracle

By default, LUNs and volumes are provisioned for one database of 2 TB. Space is provisioned for LUNs and volumes according to specific rules for Oracle. By default, Oracle Real Application Clusters (RAC) is not selected.

The following table provides information about how space is provisioned for the default values of Oracle:

| Node | Aggregate | LUN type | Volume name | LUN name | Formula for LUN size | LUN size (GB) |
|---|---|---|---|---|---|---|
| node1 | node1_aggr1 | data | ora_vol01 | ora_lundata01 | Database size ÷ 8 | 250 |
| | | data | ora_vol02 | ora_lundata02 | Database size ÷ 8 | 250 |
| | | data | ora_vol03 | ora_lundata03 | Database size ÷ 8 | 250 |
| | | data | ora_vol04 | ora_lundata04 | Database size ÷ 8 | 250 |
| | | log | ora_vol05 | ora_lunlog1 | Database size ÷ 40 | 50 |
| | | binaries | ora_vol06 | ora_orabin1 | Database size ÷ 40 | 50 |
| node2 | node2_aggr1 | data | ora_vol07 | ora_lundata05 | Database size ÷ 8 | 250 |
| | | data | ora_vol08 | ora_lundata06 | Database size ÷ 8 | 250 |
| | | data | ora_vol09 | ora_lundata07 | Database size ÷ 8 | 250 |
| | | data | ora_vol10 | ora_lundata08 | Database size ÷ 8 | 250 |
| | | log | ora_vol11 | ora_lunlog2 | Database size ÷ 40 | 50 |

For Oracle RAC, LUNs are provisioned for grid files. Only two RAC nodes are supported for Oracle RAC.

The following table provides information about how space is provisioned for the default values of Oracle RAC:

| Node | Aggregate | LUN type | Volume name | LUN name | Formula for LUN size | LUN size (GB) |
|---|---|---|---|---|---|---|
| node1 | node1_aggr1 | data | ora_vol01 | ora_lundata01 | Database size ÷ 8 | 250 |
| | | data | ora_vol02 | ora_lundata02 | Database size ÷ 8 | 250 |
| | | data | ora_vol03 | ora_lundata03 | Database size ÷ 8 | 250 |

| Node | Aggregate | LUN type | Volume name | LUN name | Formula for LUN size | LUN size (GB) |
|---|---|---|---|---|---|---|
| | | data | ora_vol04 | ora_lundata04 | Database size ÷ 8 | 250 |
| | | log | ora_vol05 | ora_lunlog1 | Database size ÷ 40 | 50 |
| | | binaries | ora_vol06 | ora_orabin1 | Database size ÷ 40 | 50 |
| | | grid | ora_vol07 | ora_lungrid1 | 10 GB | 10 |
| node2 | node2_aggr1 | data | ora_vol08 | ora_lundata05 | Database size ÷ 8 | 250 |
| | | data | ora_vol09 | ora_lundata06 | Database size ÷ 8 | 250 |
| | | data | ora_vol10 | ora_lundata07 | Database size ÷ 8 | 250 |
| | | data | ora_vol11 | ora_lundata08 | Database size ÷ 8 | 250 |
| | | log | ora_vol12 | ora_lunlog2 | Database size ÷ 40 | 50 |
| | | binaries | ora_vol13 | ora_orabin2 | Database size ÷ 40 | 50 |

### LUN settings for other application type

Each LUN is provisioned in a volume. The space is provisioned in the LUNs based on the specified size. Load balancing is performed across the nodes for all the LUNs.

## Creating LUNs

You can use System Manager to create LUNs for an existing aggregate, volume, or qtree when there is available free space. You can create a LUN in an existing volume or create a new FlexVol volume for the LUN. You can also enable storage Quality of Service (QoS) to manage the workload performance.

### About this task

If you specify the LUN ID, System Manager checks the validity of the LUN ID before adding it. If you do not specify a LUN ID, ONTAP software automatically assigns one.

While selecting the LUN multiprotocol type, you should have considered the guidelines for using each type. The LUN Multiprotocol Type, or operating system type, determines the layout of data on the LUN, and the minimum and maximum sizes of the LUN. After the LUN is created, you cannot modify the LUN host operating system type.

In a MetroCluster configuration, System Manager displays only the following aggregates for creating FlexVol volumes for the LUN:

- In normal mode, when you create volumes on sync-source SVMs or data-serving SVMs in the primary site, only those aggregates that belong to the cluster in the primary site are displayed.
- In switched-over mode, when you create volumes on sync-destination SVMs or data-serving SVMs in the surviving site, only switched-over aggregates are displayed.

### Steps

1. Click **Storage** > **LUNs**.
2. In the **LUN Management** tab, click **Create**.
3. Browse and select an SVM in which you want to create the LUNs.
4. In the **Create LUN Wizard**, specify the name, size, type, description for the LUN, and select the **Space Reserve**, and then click **Next**.
5. Create a new FlexVol volume for the LUN or select an existing volume or qtree, and then click **Next**.
6. Add initiator groups if you want to control host access to the LUN, and then click **Next**.
7. Select the **Manage Storage Quality of Service** check box if you want to manage the workload performance of the LUN.
8. Create a new storage QoS policy group or select an existing policy group to control the input/output (I/O) performance of the LUN:

| If you want to... | Do this... |
|---|---|
| Create a new policy group | **a.** Select **New Policy Group**<br>**b.** Specify the policy group name.<br>**c.** Specify the minimum throughput limit.<br>• You can set the minimum throughput limit only on an AFF platform.<br>• You cannot set the minimum throughput limit for volumes on a FabricPool-enabled aggregate.<br>• If you do not specify the minimum throughput value or if the minimum throughput value is set to 0, the system automatically displays "None" as the value.<br>This value is case-sensitive.<br>**d.** Specify the maximum throughput limit to ensure that the workload of the objects in the policy group does not exceed the specified throughput limit.<br>• The minimum throughput limit and the maximum throughput limit must be of the same unit type.<br>• If you do not specify the minimum throughput limit, then you can set the maximum throughput limit in IOPS and B/s, KB/s, MB/s, and so on.<br>• If you do not specify the maximum throughput value, the system automatically displays "Unlimited" as the value and this value is case-sensitive.<br>The unit that you specify does not affect the maximum throughput. |

| If you want to... | Do this... |
|---|---|
| Select an existing policy group | **a.** Select **Existing Policy Group**, and then click **Choose** to select an existing policy group from the Select Policy Group dialog box.<br>**b.** Specify the minimum throughput limit.<br>• You can set the minimum throughput limit only on an AFF platform.<br>• You cannot set the minimum throughput limit for volumes on a FabricPool-enabled aggregate.<br>• If you do not specify the minimum throughput value or if the minimum throughput value is set to 0, the system automatically displays "None" as the value.<br>This value is case-sensitive.<br>**c.** Specify the maximum throughput limit to ensure that the workload of the objects in the policy group does not exceed the specified throughput limit.<br>• The minimum throughput limit and the maximum throughput limit must be of the same unit type.<br>• If you do not specify the minimum throughput limit, then you can set the maximum throughput limit in IOPS and B/s, KB/s, MB/s, and so on.<br>• If you do not specify the maximum throughput value, the system automatically displays "Unlimited" as the value and this value is case-sensitive.<br>The unit that you specify does not affect the maximum throughput.<br>If the policy group is assigned to more than one object, the maximum throughput that you specify is shared among the objects. |

9. Review the specified details in the **LUN summary** window, and then click **Next**.
10. Confirm the details, and then click **Finish** to complete the wizard.

### Related concepts

*Guidelines for using LUN multiprotocol type* on page 231
The LUN multiprotocol type, or operating system type, specifies the operating system of the host accessing the LUN. It also determines the layout of data on the LUN, and the minimum and maximum size of the LUN.

### Related reference

*LUNs window* on page 233
You can use the LUNs window to create and manage LUNs and to display information about LUNs. You can also add, edit, or delete initiator groups and initiator IDs.

## Deleting LUNs

You can use System Manager to delete LUNs and return the space used by the LUNs to their containing aggregates or volumes.

### Before you begin

- The LUN must be offline.
- The LUN must be unmapped from all initiator hosts.

### Steps

1. Click **Storage** > **LUNs**.
2. In the **LUN Management** tab, select one or more LUNs that you want to delete, and then click **Delete**.

3. Select the confirmation check box, and then click **Delete**.

### Related reference

*LUNs window* on page 233
You can use the LUNs window to create and manage LUNs and to display information about LUNs. You can also add, edit, or delete initiator groups and initiator IDs.

## Creating initiator groups

You can use System Manager to create an initiator group. Initiator groups enable you to control host access to specific LUNs. You can use port sets to limit which LIFs an initiator can access.

### Steps

1. Click **Storage** > **LUNs**.
2. In the **Initiator Groups** tab, click **Create**.
3. In the **General** tab of the **Create Initiator Group** dialog box, specify the initiator group name, operating system, port set, and supported protocol for the group.
4. Click **Create**.

### Related reference

*LUNs window* on page 233
You can use the LUNs window to create and manage LUNs and to display information about LUNs. You can also add, edit, or delete initiator groups and initiator IDs.

## Deleting initiator groups

You can use the Initiator Groups tab in System Manager to delete initiator groups.

### Before you begin

All the LUNs mapped to the initiator group must be manually unmapped.

### Steps

1. Click **Storage** > **LUNs**.
2. In the **Initiator Groups** tab, select one or more initiator groups that you want to delete, and then click **Delete**.
3. Click **Delete**.
4. Verify that the initiator groups you deleted are no longer displayed in the **Initiator Groups** tab.

### Related reference

*LUNs window* on page 233
You can use the LUNs window to create and manage LUNs and to display information about LUNs. You can also add, edit, or delete initiator groups and initiator IDs.

## Adding initiators

You can use System Manager to add initiators to an initiator group. An initiator provides access to a LUN when the initiator group that it belongs to is mapped to that LUN.

### Steps

1. Click **Storage** > **LUNs**.
2. In the **LUN Management** tab, select the initiator group to which you want to add initiators and click **Edit**.
3. In the **Edit Initiator Group** dialog box, click **Initiators**.
4. Click **Add**.

5. Specify the initiator name and click **OK**.
6. Click **Save and Close**.

### Related reference

*LUNs window* on page 233
You can use the LUNs window to create and manage LUNs and to display information about LUNs. You can also add, edit, or delete initiator groups and initiator IDs.

## Deleting initiators from an initiator group

You can use the Initiator Groups tab in System Manager to delete an initiator. To delete an initiator from an initiator group, you must disassociate the initiator from the initiator group.

### Before you begin

All of the LUNs that are mapped to the initiator group that contains the initiator that you want to delete must be manually unmapped.

### Steps

1. Click **Storage** > **LUNs**.
2. In the **Initiator Groups** tab, select the initiator group from which you want to delete the initiator, and then click **Edit**.
3. In the **Edit Initiator Group** dialog box, click the **Initiators** tab.
4. Select and delete the initiator from the text box, and click **Save**.
   The initiator is disassociated from the initiator group.

### Related reference

*LUNs window* on page 233
You can use the LUNs window to create and manage LUNs and to display information about LUNs. You can also add, edit, or delete initiator groups and initiator IDs.

## Creating port sets

You can use System Manager to create port sets to limit access to your LUNs.

### Steps

1. Click **Storage** > **LUNs**.
2. In the **Portsets** tab, click **Create**.
3. In the **Create Portset** dialog box, select the type of protocol.
4. Choose the network interface that you want to associate with the port set.
5. Click **Create**.

## Deleting port sets

You can use System Manager to delete a port set when it is no longer required.

### Steps

1. Click **Storage** > **LUNs**.
2. In the **Portsets** tab, select one or more port sets and click **Delete**.
3. Confirm the deletion by clicking **Delete**.

## Cloning LUNs

LUN clones enable you to create multiple readable and writable copies of a LUN. You can use System Manager to create a temporary copy of a LUN for testing or to make a copy of your data available to additional users without providing them access to the production data.

### Before you begin

- You must have installed the FlexClone license on the storage system.
- When space reservation is disabled on a LUN, the volume that contains the LUN must have enough space to accommodate changes to the clone.

### About this task

- When you create a LUN clone, automatic deletion of the LUN clone is enabled by default in System Manager.

  The LUN clone is deleted when ONTAP triggers automatic deletion to conserve space.
- You cannot clone LUNs that are on SnapLock volumes.

### Steps

1. Click **Storage** > **LUNs**.
2. In the **LUN Management** tab, select the LUN that you want to clone, and then click **Clone**.
3. Optional: If you want to change the default name, specify a new name for the LUN clone.
4. Click **Clone**.
5. Verify that the LUN clone that you created is listed in the **LUNs** window.

### Related reference

*LUNs window* on page 233

You can use the LUNs window to create and manage LUNs and to display information about LUNs. You can also add, edit, or delete initiator groups and initiator IDs.

## Editing LUNs

You can use the LUN properties dialog box in System Manager to change the name, description, size, space reservation setting, or the mapped initiator hosts of a LUN.

### About this task

When you resize a LUN, you have to perform the steps on the host side that are recommended for the host type and the application that is using the LUN.

### Steps

1. Click **Storage** > **LUNs**.
2. In the **LUN Management** tab, select the LUN that you want to edit from the list of LUNs, and click **Edit**.
3. Make the required changes.
4. Click **Save and Close**.

### Related reference

*LUNs window* on page 233

You can use the LUNs window to create and manage LUNs and to display information about LUNs. You can also add, edit, or delete initiator groups and initiator IDs.

## Bringing LUNs online

You can use the **LUN Management** tab in System Manager to bring selected LUNs online and make them available to the host.

### Before you begin

Any host application accessing the LUN must be quiesced or synchronized.

### Steps

1. Click **Storage** > **LUNs**.
2. In the **LUN Management** tab, select one or more LUNs that you want to bring online.
3. Click **Status** > **Online**.

### Related reference

*LUNs window* on page 233
You can use the LUNs window to create and manage LUNs and to display information about LUNs. You can also add, edit, or delete initiator groups and initiator IDs.

## Taking LUNs offline

You can use the **LUN Management** tab in System Manager to take selected LUNs offline and make them unavailable for block protocol access.

### Before you begin

Any host application accessing the LUN must be quiesced or synchronized.

### Steps

1. Click **Storage** > **LUNs**.
2. In the **LUN Management** tab, select one or more LUNs that you want to take offline.
3. Click **Status** > **Offline**.

### Related reference

*LUNs window* on page 233
You can use the LUNs window to create and manage LUNs and to display information about LUNs. You can also add, edit, or delete initiator groups and initiator IDs.

## Moving LUNs

You can use System Manager to move a LUN from its containing volume to another volume or qtree within a storage virtual machine (SVM). You can move the LUN to a volume that is hosted on an aggregate containing high-performance disks, thereby improving the performance when accessing the LUN.

### About this task

- You cannot move a LUN to a qtree within the same volume.
- If you have created a LUN from a file using the command-line interface (CLI), you cannot move the LUN using System Manager.
- The LUN move operation is nondisruptive; it can be performed when the LUN is online and serving data.
- You cannot use System Manager to move the LUN if the allocated space in the destination volume is not sufficient to contain the LUN, and even if autogrow is enabled on the volume. You should use the CLI instead.

- You cannot move LUNs on SnapLock volumes.

### Steps

1. Click **Storage** > **LUNs**.
2. In the **LUN Management** tab, select the LUN that you want to move from the list of LUNs, and then click **Move**.
3. Optional: In the **Move Options** area of the **Move LUN** dialog box, specify a new name for the LUN if you want to change the default name.
4. Select the storage object to which you want to move the LUN and perform one of the following actions:

| If you want to move the LUN to... | Then... |
|---|---|
| A new volume | **a.** Select an aggregate in which you want to create the new volume.<br>**b.** Specify a name for the volume. |
| An existing volume or qtree | **a.** Select a volume to which you want to move the LUN.<br>**b.** If the selected volume contains any qtrees, select the qtree to which you want to move the LUN. |

5. Click **Move**.
6. Confirm the LUN move operation, and click **Continue**.

   For a brief period of time, the LUN is displayed on both the origin and destination volume. After the move operation is complete, the LUN is displayed on the destination volume.

   The destination volume or qtree is displayed as the new container path for the LUN.

## Assigning LUNs to storage QoS

You can use System Manager to limit the throughput of LUNs by assigning them to storage Quality of Service (QoS) policy groups. You can assign storage QoS for new LUNs or modify storage QoS details for LUNs that are already assigned to a policy group.

### About this task

- You cannot assign storage QoS to a LUN if the following storage objects are assigned to a policy group:
  - Parent volume of the LUN
  - Parent storage virtual machine (SVM) of the LUN
- You can assign storage QoS or modify the QoS details for a maximum of 10 LUNs simultaneously.

### Steps

1. Click **Storage** > **LUNs**.
2. In the **LUN Management** tab, select one or more LUNs for which you want to assign storage QoS.
3. Click **Storage QoS**.
4. In the **Quality of Service Details** dialog box, select the **Manage Storage Quality of Service** check box if you want to manage the workload performance of the LUN.

   If some of the LUNs that you selected are already assigned to a policy group, the changes that you make might affect the performance of these LUNs.

5. Create a new storage QoS policy group or select an existing policy group to control the input/output (I/O) performance of the LUN:

| If you want to... | Do this... |
|---|---|
| Create a new policy group | **a.** Select **New Policy Group**.<br>**b.** Specify the policy group name.<br>**c.** Specify the minimum throughput limit.<br>• You can set the minimum throughput limit only on an AFF platform.<br>• You cannot set the minimum throughput limit for volumes on a FabricPool-enabled aggregate.<br>• If you do not specify the minimum throughput value or if the minimum throughput value is set to 0, the system automatically displays "None" as the value.<br>This value is case-sensitive.<br>**d.** Specify the maximum throughput limit to ensure that the workload of the objects in the policy group does not exceed the specified throughput limit.<br>• The minimum throughput limit and the maximum throughput limit must be of the same unit type.<br>• If you do not specify the minimum throughput limit, then you can set the maximum throughput limit in IOPS and B/s, KB/s, MB/s, and so on.<br>• If you do not specify the maximum throughput value, the system automatically displays "Unlimited" as the value and this value is case-sensitive.<br>The unit that you specify does not affect the maximum throughput. |
| Select an existing policy group | **a.** Select **Existing Policy Group**, and then click **Choose** to select an existing policy group from the Select Policy Group dialog box.<br>**b.** Specify the minimum throughput limit.<br>• You can set the minimum throughput limit only on an AFF platform.<br>• You cannot set the minimum throughput limit for volumes on a FabricPool-enabled aggregate.<br>• If you do not specify the minimum throughput value or if the minimum throughput value is set to 0, the system automatically displays "None" as the value.<br>This value is case-sensitive.<br>**c.** Specify the maximum throughput limit to ensure that the workload of the objects in the policy group does not exceed the specified throughput limit.<br>• The minimum throughput limit and the maximum throughput limit must be of the same unit type.<br>• If you do not specify the minimum throughput limit, then you can set the maximum throughput limit in IOPS and B/s, KB/s, MB/s, and so on.<br>• If you do not specify the maximum throughput value, the system automatically displays "Unlimited" as the value and this value is case-sensitive.<br>The unit that you specify does not affect the maximum throughput.<br>If the policy group is assigned to more than one object, the maximum throughput that you specify is shared among the objects. |

6. Optional: Click the link that specifies the number of LUNs to review the list of selected LUNs, and click **Discard** if you want to remove any LUNs from the list.

   The link is displayed only when multiple LUNs are selected.
7. Click **OK**.

## Editing initiator groups

You can use the Edit Initiator Group dialog box in System Manager to change the name of an existing initiator group and its operating system. You can add initiators to or remove initiators from the initiator group. You can also change the port set associated with the initiator group.

### Steps

1. Click **Storage** > **LUNs**.
2. In the **Initiator Groups** tab, select the initiator group that you want to modify, and then click **Edit**.
3. Make the necessary changes.
4. Click **Save and Close**.
5. Verify the changes you made to the initiator group in the **Initiator Groups** tab.

### Related reference

*LUNs window* on page 233

You can use the LUNs window to create and manage LUNs and to display information about LUNs. You can also add, edit, or delete initiator groups and initiator IDs.

## Editing initiators

You can use the Edit Initiator Group dialog box in System Manager to change the name of an existing initiator in an initiator group.

### Steps

1. Click **Storage** > **LUNs**.
2. In the **Initiator Groups** tab, select the initiator group to which the initiator belongs, and then click **Edit**.
3. In the **Edit Initiator Group** dialog box, click **Initiators**.
4. Select the initiator that you want to edit and click **Edit**.
5. Change the name and click **OK**.
6. Click **Save and Close**.

### Related reference

*LUNs window* on page 233

You can use the LUNs window to create and manage LUNs and to display information about LUNs. You can also add, edit, or delete initiator groups and initiator IDs.

## Editing port sets

You can use the Portsets tab in System Manager to edit settings related to port sets.

### Steps

1. Click **Storage** > **LUNs**.
2. In the **Portsets** tab, select the port set you want to edit and click **Edit**.
3. In the **Edit Portset** dialog box, make the necessary changes.
4. Click **Save and Close**.

### Related tasks

*Configuring iSCSI protocol on SVMs* on page 46

You can configure the iSCSI protocol on a storage virtual machine (SVM) to provide block-level data access by using System Manager. You can create iSCSI LIFs and portsets and then add the LIFs to the portsets. LIFs are created on the most suitable adapters and are assigned to portsets to ensure data path redundancy.

## Viewing LUN information

You can use the LUN Management tab in System Manager to view details about a LUN, such as its name, status, size, and type.

### Steps

1. Click **Storage** > **LUNs**.
2. In the **LUN Management** tab, select the LUN that you want to view information about from the displayed list of LUNs.
3. Review the LUN details in the **LUNs** window.

## Viewing initiator groups

You can use the Initiator Groups tab in System Manager to view all the initiator groups and the initiators mapped to these initiator groups, and the LUNs and LUN ID mapped to the initiator groups.

### Steps

1. Click **Storage** > **LUNs**.
2. Click **Initiator Groups** and review the initiator groups that are listed in the upper pane.
3. Select an initiator group to view the initiators that belong to it, which are listed in the **Initiators** tab in the lower pane.
4. Select an initiator group to view the LUNs mapped to it, which are listed in the **Mapped LUNs** in the lower pane.

## Guidelines for working with FlexVol volumes that contain LUNs

When you work with FlexVol volumes that contain LUNs, you must change the default settings for Snapshot copies. You can also optimize the LUN layout to simplify administration.

Snapshot copies are required for many optional features such as SnapMirror, SyncMirror, dump and restore, and ndmpcopy.

When you create a volume, ONTAP automatically performs the following:

- Reserves 5 percent of the space for Snapshot copies
- Schedules Snapshot copies

Because the internal scheduling mechanism for creating Snapshot copies within ONTAP does not ensure that the data within a LUN is in a consistent state, you should change these Snapshot copy settings by performing the following tasks:

- Turn off the automatic Snapshot copy schedule.
- Delete all of the existing Snapshot copies.
- Set the percentage of space reserved for Snapshot copies to zero.

You should use the following guidelines to create volumes that contain LUNs:

- Do not create any LUNs in the system's root volume.

  ONTAP uses this volume to administer the storage system. The default root volume is /vol/vol0.

- You should use a SAN volume to contain the LUN.
- You should ensure that no other files or directories exist in the volume that contains the LUN.

  If this is not possible and you are storing LUNs and files in the same volume, you should use a separate qtree to contain the LUNs.
- If multiple hosts share the same volume, you should create a qtree on the volume to store all of the LUNs for the same host.

  This is a best practice that simplifies LUN administration and tracking.
- To simplify management, you should use naming conventions for LUNs and volumes that reflect their ownership or the way that they are used.

**Related information**

*NetApp Documentation: ONTAP 9*

## Understanding space reservations for LUNs

Understanding how the space reservation setting (combined with the volume guarantee) affects how space is set aside for LUNs helps you to understand the ramifications of disabling space reservations, and why certain combinations of LUN and volume settings are not useful.

When a LUN has space reservations enabled (a space-reserved LUN), and its containing volume has a volume guarantee, free space from the volume is set aside for the LUN at creation time; the size of this reserved space is governed by the size of the LUN. Other storage objects in the volume (other LUNs, files, Snapshot copies, and so on) are prevented from using this space.

When a LUN has space reservations disabled (a non-space-reserved LUN), no space is set aside for that LUN at creation time. The storage required by any write operation to the LUN is allocated from the volume when it is needed, provided sufficient free space is available.

If a space-reserved LUN is created in a none-guaranteed volume, the LUN behaves the same as a non-space-reserved LUN. This is because a none-guaranteed volume has no space to allocate to the LUN; the volume itself can only allocate space as it is written to, due to its none guarantee. Therefore, creating a space-reserved LUN in a none-guaranteed volume is not recommended; employing this configuration combination might provide write guarantees that are in fact impossible.

When the space reserve is set to "Default", the ONTAP space reservation settings apply to the LUNs. ONTAP space reservation settings also apply to the container volumes if new volumes are created.

## Guidelines for using LUN multiprotocol type

The LUN multiprotocol type, or operating system type, specifies the operating system of the host accessing the LUN. It also determines the layout of data on the LUN, and the minimum and maximum size of the LUN.

**Note:** Not all ONTAP versions support all LUN multiprotocol types. For the latest information, see the Interoperability Matrix Tool.

The following table describes the LUN multiprotocol type values and the guidelines for using each type:

| LUN multiprotocol type | When to use |
|---|---|
| AIX | If your host operating system is AIX. |
| HP-UX | If your host operating system is HP-UX. |

| LUN multiprotocol type | When to use |
|---|---|
| Hyper-V | If you are using Windows Server 2008 or Windows Server 2012 Hyper-V and your LUNs contain virtual hard disks (VHDs). If you are using hyper_v for your LUN type, you should also use hyper_v for your igroup OS type.<br>**Note:** For raw LUNs, you can use the type of child operating system that the LUN multiprotocol type uses. |
| Linux | If your host operating system is Linux. |
| NetWare | If your host operating system is NetWare. |
| OpenVMS | If your host operating system is OpenVMS. |
| Solaris | If your host operating system is Solaris and you are not using Solaris EFI labels. |
| Solaris EFI | If you are using Solaris EFI labels.<br>**Note:** Using any other LUN multiprotocol type with Solaris EFI labels might result in LUN misalignment problems. |
| VMware | If you are using an ESX Server and your LUNs will be configured with VMFS.<br>**Note:** If you configure the LUNs with RDM, you can use the guest operating system as the LUN multiprotocol type. |
| Windows 2003 MBR | If your host operating system is Windows Server 2003 using the MBR partitioning method. |
| Windows 2003 GPT | If you want to use the GPT partitioning method and your host is capable of using it. Windows Server 2003, Service Pack 1 and later are capable of using the GPT partitioning method, and all 64-bit versions of Windows support it. |
| Windows 2008 or later | If your host operating system is Windows Server 2008 or later; both MBR and GPT partitioning methods are supported. |
| Xen | If you are using Xen and your LUNs will be configured with Linux LVM with Dom0.<br>**Note:** For raw LUNs, you can use the type of guest operating system that the LUN multiprotocol type uses. |

### Related tasks

*Creating LUNs* on page 220

You can use System Manager to create LUNs for an existing aggregate, volume, or qtree when there is available free space. You can create a LUN in an existing volume or create a new FlexVol volume for the LUN. You can also enable storage Quality of Service (QoS) to manage the workload performance.

### Related information

*NetApp Interoperability*

*Solaris Host Utilities 6.1 Installation and Setup Guide*

*Solaris Host Utilities 6.1 Quick Command Reference*

*Solaris Host Utilities 6.1 Release Notes*

## Understanding LUN clones

LUN clones are writable, space-efficient clones of parent LUNs. Creating LUN clones is highly space-efficient and time-efficient because the cloning operation does not involve physically copying any data. Clones help in space storage utilization of the physical aggregate space.

You can clone a complete LUN without the need of a backing Snapshot copy in a SAN environment. The cloning operation is instantaneous and clients that are accessing the parent LUN do not experience any disruption or outage. Clients can perform all normal LUN operations on both parent entities and clone entities. Clients have immediate read/write access to both the parent and cloned LUN.

Clones share the data blocks of their parent LUNs and occupy negligible storage space until clients write new data either to the parent LUN, or to the clone. By default, the LUN clone inherits the space reserved attribute of the parent LUN. For example, if space reservation is disabled on the parent LUN, then space reservation is also disabled on the LUN clone.

**Note:** When you clone a LUN, you must ensure that the volume has enough space to contain the LUN clone.

## Initiator hosts

Initiator hosts can access the LUNs mapped to them. When you map a LUN on a storage system to the igroup, you grant all the initiators in that group access to that LUN. If a host is not a member of an igroup that is mapped to a LUN, that host does not have access to the LUN.

## igroup name

The igroup name is a case-sensitive name that must satisfy several requirements.

The igroup name:

- Contains 1 to 96 characters. Spaces are not allowed.
- Can contain the letters A through Z, a through z, numbers 0 through 9, hyphen (“-”), underscore (“_”), colon (“:”), and period (“.”).
- Must start with a letter or number.

The name you assign to an igroup is independent of the name of the host that is used by the host operating system, host files, or Domain Name Service (DNS). If you name an igroup aix1, for example, it is not mapped to the actual IP host name (DNS name) of the host.

**Note:** You might find it useful to provide meaningful names for igroups, ones that describe the hosts that can access the LUNs mapped to them.

## igroup type

The igroup type can be mixed type, iSCSI, or FC/FCoE.

## igroup ostype

The ostype indicates the type of host operating system used by all of the initiators in the igroup. All initiators in an igroup must be of the same ostype. The ostypes of initiators are **`solaris`**, **`windows`**, **`hpux`**, **`aix`**, **`netware`**, **`xen`**, **`hyper_v`**, **`vmware`**, and **`linux.`**

You must select an ostype for the igroup.

## LUNs window

You can use the LUNs window to create and manage LUNs and to display information about LUNs. You can also add, edit, or delete initiator groups and initiator IDs.

- *LUN Management tab*

- *Initiator Groups tab*
- *Portsets tab*

### LUN Management tab

This tab enables you to create, clone, delete, move, or edit the settings of LUNs. You can also assign LUNs to a Storage Quality of Service (QoS) policy group.

### Command buttons

**Create**

Opens the Create LUN wizard, which enables you to create LUNs.

In a cluster on an AFF platform that does not contain any existing LUNs, the Create FC SAN optimized LUNs dialog box is opened, which enables you to set up one or more FC SAN optimized LUNs.

**Clone**

Opens the Clone LUN dialog box, which enables you to clone the selected LUNs.

**Edit**

Opens the Edit LUN dialog box, which enables you to edit the settings of the selected LUN.

**Delete**

Deletes the selected LUN.

**Status**

Enables you to change the status of the selected LUN to either Online or Offline.

**Move**

Opens the Move LUN dialog box, which enables you to move the selected LUN to a new volume or an existing volume or qtree within the same storage virtual machine (SVM).

**Storage QoS**

Opens the Quality of Service details dialog box, which enables you to assign one or more LUNs to a new or existing policy group.

**Refresh**

Updates the information in the window.

### LUNs list

**Name**

Displays the name of the LUN.

**SVM**

Displays the name of the storage virtual machine (SVM) in which the LUN is created.

**Container Path**

Displays the name of the file system (volume or qtree) that contains the LUN.

**Space Reservation**

Specifies whether space reservation is enabled or disabled.

**Available Size**

Displays the space available in the LUN.

**Total Size**

Displays the total space in the LUN.

**%Used**

Displays the total space (in percentage) that is used.

**Type**

Specifies the LUN type.

**Status**

Specifies the status of the LUN.

**Policy Group**

Displays the name of the Storage QoS policy group to which the LUN is assigned. By default, this column is hidden.

**Application**

Displays the name of the application that is assigned to the LUN.

### Details area

The area below the LUNs list displays details related to the selected LUN.

**Details tab**

Displays details related to the LUN such as the LUN serial number, whether the LUN is a clone, LUN description, the policy group to which the LUN is assigned, minimum throughput of the policy group, maximum throughput of the policy group, details about the LUN move operation, and the application assigned to the LUN. You can also view details about the initiator groups and initiators that are associated with the selected LUN.

**Performance tab**

Displays performance metrics graphs of the LUNs, including data rate, IOPS, and response time.

Changing the client time zone or the cluster time zone impacts the performance metrics graphs. Refresh your browser to see the updated graphs.

### Initiator Groups tab

This tab enables you to create, delete, or edit the settings of initiator groups and initiator IDs.

### Command buttons

**Create**

Opens the Create Initiator Group dialog box, which enables you to create initiator groups to control host access to specific LUNs.

**Edit**

Opens the Edit Initiator Group dialog box, which enables you to edit the settings of the selected initiator group.

**Delete**

Deletes the selected initiator group.

**Refresh**

Updates the information in the window.

### Initiator Groups list

**Name**

Displays the name of the initiator group.

**Type**

Specifies the type of protocol supported by the initiator group. The supported protocols are iSCSI, FC/FCoE, or Mixed (iSCSI and FC/FCoE).

**Operating System**

Specifies the operating system for the initiator group.

**Portset**

Displays the port set that is associated with the initiator group.

**Initiator Count**

Displays the number of initiators added to the initiator group.

### Details area

The area below the Initiator Groups list displays details about the initiators that are added to the selected initiator group and the LUNs that are mapped to the initiator group.

### Portsets tab

This tab enables you to create, delete, or edit the settings of port sets.

### Command buttons

**Create**

Opens the Create Portset dialog box, which enables you to create port sets to limit access to your LUNs.

**Edit**

Opens the Edit Portset dialog box, which enables you to select the network interfaces that you want to associate with the port set.

**Delete**

Deletes the selected port set.

**Refresh**

Updates the information in the window.

### Portsets list

**Portset Name**

Displays the name of the port set.

**Type**

Specifies the type of protocol supported by the port set. The supported protocols are iSCSI, FC/FCoE, or Mixed (iSCSI and FC/FCoE).

**Interface Count**

Displays the number of network interfaces that are associated with the port set.

**Initiator Group Count**

Displays the number of initiator groups that are associated with the port set.

### Details area

The area below the Portsets list displays details about the network interfaces and initiator groups associated with the selected port set.

### Related tasks

*Creating LUNs* on page 220

You can use System Manager to create LUNs for an existing aggregate, volume, or qtree when there is available free space. You can create a LUN in an existing volume or create a new FlexVol

volume for the LUN. You can also enable storage Quality of Service (QoS) to manage the workload performance.

*Deleting LUNs* on page 222
You can use System Manager to delete LUNs and return the space used by the LUNs to their containing aggregates or volumes.

*Creating initiator groups* on page 223
You can use System Manager to create an initiator group. Initiator groups enable you to control host access to specific LUNs. You can use port sets to limit which LIFs an initiator can access.

*Deleting initiator groups* on page 223
You can use the Initiator Groups tab in System Manager to delete initiator groups.

*Adding initiators* on page 223
You can use System Manager to add initiators to an initiator group. An initiator provides access to a LUN when the initiator group that it belongs to is mapped to that LUN.

*Deleting initiators from an initiator group* on page 224
You can use the Initiator Groups tab in System Manager to delete an initiator. To delete an initiator from an initiator group, you must disassociate the initiator from the initiator group.

*Editing LUNs* on page 225
You can use the LUN properties dialog box in System Manager to change the name, description, size, space reservation setting, or the mapped initiator hosts of a LUN.

*Editing initiator groups* on page 229
You can use the Edit Initiator Group dialog box in System Manager to change the name of an existing initiator group and its operating system. You can add initiators to or remove initiators from the initiator group. You can also change the port set associated with the initiator group.

*Editing initiators* on page 229
You can use the Edit Initiator Group dialog box in System Manager to change the name of an existing initiator in an initiator group.

*Bringing LUNs online* on page 226
You can use the **LUN Management** tab in System Manager to bring selected LUNs online and make them available to the host.

*Taking LUNs offline* on page 226
You can use the **LUN Management** tab in System Manager to take selected LUNs offline and make them unavailable for block protocol access.

*Cloning LUNs* on page 225
LUN clones enable you to create multiple readable and writable copies of a LUN. You can use System Manager to create a temporary copy of a LUN for testing or to make a copy of your data available to additional users without providing them access to the production data.

## Qtrees

You can use System Manager create, edit, and delete Qtrees.

### Related information

*ONTAP concepts*

*Logical storage management*

*NFS management*

*SMB/CIFS management*

## Creating qtrees

Qtrees enable you to manage and partition your data within a volume. You can use the Create Qtree dialog box in System Manager to add a new qtree to a volume on your storage system.

### Steps

1. Click **Storage** > **Qtrees**.
2. From the **Qtrees on SVM** list, select the storage virtual machine (SVM) on which you want to create a qtree.
3. Click **Create**.
4. In the **Details** tab of the **Create Qtree** dialog box, type a name for the qtree.
5. Select the volume to which you want to add the qtree.

   The Volume browse list includes only the volumes that are online.
6. If you want to disable opportunistic locks (oplocks) for the qtree, clear the **Enable Oplocks for files and directories in this Qtree** check box.

   By default, oplocks are enabled for each qtree.
7. If you want to change the default inherited security style, select a new security style.

   The default security style of the qtree is the security style of the volume that contains the qtree.
8. If you want to change the default inherited export policy, either select an existing export policy or create an export policy.

   The default export policy of the qtree is the export policy that is assigned to the volume that contains the qtree.
9. If you want to restrict the disk space usage, click the **Quotas** tab.
   a. If you want to apply quotas on the qtree, click **Qtree quota**, and then specify the disk space limit.
   b. If you want to apply quotas for all the users on the qtree, click **User quota**, and then specify the disk space limit.
10. Click **Create**.
11. Verify that the qtree that you created is included in the list of qtrees in the **Qtrees** window.

### Related reference

*Qtrees window* on page 240

You can use the Qtrees window to create, display, and manage information about qtrees.

## Deleting qtrees

You can delete a qtree and reclaim the disk space that the qtree uses within a volume by using System Manager. When you delete a qtree, all of the quotas that are applicable to that qtree are no longer applied by ONTAP.

### Before you begin

- The qtree status must be normal.
- The qtree must not contain any LUN.

### Steps

1. Click **Storage** > **Qtrees**.
2. In the **Qtrees** window, select one or more qtrees that you want to delete, and then click **Delete**.
3. Select the confirmation check box, and then click **Delete**.
4. Verify that the qtree that you deleted is no longer included in the list of qtrees in the **Qtrees** window.

### Related reference

*Qtrees window* on page 240

You can use the Qtrees window to create, display, and manage information about qtrees.

## Editing qtrees

You can use System Manager to modify the properties of a qtree such as the security style, enable or disable opportunistic locks (oplocks), and assign a new or existing export policy.

### Steps

1. Click **Storage** > **Qtrees**.
2. Select the qtree that you want to edit, and then click **Edit**.
3. In the **Edit Qtree** dialog box, edit the following properties as required:
   - Oplocks
   - Security style
   - Export policy
4. Click **Save**.
5. Verify the changes that you made to the selected qtree in the **Qtrees** window.

### Related reference

*Qtrees window* on page 240

You can use the Qtrees window to create, display, and manage information about qtrees.

## Assigning export policies to qtrees

Instead of exporting an entire volume, you can export a specific qtree on a volume to make it directly accessible to clients. You can use System Manager to export a qtree by assigning an export policy to the qtree. You can assign an export policy to one or more qtrees from the Qtrees window.

### Steps

1. Click **Storage** > **Qtrees**.
2. From the **Qtrees on SVM** list, select the storage virtual machine (SVM) on which the qtrees that you want to export reside.
3. Select one or more qtrees for which you want to assign an export policy, and then click **Change Export Policy**.
4. In the **Export Policy** dialog box, either create an export policy or select an existing export policy.

   *Creating an export policy* on page 269
5. Click **Save**.
6. Verify that the export policy and its related export rules that you assigned to the qtrees are displayed in the **Details** tab of the appropriate qtrees.

## Viewing qtree information

You can use the Qtrees window in System Manager to view the volume that contains the qtree, the name, security style, and status of the qtree, and the oplocks status.

### Steps

1. Click **Storage** > **Qtrees**.
2. From the **Qtrees on SVM** list, select the storage virtual machine (SVM) on which the qtree about which you want to view information resides.

3. Select the qtree from the displayed list of qtrees.
4. Review the qtree details in the **Qtrees** window.

## Qtree options

A qtree is a logically defined file system that can exist as a special subdirectory of the root directory within a FlexVol volume and are used to manage and partition data within the volume.

You can specify the following options when creating a qtree:

- Name of the qtree
- Volume in which you want the qtree to reside
- Oplocks
  By default, oplocks are enabled for the qtree. If you disable oplocks for the entire storage system, oplocks are not set even if you enable oplocks on a per-qtree basis
- Security style
  The security style can be UNIX, NTFS, or Mixed (UNIX and NTFS). By default, the security style of the qtree is the same as that of the selected volume.
- Export policy
  Create a new export policy or select an existing policy. By default, the export policy of the qtree is same as that of the selected volume.
- Space usage limits for qtree and user quotas

## Qtrees window

You can use the Qtrees window to create, display, and manage information about qtrees.

- *Command buttons* on page 240
- *Qtree list* on page 240
- *Details area* on page 241

### Command buttons

**Create**

Opens the Create Qtree dialog box, which enables you to create a new qtree.

**Edit**

Opens the Edit Qtree dialog box, which enables you to change the security style and to enable or disable oplocks (opportunistic locks) on a qtree.

**Change Export Policy**

Opens the Export Policy dialog box, which enables you to assign one or more qtrees to new or existing export policies.

**Delete**

Deletes the selected qtree.

This button is disabled unless the status of the selected qtree is normal.

**Refresh**

Updates the information in the window.

### Qtree list

The qtree list displays the volume in which the qtree resides and the qtree name.

**Name**

Displays the name of the qtree.

**Volume**

Displays the name of the volume in which the qtree resides.

**Security Style**

Specifies the security style of the qtree.

**Status**

Specifies the current status of the qtree.

**Oplocks**

Specifies whether the oplocks setting is enabled or disabled for the qtree.

**Export Policy**

Displays the name of the export policy to which the qtree is assigned.

### Details area

**Details tab**

Displays detailed information about the selected qtree, such as the mount path of the volume containing the qtree, details about the export policy, and the export policy rules.

### Related tasks

*Creating qtrees* on page 238
Qtrees enable you to manage and partition your data within a volume. You can use the Create Qtree dialog box in System Manager to add a new qtree to a volume on your storage system.

*Deleting qtrees* on page 238
You can delete a qtree and reclaim the disk space that the qtree uses within a volume by using System Manager. When you delete a qtree, all of the quotas that are applicable to that qtree are no longer applied by ONTAP.

*Editing qtrees* on page 239
You can use System Manager to modify the properties of a qtree such as the security style, enable or disable opportunistic locks (oplocks), and assign a new or existing export policy.

# Quotas

You can use System Manager to create, edit, and delete quotas.

### Related information

*Logical storage management*

## Creating quotas

Quotas enable you to restrict or track the disk space and number of files that are used by a user, group, or qtree. You can use the Add Quota wizard in System Manager to create a quota and to apply the quota to a specific volume or qtree.

### About this task

Using System Manager, the minimum value that you can specify for the hard limit and soft limit on the number of files that the quota can own is 1000. If you want to specify a value lower than 1000, you should use the command-line interface (CLI).

### Steps

1. Click **Storage** > **Quotas**.
2. From the **Quotas on SVM** list, select the storage virtual machine (SVM) on which you want to create a quota.
3. In the **User Defined Quotas** tab, click **Create**.
   The Create Quota Wizard is displayed.

4. Type or select information as prompted by the wizard.
5. Confirm the details, and then click **Finish** to complete the wizard.

### After you finish

You can use the local user name or RID to create user quotas. If you create the user quota or group quota by using the user name or group name, then the `/etc/passwd` file and the `/etc/group` file must be updated, respectively.

#### Related reference

*Quotas window* on page 245
You can use the Quotas window to create, display, and manage information about quotas.

## Deleting quotas

You can use System Manager to delete one or more quotas as your users and their storage requirements and limitations change.

### Steps

1. Click **Storage** > **Quotas**.
2. From the **Quotas on SVM** list, select the storage virtual machine (SVM) on which the quotas that you want to delete reside.
3. Select one or more quotas that you want to delete, and then click **Delete**.
4. Select the confirmation check box, and then click **Delete**.

#### Related reference

*Quotas window* on page 245
You can use the Quotas window to create, display, and manage information about quotas.

## Editing quota limits

You can use System Manager to edit the disk space threshold, the hard limit and soft limit on the amount of disk space that the quota target can use, and the hard limit and soft limit on the number of files that the quota target can own.

### Steps

1. Click **Storage** > **Quotas**.
2. From the **Quotas on SVM** list, select the storage virtual machine (SVM) on which the quota that you want to edit resides.
3. Select the quota that you want to edit, and click **Edit Limits**.
4. In the **Edit Limits** dialog box, edit the quota settings as required.

   One hundred (100) is the minimum value that you can specify for the hard limit and soft limit on the number of files that the quota can own. If you want to specify a value lower than 100, you should use the command-line interface (CLI).
5. Click **Save and Close**.
6. Verify the changes that you made to the selected quota in the **User Defined Quotas** tab.

#### Related reference

*Quotas window* on page 245

You can use the Quotas window to create, display, and manage information about quotas.

## Activating or deactivating quotas

You can use System Manager to activate or deactivate quotas on one or more selected volumes on your storage system, as your users and their storage requirements and limitations change.

### Steps

1. Click **Storage** > **Quotas**.
2. From the **Quotas on SVM** list, select the storage virtual machine (SVM) on which the quotas that you want to activate or deactivate reside.
3. In the **Quota Status on Volumes** tab, select one or more volumes for which you want to activate or deactivate quotas.
4. Click **Activate** or **Deactivate**, as required.
5. If you are deactivating a quota, select the confirmation check box, and then click **OK**.
6. Verify the quota status on the volumes in the **Status** column.

### Related reference

*Quotas window* on page 245
You can use the Quotas window to create, display, and manage information about quotas.

## Resizing quotas

You can use the Resize Quota dialog box in System Manager to adjust the active quotas in the specified volume so that they reflect the changes that you have made to a quota.

### Before you begin

Quotas must be enabled for the volumes for which you want to resize quotas.

### Steps

1. Click **Storage** > **Quotas**.
2. In the **Quota Status on Volumes** tab of the **Quotas** window, select one or more volumes for which you want to resize the quotas.
3. Click **Resize**.

### Related reference

*Quotas window* on page 245
You can use the Quotas window to create, display, and manage information about quotas.

## Viewing quota information

You can use the Quotas window in System Manager to view quota details such as the volume and qtrees to which the quota is applied, the type of quota, the user or group to which the quota is applied, and the space and file usage.

### Steps

1. Click **Storage** > **Quotas**.
2. From the **Quotas on SVM** list, select the storage virtual machine (SVM) on which the quota that you want to view information about resides.
3. Perform the appropriate action:

| If... | Then... |
|---|---|
| You want to view details of all of the quotas that you created | Click the **User Defined Quotas** tab. |

| If... | Then... |
|---|---|
| You want to view details of the quotas that are currently active | Click the **Quota Report** tab. |

4. Select the quota that you want to view information about from the displayed list of quotas.
5. Review the quota details.

## Types of quotas

Quotas can be classified on the basis of the targets to which they are applied.

The following are the types of quotas based on the targets to which they are applied:

**User quota**

The target is a user.

The user can be represented by a UNIX user name, UNIX UID, a Windows SID, a file or directory whose UID matches the user, Windows user name in pre-Windows 2000 format, and a file or directory with an ACL owned by the user's SID. You can apply it to a volume or a qtree.

**Group quota**

The target is a group.

The group is represented by a UNIX group name, a GID, or a file or directory whose GID matches the group. ONTAP does not apply group quotas based on a Windows ID. You can apply a quota to a volume or a qtree.

**Qtree quota**

The target is a qtree, specified by the path name to the qtree.

You can determine the size of the target qtree.

**Default quota**

Automatically applies a quota limit to a large set of quota targets without creating separate quotas for each target.

Default quotas can be applied to all three types of quota target (users, groups, and qtrees). The quota type is determined by the value of the type field.

## Quota limits

You can apply a disk space limit or limit the number of files for each quota type. If you do not specify a limit for a quota, none is applied.

**Disk space soft limit**

Disk space limit applied to soft quotas.

**Disk space hard limit**

Disk space limit applied to hard quotas.

**Threshold limit**

Disk space limit applied to threshold quotas.

**Files soft limit**

The maximum number of files on a soft quota.

**Files hard limit**

The maximum number of files on a hard quota.

## Quota management

System Manager includes several features that help you to create, edit, or delete quotas. You can create a user, group, or tree quota and you can specify quota limits at the disk and file levels. All quotas are established on a per-volume basis.

After creating a quota, you can perform the following tasks:

- Enable and disable quotas
- Resize quotas

## Quotas window

You can use the Quotas window to create, display, and manage information about quotas.

- *Tabs*
- *Command buttons*
- *User Defined Quotas list*
- *Details area*

### Tabs

**User Defined Quotas**

You can use the **User Defined Quotas** tab to view details of the quotas that you create and to create, edit, or delete quotas.

**Quota Report**

You can use the Quota Report tab to view the space and file usage and to edit the space and file limits of quotas that are active.

**Quota Status on Volumes**

You can use the Quota Status on Volumes tab to view the status of a quota and to turn quotas on or off and to resize quotas.

### Command buttons

**Create**

Opens the Create Quota wizard, which enables you to create quotas.

**Edit Limits**

Opens the Edit Limits dialog box, which enables you to edit settings of the selected quota.

**Delete**

Deletes the selected quota from the quotas list.

**Refresh**

Updates the information in the window.

### User Defined Quotas list

The quotas list displays the name and storage information for each quota.

**Volume**

Specifies the volume to which the quota is applied.

**Qtree**

Specifies the qtree associated with the quota. "All Qtrees" indicates that the quota is associated with all the qtrees.

**Type**

Specifies the quota type: user, or group, or tree.

**User/Group**

Specifies a user or a group associated with the quota. "All Users" indicates that the quota is associated with all the users. "All groups" indicates that the quota is associated with all the groups.

**Quota Target**

Specifies the type of target that the quota is assigned to. The target can be qtree, user, or group.

**Space Hard Limit**

Specifies the disk space limit applied to hard quotas.

This field is hidden by default.

**Space Soft Limit**

Specifies the disk space limit applied to soft quotas.

This field is hidden by default.

**Threshold**

Specifies the disk space limit applied to threshold quotas.

This field is hidden by default.

**File Hard Limit**

Specifies the maximum number of files in a hard quota.

This field is hidden by default.

**File Soft Limit**

Specifies the maximum number of files in a soft quota.

This field is hidden by default.

## Details area

The area below the quotas list displays quota details such as the quota error, space usage and limits, and file usage and limits.

### Related tasks

*Creating quotas* on page 241
Quotas enable you to restrict or track the disk space and number of files that are used by a user, group, or qtree. You can use the Add Quota wizard in System Manager to create a quota and to apply the quota to a specific volume or qtree.

*Deleting quotas* on page 242
You can use System Manager to delete one or more quotas as your users and their storage requirements and limitations change.

*Editing quota limits* on page 242
You can use System Manager to edit the disk space threshold, the hard limit and soft limit on the amount of disk space that the quota target can use, and the hard limit and soft limit on the number of files that the quota target can own.

*Activating or deactivating quotas* on page 243
You can use System Manager to activate or deactivate quotas on one or more selected volumes on your storage system, as your users and their storage requirements and limitations change.

*Resizing quotas* on page 243

You can use the Resize Quota dialog box in System Manager to adjust the active quotas in the specified volume so that they reflect the changes that you have made to a quota.

# CIFS protocol

You can use System Manager to enable and configure CIFS servers to allow CIFS clients to access files on the cluster.

### Related information

*SMB/CIFS management*

## Setting up CIFS

You can use System Manager to enable and configure CIFS servers to allow CIFS clients to access the files on the cluster.

### Before you begin

- The CIFS license must be installed on your storage system.
- While configuring CIFS in the Active Directory domain, the following requirements must be met:
  - DNS must be enabled and configured correctly.
  - The storage system must be able to communicate with the domain controller by using the fully qualified domain name (FQDN).
  - The time difference (clock skew) between the cluster and the domain controller must not be more than five minutes.
- If CIFS is the only protocol that is configured on the storage virtual machine (SVM), the following requirements must be met:
  - The root volume security style must be NTFS.
    By default, System Manager sets the security style as UNIX.
  - Superuser access must be set to **Any** for the CIFS protocol.

### Steps

1. Click **Storage** > **SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **Configuration** tab, click **Set up**.
4. In the **General** tab of the **CIFS Server Setup** dialog box, specify the NetBIOS name and the Active Directory domain details.
5. Click the **Options** tab, and then perform the following actions:
   - In the SMB settings area, select or clear the SMB signing check box and the SMB encryption check box, as required.
   - Specify the default UNIX user.
   - In the WINS Servers area, add the required IP address.
6. Click **Set up**.

### Related tasks

*Creating a CIFS share* on page 211
You can use System Manager to create a CIFS share that enables you to specify the folder, qtree, or volume that CIFS users can access.

*Editing volume properties* on page 174

You can modify volume properties such as the volume name, security style, fractional reserve, and space guarantee by using System Manager. You can modify storage efficiency settings (deduplication schedule, deduplication policy, and compression) and space reclamation settings.-

*Modifying export policy rules* on page 271
You can use System Manager to modify the specified client, access protocols, and access permissions of an export policy rule.

### Related reference

*CIFS window* on page 253
You can use the CIFS window to configure the CIFS server, to manage domain controllers, to manage symbolic UNIX mappings, and to configure BranchCache.

## Editing the general properties for CIFS

You can modify the general properties for CIFS such as the default UNIX user and default Windows user by using System Manager. You can also enable or disable SMB signing for the CIFS server.

### Steps

1. Click **Storage** > **SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **Configuration** tab, click **Options**.
4. In the **CIFS Options** dialog box, modify the following CIFS server properties, as required:
   - UNIX user
   - Windows user
   - IP address
   - Enable or disable SMB signing

     Enabling SMB signing prevents the data from being compromised. However, you might encounter performance degradation in the form of increased CPU usage on both the clients and the server, although the network traffic remains the same. You can disable SMB signing on any of your Windows clients that do not require protection against replay attacks.

     For information about disabling SMB signing on Windows clients, see the Microsoft Windows documentation.
   - Enable or disable SMB 3.0 encryption

     You should enable SMB Multichannel to establish multiple channels between an SMB 3.0 session and transport connections.
5. Click either **Save** or **Save and Close**.

### Related reference

*CIFS window* on page 253
You can use the CIFS window to configure the CIFS server, to manage domain controllers, to manage symbolic UNIX mappings, and to configure BranchCache.

## Adding home directory paths

You can use System Manager to specify one or more paths that can be used by the storage system to resolve the location of the CIFS home directories of users.

### Steps

1. Click **Storage** > **SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **Home Directories** area of the **Configuration** tab, click **Manage**.

4. In the **Manage Home Directories** dialog box, specify the paths that are to be used by the storage system to search for the CIFS home directories of users.
5. Click **Add**, and then click **Save and Close**.

### Related reference

*CIFS window* on page 253
You can use the CIFS window to configure the CIFS server, to manage domain controllers, to manage symbolic UNIX mappings, and to configure BranchCache.

## Deleting home directory paths

You can use System Manager to delete a home directory path when you do not want the storage system to use the path to resolve the location of the CIFS home directories of users.

### Steps

1. Click **Storage** > **SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **Home Directories** area of the **Configuration** tab, click **Manage**.
4. In the **Manage Home Directories** dialog box, select the home directory path that you want to delete, and then click **Delete**.
5. Click **Save and Close**.

### Related reference

*CIFS window* on page 253
You can use the CIFS window to configure the CIFS server, to manage domain controllers, to manage symbolic UNIX mappings, and to configure BranchCache.

## Resetting CIFS domain controllers

You can use System Manager to reset the CIFS connection to domain controllers for the specified domain. Failure to reset the domain controller information can cause a connection failure.

### About this task

You have to update the discovery information of the storage system's available domain controller after you add or delete a domain from the list of preferred domain controllers. You can update the storage system's available domain controller discovery information in ONTAP through the command-line interface (CLI).

### Steps

1. Click **Storage** > **SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **Domain** tab, click **Reset**.

### Related reference

*CIFS window* on page 253

You can use the CIFS window to configure the CIFS server, to manage domain controllers, to manage symbolic UNIX mappings, and to configure BranchCache.

## Updating the CIFS group policy configuration

You have to update the group policy after the policy configuration is changed through the command-line interface (CLI). You can use the CIFS window in System Manager to update the group policy.

### Steps

1. Click **Storage** > **SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. Click the **Domain** tab.
4. In the **Group Policy** area, select the group policy configuration that you want to update, and then click **Update**.

## Enabling or disabling a CIFS group policy configuration

You can enable or disable the CIFS group policy configuration from the CIFS window in System Manager.

### Steps

1. Click **Storage** > **SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. Click the **Domain** tab.
4. In the **Group Policy** area, select the group policy configuration that you want to enable or disable, and then click **Enable** or **Disable**, as required.

## Reloading CIFS group policy

You have to reload a CIFS group policy if the status of the policy is changed. You can use the CIFS window in System Manager to reload the group policy.

### Steps

1. Click **Storage** > **SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. Click the **Domain** tab.
4. In the **Group Policy** area, select the group policy configuration that you want to reload, and then click **Reload**.

## Setting up BranchCache

You can use System Manager to configure BranchCache on a CIFS-enabled storage virtual machine (SVM) to enable the caching of content on computers that are local to the requesting clients.

### Before you begin

- CIFS must be licensed, and a CIFS server must be configured.
- For BranchCache version 1, SMB 2.1 or later must be enabled.
- For BranchCache version 2, SMB 3.0 must be enabled, and the remote Windows clients must support BranchCache 2.

### About this task

- You can configure BranchCache on SVMs.

- You can create an all-shares BranchCache configuration if you want to offer caching services for all of the content that is contained within all of the SMB shares on the CIFS server.
- You can create a per-share BranchCache configuration if you want to offer caching services for the content that is contained within selected SMB shares on the CIFS server.

### Steps

1. Click **Storage** > **SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **BranchCache** tab, click **Set Up**.
4. In the **BranchCache Setup** dialog box, enter the following information:
   a. Specify the path to the hash store.

      The path can be to an existing directory where you want the hash data to be stored. The destination path must be read-writable. Read-only paths such as Snapshot directories are not allowed.
   b. Specify the maximum size (in KB, MB, GB, TB, or PB) for a hash data store.

      If the hash data exceeds this value, older hashes are deleted to provide space for newer hashes. The default size for a hash store is 1 GB.
   c. Specify the operating mode for the BranchCache configuration.

      The default operating mode is set to all shares.
   d. Specify a server key to prevent clients from impersonating the BranchCache server.

      You can set the server key to a specific value so that if multiple servers are providing BranchCache data for the same files, clients can use hashes from any server using that same server key. If the server key contains any spaces, you must enclose the server key in quotation marks.
   e. Select the required BranchCache version.

      By default, all of the versions that are supported by the client are selected.
5. Click **Set Up**.

## Modifying the BranchCache settings

You can use the CIFS window in System Manager to modify the BranchCache settings that are configured for a CIFS-enabled storage virtual machine (SVM). You can change the hash store path, the hash store size, the operating mode, and the BranchCache versions that are supported.

### Steps

1. Click **Storage** > **SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **BranchCache** tab, click **Edit**.
4. In the **Modify BranchCache Settings** dialog box, modify the required information:
   - Hash store path

     If you modify the hash store path, you are provided with an option to retain the cached hash data from the previous hash store.
   - Hash store size
   - Operating mode
   - BranchCache version
5. Click **Modify**.

## Deleting the BranchCache configuration

You can use System Manager to delete the BranchCache configuration if you no longer want to offer caching services on the storage virtual machine (SVM) that is configured for BranchCache.

### Steps

1. Click **Storage** > **SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **BranchCache** tab, click **Delete**.
4. Select the confirmation check box, and then click **Delete**.

   You can also remove existing hashes from the hash store.

## Adding preferred domain controllers

System Manager automatically discovers domain controllers through DNS. Optionally, you can add one or more domain controllers to the list of preferred domain controllers for a specific domain.

### Steps

1. Click **Storage** > **SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **Domain** tab, click **Add** in the **Preferred Domain Controllers** area.
4. Enter the fully qualified domain name (FQDN) and the IP addresses of the domain controllers that you want to add.

   You can add multiple domain controllers by entering the IP addresses of the domain controllers, separated by commas.
5. Click **Save**.
6. Verify that the domain controller that you added is displayed in the list of preferred domain controllers.

## Editing preferred domain controllers

You can use System Manager to modify the IP address of the preferred domain controllers that are configured for a specific domain.

### Steps

1. Click **Storage** > **SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **Preferred Domain Controllers** area of the **Domain** tab, double-click the domain controller that you want to edit.
4. Modify the IP addresses of the domain controller, and then click **Save**.

## Deleting preferred domain controllers

You can use System Manager to delete a preferred domain controller to which the storage virtual machine (SVM) computer account is associated. You can do this when you no longer want to use a particular domain controller.

### Steps

1. Click **Storage** > **SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **Domain** tab, select the domain that you want to delete from the **Preferred Domain Controllers** area, and then click **Delete**.

4. Select the confirmation check box, and then click **Delete**.

## Viewing CIFS domain information

You can use System Manager to view information about the domain controllers and servers that are connected to the storage system.

### Steps

1. Click **Storage** > **SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. Click the **Domain** tab.
4. Review the information about the connected domain controllers and servers.

## CIFS window

You can use the CIFS window to configure the CIFS server, to manage domain controllers, to manage symbolic UNIX mappings, and to configure BranchCache.

### Configuration tab

The Configuration tab enables you to create and manage the CIFS server.

**Server**

Specifies the status of the CIFS server, name of the server, authentication mode, name of the active directory domain, and status of SMB multichannel.

**Home Directories**

Specifies home directory paths and the style for determining how PC user names are mapped to home directory entries.

**Command buttons**

- Setup
  Opens the CIFS Setup wizard, which enables you to set up CIFS on your storage virtual machine (SVM).
- Options
  Displays the CIFS Options dialog box, which enables you to enable or disable SMB 3.0 signing, to enable or disable SMB 3.0 encryption, and to add Windows Internet Name Service (WINS) servers.
  SMB signing prevents the network traffic between the CIFS server and the client from being compromised.
- Delete
  Enables you to delete the CIFS server.
- Refresh
  Updates the information in the window.

### Domain tab

The Domain tab enables you to view and reset your CIFS domain controllers, and to add or delete preferred domain controllers. You can also use this tab to manage CIFS group policy configurations.

**Servers**

Displays information about discovered authentication servers and your preferred domain controllers on the CIFS-enabled SVM.

You can also reset the information about the discovered servers, add a preferred domain controller, delete a domain controller, or refresh the list of domain controllers.

**Group Policy**

Enables you to view, enable, or disable group policy configurations on the CIFS server. You can also reload a group policy if the status of the policy is changed.

### Symlinks tab

The Symlinks tab enables you to manage the mappings of UNIX symbolic links for CIFS users.

**Path Mappings**

Displays the list of symbolic link mappings for CIFS.

**Command buttons**

- Create
  Opens the Create New Symlink Path Mappings dialog box, which enables you to create a UNIX symbolic link mapping.
- Edit
  Opens the Edit Symlink Path Mappings dialog box, which enables you to modify the CIFS share and path.
- Delete
  Enables you to delete the symbolic link mapping.
- Refresh
  Updates the information in the window.

### BranchCache tab

The BranchCache tab enables you to set up and manage BranchCache settings on CIFS-enabled SVMs.

You can view the status of the BranchCache service, the path to the hash store, the size of the hash store, and the operating mode, server key, and version of BranchCache.

**Command buttons**

- Setup
  Opens the BranchCache Setup dialog box, which enables you to configure BranchCache for the CIFS server.
- Edit
  Opens the Modify BranchCache Settings dialog box, which enables you to modify the properties of the BranchCache configuration.
- Delete
  Enables you to delete the BranchCache configuration.
- Refresh
  Updates the information in the window.

### Related tasks

*Setting up CIFS* on page 247
You can use System Manager to enable and configure CIFS servers to allow CIFS clients to access the files on the cluster.

*Editing the general properties for CIFS* on page 248
You can modify the general properties for CIFS such as the default UNIX user and default Windows user by using System Manager. You can also enable or disable SMB signing for the CIFS server.

*Adding home directory paths* on page 248
You can use System Manager to specify one or more paths that can be used by the storage system to resolve the location of the CIFS home directories of users.

*Deleting home directory paths* on page 249

You can use System Manager to delete a home directory path when you do not want the storage system to use the path to resolve the location of the CIFS home directories of users.

*Resetting CIFS domain controllers* on page 249
You can use System Manager to reset the CIFS connection to domain controllers for the specified domain. Failure to reset the domain controller information can cause a connection failure.

# NFS protocol

You can use System Manager to authenticate NFS clients to access data on the SVM.

### Related information

*NFS management*

## Editing NFS settings

You can use System Manager to edit the NFS settings such as enabling or disabling NFSv3, NFSv4, and NFSv4.1, enabling or disabling read and write delegations for NFSv4 clients, and enabling NFSv4 ACLs. You can also edit the default Windows user.

### Steps

1. Click **Storage** > **SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **Protocols** pane, click **NFS**.
4. In the **NFS** window, click **Edit**.
5. In the **Edit NFS Settings** dialog box, make the required changes.
6. Click **Save and Close**.

### Related reference

*NFS window* on page 255
You can use the NFS window to display and configure your NFS settings.

## NFS window

You can use the NFS window to display and configure your NFS settings.

**Server Status**

Displays the status of the NFS service. The service is enabled if the NFS protocol is configured on the storage virtual machine (SVM).

**Note:** If you have upgraded to ONTAP 8.3 or later from an NFS-enabled storage system running Data ONTAP 8.1.x, the NFS service is enabled in ONTAP 8.3 or later. However, you must enable support for NFSv3 or NFSv4 because NFSv2 is no longer supported.

### Command buttons

**Enable**

Enables the NFS service.

**Disable**

Disables the NFS service.

**Edit**

Opens the Edit NFS Settings dialog box, which enables you to edit NFS settings.

**Refresh**

Updates the information in the window.

**Related tasks**

*Editing NFS settings* on page 255

You can use System Manager to edit the NFS settings such as enabling or disabling NFSv3, NFSv4, and NFSv4.1, enabling or disabling read and write delegations for NFSv4 clients, and enabling NFSv4 ACLs. You can also edit the default Windows user.

# NVMe protocol

You can use System Manager to configure the NVMe protocol. The NVMe is a transport protocol that provides high speed access to flash-based network storage. Systems that use NVMe protocol have a subsystem consisting of specific NVME controllers, namespaces, nonvolatile storage medium, hosts, ports and interface between the controller and storage medium.

## Setting up NVMe

You can set up the NVMe protocol for an SVM using System Manager. When the NVMe protocol is enabled on the SVM, you can then provision a namespace or namespaces and assign them to a host and a subsystem.

The following illustration shows the workflow for setting up NVMe:

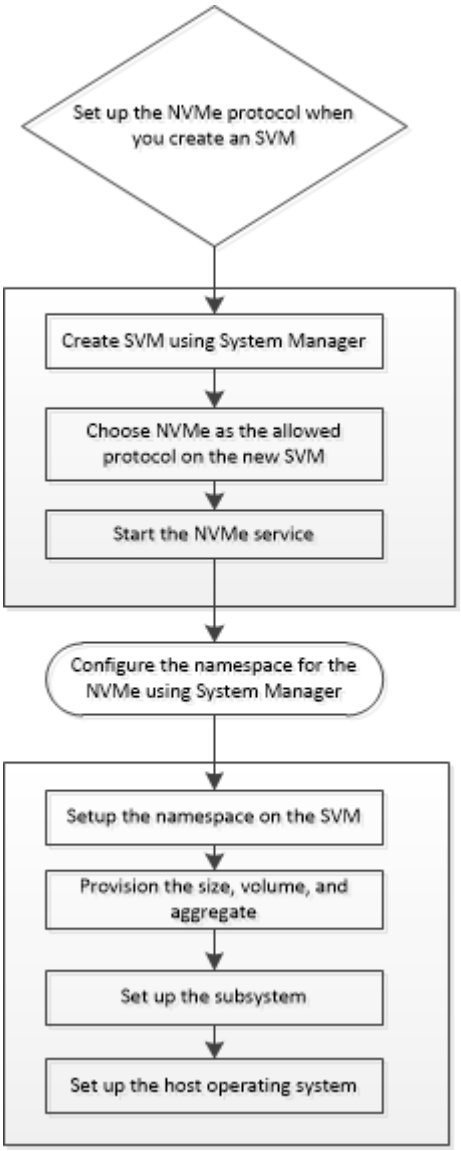


## Creating an NVMe namespace

You can use System Manager to create one or more NVMe namespaces and connect each to a host or set of hosts in a storage virtual machine (SVM). The NVMe namespace is a quantity of memory that can be formatted into logical blocks. Each namespace can be mapped to an NVMe subsystem.

### Before you begin

The SVM must already be configured with the NVMe protocol. To map a namespace, at least one LIF with the data protocol NVMe must exist in the node that owns the namespace.

### Steps

1. Click **Storage** > **NVMe** > **NVMe namespaces**.
2. Select the SVM that will contain the namespace.
3. Configure the size of the namespace (between 1MB and 16TB).
4. Select the existing volume or create a new volume by choosing the aggregate.

   Click on the + symbol to set up additional namespaces (max 250) within the SVM.
5. Select the NVMe subsystem that will be associated with this namespace.

You can choose from the following options:

- None: No subsystems are mapped.
- Use an existing subsystem: The subsystems listed are based on the selected SVM.
- Create a new subsystem: You can choose to create a new subsystem and map to all the new namespaces.

6. Select the host operating system.
7. Click **Submit**.

### Related reference

*NVMe namespaces window* on page 260

You can use the NVMe namepaces window to set up and manage your namespaces and associated subsytems for the NVMe protocol. You can search for an existing namespace using the namespace path.

## Editing an NVMe namespace

You can use System Manager to edit the namespace by changing the subsystem that the namespace is mapped to.

### About this task

You can only modify the NVMe subsystem settings in this window, you cannot edit the other namespace details.

### Steps

1. Click **NVMe** > **NVMe namespaces**.
2. In the **NVMe namespaces window**, select the namespace you want to edit.
3. Select a subsystem option:
   - None: Choosing this option unmaps the existing subsystem mapping for this namespace only. This option is preselected if no subsystem mapping is present for the selected namespace.
   - Use an existing subsystem: This option is preselected if subsystem-to-namespace mapping is present. Choosing a different subsystem maps the new subsystem by unmapping the previously mapped subsystem.

## Cloning an NVMe namespace

You can use System Manager to quickly create another namespace of the same configuration by choosing to clone a namespace. You can map the newly cloned namespace to another host NQN.

### Before you begin

You must have a FlexClone license to clone a namespace.

### About this task

You can clone a namespace with the selected host mapping and associate it with another subsystem.

### Steps

1. Click **NVMe** > **NVMe namespaces**.
2. In the **NVMe namespaces window**, select the namespace you want to clone.
3. You can rename the cloned namespace if you need a specific name but it is not required.

   The dialog provides a default name of the namespace to-be-cloned.
4. Modify the subsystem mapping for the cloned namespace.

5. Click **OK**.
   The online, mapped namespace is cloned inside the same SVM with a different name. Host mapping will not be cloned.

## Starting and stopping the NVMe service

The NVMe service enables you to manage NVMe adapters for use with namespaces. You can use System Manager to start the NVMe service to bring the adapters online. You can stop the NVMe service to take the NVMe adapters offline and to disable access to the namespaces.

### Before you begin

NVMe capable adapters must be present before you start the NVMe service.

### Steps

1. Click **Storage** > **SVMs**.
2. Select the SVM, and then click **SVM settings**.
3. In the **Protocols** menu, click **NVMe**.
4. Click **Start** or **Stop** service as required.

## What NVMe is

The nonvolatile memory express (NVMe) protocol is a transport protocol used for accessing nonvolatile storage media. This interface is used when the storage devices reside in the same physical enclosure as the host operating system or application and can be directly connected by PCIe, such as in servers or laptop computers.

NVMe is the standardized interface for applications to communicate with nonvolatile memory data storage using a PCI express (PCIe) connection.

NVMe over Fabrics (NVMeoF) is a specification-defined extension to NVMe that enables NVMe-based communication over connections other than PCIe. This interface allows for external storage enclosures to be connected to a server, either directly or through a switch, while still using NVMe as the fundamental communication mechanism.

NVMe is designed to provide efficient access to storage devices built with non-volatile memory, from flash technology to higher performing, persistent memory technologies. As such, it does not have the same limitations as storage protocols designed for hard disk drives. Flash and solid state devices (SSDs) are a type of non-volatile memory (NVM). NVM is a type of memory that keeps its content during a power outage. NVMe is a way that you can access that memory.

The benefits of NVMe include increased speeds, productivity, throughput, and capacity for data transfer. Specific characteristics include the following:

- NVMe is designed to have up to 64 thousand queues.
  Each queue in turn can have up to 64 thousand concurrent commands.
- NVMe is supported by multiple hardware and software vendors
- NMVe is more productive with Flash technologies enabling faster response times
- NVMe allows for multiple data requests for each "request" sent to the SSD.
  NVMe takes less time to decode a "request" and does not require thread locking in a multithreaded program.
- NVMe supports functionality that prevents bottlenecking at the CPU level and enables massive scalability as systems expand

## NVMe namespaces

An NVMe namespace is a quantity of non-volatile memory (NVM) that can be formatted into logical blocks. Namespaces are used when a storage virtual machine is configured with the NVMe protocol.

One or more namespaces are provisioned and connected to an NVMe host. Each namespace can support various block sizes.

The NVMe protocol provides access to namespaces through multiple controllers. Using NVMe drivers, which are supported on most operating systems, solid state drive (SSD) namespaces appear as standard-block devices on which file systems and applications can be deployed without any modification.

A namespace ID (NSID) is an identifier used by a controller to provide access to a namespace. When setting the NSID for a host or host group, you also configure the accessibility to a volume by a host. A logical block can only be mapped to a single host group at a time, and a given host group does not have any duplicate NSIDs.

## NVMe subsystems

When you create an NVMe namespace, you can associate this with a nonvolatile memory (NVM) subsystem. The subsystem includes one or more NVMe controllers, namespaces, NVM subsystem ports, an NVM storage medium and an interface between the controller and the NVM storage medium.

You provision a subsystem by providing the following details:

- The subsystem name
  The subsystem name is case sensitive. It must contain 1-64 characters and special characters are allowed.
- Host OS
  The host OS type that the subsystem is being created on.
- Host NQN
  The host NVMe qualification name attached to the controller. This column can be a comma separated value as there can be one to many hosts attached to a subsystem.

## NVMe namespaces window

You can use the NVMe namepaces window to set up and manage your namespaces and associated subsytems for the NVMe protocol. You can search for an existing namespace using the namespace path.

### Command Buttons

**Create**

Opens the NVMe namespace create dialog box, which allows you to set up a new namespace and map it to an NVMe subsystem.

**Edit**

Enables you to edit the namespace mapping.

**Delete**

Deletes the selected namespace.

**More Actions**

Allows you to create a clone of the selected namespace, which can be associated with an existing subsystem, or you can choose not to map it to a subsystem.

**Refresh**

Updates the information in the window.

### NVMe List

**Status**

Displays if the namespace is online or offline.

**Namespace Path**

The path to the new namespace in the `/vol/volume'/file` format. The namespace path is a clickable link. Clicking the link takes you to the namespace details page.

**NVMe Subsystem**

The name of the subsystem attached to a namespace. If no subsystems are attached, the value of this column is shown as `None`. You can see the list of unmapped namespaces by filtering this column for NVMe subsystem contains `None`.

**SVMs**

The SVM name on which the namespace is created. The SVM name is a clickable link. Clicking the link takes you to the existing SVM dashboard page.

**Namespace ID**

A unique identifier used by the controller to provide access to a namespace. This is not a user input; it is generated by the system when the new namespace is created.

**Total Space**

Displays the total size of the namespace.

**Used Space**

Displays the amount of used space in the namespace.

**%Used**

Displays the amount of space (in percentage) that is used in the namespace. The value for this field is calculated using total and used space.

### Details Area

You can select a namespace to view information about the selected namespace. From this area, you can also edit, delete or clone the namespace.

**Overview tab**

Displays general information about the selected namespace, and displays a pictorial representation of the space allocation of the namespace and the performance of the namespace.

In the Overview tab, the SVM and volume names are clickable links. Clicking the link takes you to the SVM and volume pages, respectively. The number of hosts can be one or more; by default two host names are shown. If more than two host names are shown, you can click a link to access the additional hosts.

The Overview tab also displays a space chart that shows the total and used space details for the namespace and a performance chart that shows details such as latency, IOPS, and throughput.

**Status:**

The status of the namespace; the value can be online or offline.

**Host NQN**

The host NVMe Qualified Names (NQNs) uniquely describes the host for the purposes of identification and authentication. This field can accept comma separated NVMe

qualification name (NQN) values. The host NQN starts with nqn and rest of the validation is the same as the initiator qualification name (IQN).

**Host OS**

The host operating system for the namespace: Hyper-V, Linux, VMware, Windows or Xen.

**Volume**

Displays the volume name on which the namespace is hosted.

**Read-Only**

Displays whether the namespace is read-only or not.

**Node**

The node that owns the namespace.

**Block Size**

The size of the storage block.

**Restore Inaccessible**

If unmapping a subsystem fails and partial data remains, unmapped namespaces cannot be restored.

# iSCSI protocol

You can use System Manager to configure the iSCSI protocol that enables you to transfer block data to hosts using SCSI protocol over TCP/IP.

### Related information

*SAN administration*

## Creating iSCSI aliases

An iSCSI alias is a user-friendly identifier that you assign to an iSCSI target device (in this case, the storage system) to make it easier to identify the target device in user interfaces. You can use System Manager to create an iSCSI alias.

### About this task

An iSCSI alias is a string of 1 to 128 printable characters. An iSCSI alias must not include spaces.

### Steps

1. Click **Storage** > **SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **Protocols** pane, click **iSCSI**.
4. In the **Service** tab of the **iSCSI** window, click **Edit**.
5. In the **Edit iSCSI Service Configuration** dialog box, enter an iSCSI alias in the **Target Alias** field, and then click **OK**.

### Related reference

*iSCSI window* on page 266

You can use the iSCSI window to start or stop the iSCSI service, change a storage system iSCSI node name, and create or change the iSCSI alias of a storage system. You can also add or change the initiator security setting for an iSCSI initiator that is connected to your storage system.

## Enabling or disabling the iSCSI service on storage system interfaces

You can use System Manager to control which network interfaces are used for iSCSI communication by enabling or disabling the interfaces. When the iSCSI service is enabled, iSCSI

connections and requests are accepted over those network interfaces that are enabled for iSCSI, but not over disabled interfaces.

### Before you begin

You must have terminated any outstanding iSCSI connections and sessions that are currently using the interface. By default, the iSCSI service is enabled on all of the Ethernet interfaces after you enable the iSCSI license.

### Steps

1. Click **Storage** > **SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **Protocols** pane, click **iSCSI**.
4. In the **iSCSI Interfaces** area, select the interface on which you want to enable or disable the iSCSI service.
5. Click **Enable** or **Disable**, as required.

### Related tasks

*Configuring iSCSI protocol on SVMs* on page 46

You can configure the iSCSI protocol on a storage virtual machine (SVM) to provide block-level data access by using System Manager. You can create iSCSI LIFs and portsets and then add the LIFs to the portsets. LIFs are created on the most suitable adapters and are assigned to portsets to ensure data path redundancy.

### Related reference

*iSCSI window* on page 266

You can use the iSCSI window to start or stop the iSCSI service, change a storage system iSCSI node name, and create or change the iSCSI alias of a storage system. You can also add or change the initiator security setting for an iSCSI initiator that is connected to your storage system.

## Adding the security method for iSCSI initiators

You can use System Manager to add an initiator and to specify the security method that is used to authenticate the initiator.

### Steps

1. Click **Storage** > **SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **Protocols** pane, click **iSCSI**.
4. In the **iSCSI** window, click the **Initiator Security** tab.
5. Click **Add** in the **Initiator Security** area.
6. Specify the initiator name and the security method for authenticating the initiator.

   For CHAP authentication, you must provide the user name and password, and then confirm your password for inbound settings. For outbound settings, this login information is optional.
7. Click **OK**.

### Related reference

*iSCSI window* on page 266

You can use the iSCSI window to start or stop the iSCSI service, change a storage system iSCSI node name, and create or change the iSCSI alias of a storage system. You can also add or change the initiator security setting for an iSCSI initiator that is connected to your storage system.

## Editing default security settings

You can use the Edit Default Security dialog box in System Manager to edit the default security settings for the iSCSI initiators that are connected to the storage system.

### Steps

1. Click **Storage** > **SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **Protocols** pane, click **iSCSI**.
4. In the **Default Security** area of the **Initiator Security** tab, click **Edit**.
5. In the **Edit Default Security** dialog box, change the security type.

   For CHAP authentication, you must provide the user name and password, and then confirm your password for inbound settings. For outbound settings, this login information is optional.
6. Click **OK**.

### Related reference

*iSCSI window* on page 266

You can use the iSCSI window to start or stop the iSCSI service, change a storage system iSCSI node name, and create or change the iSCSI alias of a storage system. You can also add or change the initiator security setting for an iSCSI initiator that is connected to your storage system.

## Editing initiator security

The security style that is configured for an initiator specifies how authentication is done for that initiator during the iSCSI connection login phase. You can use System Manager to change the security for selected iSCSI initiators by changing the authentication method.

### Steps

1. Click **Storage** > **SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **Protocols** pane, click **iSCSI**.
4. In the **Initiator Security** tab, select one or more initiators from the initiator list, and then click **Edit** in the **Initiator Security** area.
5. Change the security type.

   For CHAP authentication, you must provide the user name and password, and then confirm your password for inbound settings. For outbound settings, this login information is optional.
6. Click **OK**.
7. Verify the changes that you made in the **Initiator Security** tab.

### Related reference

*iSCSI window* on page 266

You can use the iSCSI window to start or stop the iSCSI service, change a storage system iSCSI node name, and create or change the iSCSI alias of a storage system. You can also add or change the initiator security setting for an iSCSI initiator that is connected to your storage system.

## Changing the default iSCSI initiator authentication method

You can use System Manager to change the default iSCSI authentication method, which is the authentication method that is used for any initiator that is not configured with a specific authentication method.

### Steps

1. Click **Storage** > **SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **Protocols** pane, click **iSCSI**.
4. In the **Initiator Security** tab, click **Edit** in the **Default Security** area.
5. Change the security type.

   For CHAP authentication, you must provide the user name and password, and then confirm your password for inbound settings. For outbound settings, this login information is optional.
6. Click **OK**.

### Related reference

*iSCSI window* on page 266

You can use the iSCSI window to start or stop the iSCSI service, change a storage system iSCSI node name, and create or change the iSCSI alias of a storage system. You can also add or change the initiator security setting for an iSCSI initiator that is connected to your storage system.

## Setting the default security for iSCSI initiators

You can use System Manager to remove the authentication settings for an initiator and to use the default security method to authenticate the initiator.

### Steps

1. Click **Storage** > **SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **Protocols** pane, click **iSCSI**.
4. In the **Initiator Security** tab, select the initiator for which you want to change the security setting.
5. Click **Set Default** in the **Initiator Security** area, and then click **Set Default** in the confirmation dialog box.

### Related reference

*iSCSI window* on page 266

You can use the iSCSI window to start or stop the iSCSI service, change a storage system iSCSI node name, and create or change the iSCSI alias of a storage system. You can also add or change the initiator security setting for an iSCSI initiator that is connected to your storage system.

## Starting or stopping the iSCSI service

You can use System Manager to start or stop the iSCSI service on your storage system.

### Steps

1. Click **Storage** > **SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **Protocols** pane, click **iSCSI**.

4. Click **Start** or **Stop**, as required.

### Related reference

*iSCSI window* on page 266

You can use the iSCSI window to start or stop the iSCSI service, change a storage system iSCSI node name, and create or change the iSCSI alias of a storage system. You can also add or change the initiator security setting for an iSCSI initiator that is connected to your storage system.

## Viewing initiator security information

You can use System Manager to view the default authentication information and all the initiator-specific authentication information.

### Steps

1. Click **Storage** > **SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **Protocols** pane, click **iSCSI**.
4. In the **Initiator Security** tab of the **iSCSI** window, review the details.

## iSCSI window

You can use the iSCSI window to start or stop the iSCSI service, change a storage system iSCSI node name, and create or change the iSCSI alias of a storage system. You can also add or change the initiator security setting for an iSCSI initiator that is connected to your storage system.

### Tabs

**Service**

You can use the **Service** tab to start or stop the iSCSI service, change a storage system iSCSI node name, and create or change the iSCSI alias of a storage system.

**Initiator Security**

You can use the **Initiator Security** tab to add or change the initiator security setting for an iSCSI initiator that is connected to your storage system.

### Command buttons

**Edit**

Opens Edit iSCSI Service Configurations dialog box, which enables you to change iSCSI node name and iSCSI alias of the storage system.

**Start**

Starts the iSCSI service.

**Stop**

Stops the iSCSI service.

**Refresh**

Updates the information in the window.

### Details area

The details area displays information about the status of the iSCSI service, iSCSI target node name, and iSCSI target alias. You can use this area to enable or disable the iSCSI service on a network interface.

### Related tasks

*Creating iSCSI aliases* on page 262

An iSCSI alias is a user-friendly identifier that you assign to an iSCSI target device (in this case, the storage system) to make it easier to identify the target device in user interfaces. You can use System Manager to create an iSCSI alias.

*Enabling or disabling the iSCSI service on storage system interfaces* on page 262
You can use System Manager to control which network interfaces are used for iSCSI communication by enabling or disabling the interfaces. When the iSCSI service is enabled, iSCSI connections and requests are accepted over those network interfaces that are enabled for iSCSI, but not over disabled interfaces.

*Adding the security method for iSCSI initiators* on page 263
You can use System Manager to add an initiator and to specify the security method that is used to authenticate the initiator.

*Editing default security settings* on page 264
You can use the Edit Default Security dialog box in System Manager to edit the default security settings for the iSCSI initiators that are connected to the storage system.

*Editing initiator security* on page 264
The security style that is configured for an initiator specifies how authentication is done for that initiator during the iSCSI connection login phase. You can use System Manager to change the security for selected iSCSI initiators by changing the authentication method.

*Changing the default iSCSI initiator authentication method* on page 265
You can use System Manager to change the default iSCSI authentication method, which is the authentication method that is used for any initiator that is not configured with a specific authentication method.

*Setting the default security for iSCSI initiators* on page 265
You can use System Manager to remove the authentication settings for an initiator and to use the default security method to authenticate the initiator.

*Starting or stopping the iSCSI service* on page 265
You can use System Manager to start or stop the iSCSI service on your storage system.

# FC/FCoE protocol

You can use System Manager to configure FC/FCoE protocols.

**Related information**

*SAN administration*

## Starting or stopping the FC or FCoE service

The FC service enables you to manage FC target adapters for use with LUNs. You can use System Manager to start the FC service to bring the adapters online and to enable access to the LUNs on the storage system. You can stop the FC service to take the FC adapters offline and to disable access to the LUNs.

### Before you begin

- The FC license must be installed.
- An FC adapter must be present in the target storage system.

### Steps

1. Click **Storage** > **SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **Protocols** pane, click **FC/FCoE**.
4. Click **Start** or **Stop**, as required.

### Related reference

*FC/FCoE window* on page 268

You can use the FC/FCoE window to start or stop the FC service.

## Changing an FC or FCoE node name

If you replace a storage system chassis and reuse it in the same Fibre Channel SAN, the node name of the replaced storage system might be duplicated in certain cases. You can change the node name of the storage system by using System Manager.

### Steps

1. Click **Storage** > **SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **Protocols** pane, click **FC/FCoE**.
4. Click **Edit**.
5. Type the new name, and then click **OK**.

### Related reference

*FC/FCoE window* on page 268

You can use the FC/FCoE window to start or stop the FC service.

## The FCoE protocol

Fibre Channel over Ethernet (FCoE) is a new model for connecting hosts to storage systems. Like the traditional FC protocol, FCoE maintains existing FC management and controls, but it uses a 10-gigabit Ethernet network as the hardware transport.

Setting up an FCoE connection requires one or more supported converged network adapters (CNAs) in the host, connected to a supported data center bridging (DCB) Ethernet switch. The CNA is a consolidation point and effectively serves as both an HBA and an Ethernet adapter.

In general, you can configure and use FCoE connections the same way you use traditional FC connections.

## FC/FCoE window

You can use the FC/FCoE window to start or stop the FC service.

### Command buttons

**Edit**

Opens the Edit Node Name dialog box, which enables you to change the FC or FCoE node name.

**Start**

Starts the FC/FCoE service.

**Stop**

Stops the FC/FCoE service.

**Refresh**

Updates the information in the window.

### FC/FCoE details

The details area displays information about the status of FC/FCoE service, the node name, and the FC/FCoE adapters.

### Related tasks

*Starting or stopping the FC or FCoE service* on page 267
The FC service enables you to manage FC target adapters for use with LUNs. You can use System Manager to start the FC service to bring the adapters online and to enable access to the LUNs on the storage system. You can stop the FC service to take the FC adapters offline and to disable access to the LUNs.

*Changing an FC or FCoE node name* on page 268
If you replace a storage system chassis and reuse it in the same Fibre Channel SAN, the node name of the replaced storage system might be duplicated in certain cases. You can change the node name of the storage system by using System Manager.

*Configuring FC protocol and FCoE protocol on SVMs* on page 48
You can configure the FC protocol and the FCoE protocol on the storage virtual machine (SVM) for SAN hosts. LIFs are created on the most suitable adapters and are assigned to port sets to ensure data path redundancy. Based on your requirements, you can configure either the FC protocol or the FCoE protocols, or both the protocols by using System Manager.

# Export policies

You can use System Manager to create, edit, and manage export policies.

## Creating an export policy

You can use System Manager to create an export policy so that clients can access specific volumes.

### Steps

1. Click **Storage** > **SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **Policies** pane, click **Export Policies**.
4. Click **Create**.
5. In the **Create Export Policy** dialog box, specify a name for the export policy.
6. If you want to create an export policy by copying the rules from an existing export policy, select the **Copy Rules from** check box, and then select the storage virtual machine (SVM) and the export policy.

   You should not select the destination SVM for disaster recovery from the drop-down menu to create an export policy.
7. In the **Export Rules** area, click **Add** to add rules to the export policy.
8. Click **Create**.
9. Verify that the export policy that you created is displayed in the **Export Policies** window.

## Renaming export policies

You can use System Manager to rename an existing export policy.

### Steps

1. Click **Storage** > **SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **Policies** pane, click **Export Policies**.
4. Select the export policy that you want to rename, and then click **Rename Policy**.
5. In the **Rename Policy** dialog box, specify a new policy name, and then click **Modify**.
6. Verify the changes that you made in the **Export Policies** window.

## Deleting export policies

You can use System Manager to delete export policies that are no longer required.

### Steps

1. Click **Storage** > **SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **Policies** pane, click **Export Policies**.
4. Select the export policy that you want to delete, and then click **Delete Policy**.
5. Select the confirmation check box, and then click **Delete**.

## Adding rules to an export policy

You can use System Manager to add rules to an export policy, which enables you to define client access to data.

### Before you begin

You must have created the export policy to which you want to add the export rules.

### Steps

1. Click **Storage** > **SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **Policies** pane, click **Export Policies**.
4. Select the export policy to which you want to add rules, and from the **Export Rules** tab, click **Add**.
5. In the **Create Export Rule** dialog box, perform the following steps:
   a. Specify the client that requires access to the data.

      You can specify multiple clients as comma-separated values.

      You can specify the client in any of the following formats:

      - As a host name; for instance, host1
      - As an IPv4 address; for instance, 10.1.12.24
      - As an IPv4 address with a network mask; for instance, 10.1.16.0/255.255.255.0
      - As an IPv6 address; for instance, FE80::0202:B3FF:FE1E:8329
      - As an IPv6 address with a network mask; for instance, 2001:db8::/32
      - As a netgroup, with the netgroup name preceded by an at symbol (@); for instance, @netgroup
      - As a domain name preceded by a period (.); for instance, .example.com

      **Note:** You must not enter an IP address range, such as 10.1.12.10 through 10.1.12.70. Entries in this format are interpreted as a text string and are treated as a host name.

      You can enter the IPv4 address **`0.0.0.0/0`** to provide access to all of the hosts.

   b. If you want to modify the rule index number, select the appropriate rule index number.
   c. Select one or more access protocols.

      If you do not select any access protocol, the default value "Any" is assigned to the export rule.

   d. Select one or more security types and access rules.
6. Click **OK**.
7. Verify that the export rule that you added is displayed in the **Export Rules** tab for the selected export policy.

## Modifying export policy rules

You can use System Manager to modify the specified client, access protocols, and access permissions of an export policy rule.

### Steps

1. Click **Storage** > **SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **Policies** pane, click **Export Policies**.
4. In the **Export Policies** window, select the export policy for which you want to edit the export rule, and in the **Export Rules** tab, select the rule that you want to edit, and then click **Edit**.
5. Modify the following parameters as required:
   - Client specification
   - Access protocols
   - Access details
6. Click **OK**.
7. Verify that the updated changes for the export rule are displayed in the **Export Rules** tab.

### Related tasks

*Setting up CIFS* on page 247
You can use System Manager to enable and configure CIFS servers to allow CIFS clients to access the files on the cluster.

## Deleting export policy rules

You can use System Manager to delete export policy rules that are no longer required.

### Steps

1. Click **Storage** > **SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **Policies** pane, click **Export Policies**.
4. Select the export policy for which you want to delete the export rule.
5. In the **Export Rules** tab, select the export rule that you want to delete, and then click **Delete**.
6. In the confirmation box, click **Delete**.

## How export policies control client access to volumes or qtrees

Export policies contain one or more *export rules* that process each client access request. The result of the process determines whether the client is denied or granted access and what level of access. An export policy with export rules must exist on the storage virtual machine (SVM) for clients to access data.

You associate exactly one export policy with each volume or qtree to configure client access to the volume or qtree. The SVM can contain multiple export policies. This enables you to do the following for SVMs with multiple volumes or qtrees:

- Assign different export policies to each volume or qtree of the SVM for individual client access control to each volume or qtree in the SVM.
- Assign the same export policy to multiple volumes or qtrees of the SVM for identical client access control without having to create a new export policy for each volume or qtree.

If a client makes an access request that is not permitted by the applicable export policy, the request fails with a permission-denied message. If a client does not match any rule in the export policy, then access is denied. If an export policy is empty, then all accesses are implicitly denied.

You can modify an export policy dynamically on a system running ONTAP.

## Export Policies window

You can use the Export Policies window to create, view, and manage information about export policies and its related export rules.

### Export Policies

The Export Policies window enables you to view and manage the export policies created for the storage virtual machine (SVM).

**Command buttons**

- Create
  Opens the Create Export Policy dialog box, which enables you to create an export policy and add export rules. You can also copy export rules from an existing SVM.
- Rename
  Opens the Rename Policy dialog box, which enables you to rename the selected export policy.
- Delete
  Opens the Delete Export Policy dialog box, which enables you to delete the selected export policy.
- Refresh
  Updates the information in the window.

### Export Rules tab

The Export Rules tab enables you to view information about the export rules created for a particular export policy. You can also add, edit, and delete rules.

**Command buttons**

- Add
  Opens the Create Export Rule dialog box, which enables you to add an export rule to the selected export policy.
- Edit
  Opens the Modify Export Rule dialog box, which enables you to modify the attributes of the selected export rule.
- Delete
  Opens the Delete Export Rule dialog box, which enables you to delete the selected export rule.
- Move Up
  Moves up the rule index of the selected export rule.
- Move Down
  Moves down the rule index of the selected export rule.
- Refresh
  Updates the information in the window.

**Export rules list**

- Rule Index
  Specifies the priority based on which the export rules are processed. You can use the Move Up and Move Down buttons to choose the priority.
- Client
  Specifies the client to which the rule applies.
- Access Protocols
  Displays the access protocol that is specified for the export rule.

If you have not specified any access protocol, the default value "Any" is considered.
- Read-Only Rule
  Specifies one or more security types for read-only access.
- Read/Write Rule
  Specifies one or more security types for read/write access.
- Superuser Access
  Specifies the security type or types for superuser access.

### Assigned Objects tab

The Assigned Objects tab enables you to view the volumes and qtrees that are assigned to the selected export policy. You can also view whether the volume is encrypted or not.

# Efficiency policies

You can use System Manager to create, edit, and delete efficiency policies.

## Adding efficiency policies

You can use System Manager to add efficiency policies for running the deduplication operation on a volume on a specified schedule or when the change in volume data reaches a specified threshold value.

### Steps

1. Click **Storage** > **SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **Policies** pane, click **Efficiency Policies**.
4. Click **Add**, and then specify the policy name.
5. Specify how the storage efficiency policy should be run:
   - Select **Schedule**, and specify the schedule name and the schedule details.
     You can specify the maximum run-time duration of the efficiency policy, if required.
   - Select **ChangeLog Threshold**, and specify the threshold value (in percent) for the change in volume data.
6. Optional: Select the **Set QoS policy to background** check box to reduce performance impact on client operations.
7. Click **Add**.

## Editing efficiency policies

You can use System Manager to modify the attributes of an efficiency policy such as the policy name, schedule name, and maximum runtime.

### Steps

1. Click **Storage** > **SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **Policies** pane, click **Efficiency Policies**.
4. In the **Efficiency Policies** window, select the policy that you want to edit, and then click **Edit**.
5. In the **Edit Efficiency Policy** dialog box, make the required changes.
6. Click **Save**.

## Deleting efficiency policies

You can use System Manager to delete an efficiency policy that is no longer required.

### Before you begin

The efficiency policy must be disabled.

### Steps

1. Click **Storage** > **SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **Policies** pane, click **Efficiency Policies**.
4. Select the efficiency policy that you want to delete, and then click **Delete**.
5. Select the confirmation check box, and then click **Delete**.

## Enabling or disabling efficiency policies

You can use System Manager to enable or disable an efficiency policy.

### Steps

1. Click **Storage** > **SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **Policies** pane, click **Efficiency Policies**.
4. Select one or more efficiency policies that you want to enable or disable.
5. Click **Status** > **Enable** or **Status** > **Disable**, as required.
6. If you are disabling an efficiency policy, select the confirmation check box, and then click **OK**.

## What an efficiency policy is

An efficiency policy is a job schedule for a deduplication operation on a volume .

You can run deduplication on a volume either by scheduling the operations to start at a specific time or by specifying a threshold percentage after which the operations are triggered. You can schedule a deduplication operation by creating job schedules that are enclosed within the efficiency policies or you can specify a threshold percentage, which waits for the new data to exceed the specified percentage and then triggers the deduplication. The volume efficiency policies support only job schedules that are of type cron.

## Understanding predefined efficiency policies

You can configure a volume with efficiency policies to achieve additional space savings. You can configure a volume to run inline compression without a scheduled or manually started background efficiency operation configured on the volume.

When you create an SVM, the following efficiency policies are created automatically and cannot be deleted:

- Default
  You can configure a volume with the efficiency policy to run the scheduled deduplication operations on the volume.
- Inline-only
  You can configure a volume with the inline-only efficiency policy and enable inline compression, to run inline compression on the volume without any scheduled or manually started background efficiency operations.

For more information about the inline-only and default efficiency policies, see the man pages.

## Efficiency Policies window

You can use the Efficiency Policies window to create, display, and manage information about efficiency policies.

### Command buttons

**Add**

Opens the Add Efficiency Policy dialog box, which enables you to run a deduplication operation on a volume for a specified duration (schedule-based) or when the change in volume data reaches a specified threshold value (threshold-based).

**Edit**

Opens the Edit Efficiency Policy dialog box, which enables you to modify the schedule, threshold value, QoS type, and maximum run time for a deduplication operation.

**Delete**

Opens the Delete Efficiency Policy dialog box, which enables you to delete the selected efficiency policy.

**Status**

Open a drop-down menu, which provides options to enable or disable the selected efficiency policy.

**Refresh**

Updates the information in the window.

### Efficiency policies list

**Policy**

Specifies the name of an efficiency policy.

**Status**

Specifies the status of an efficiency policy. The status can be one of the following:

- Enabled
  Specifies that the efficiency policy can be assigned to a deduplication operation.
- Disabled
  Specifies that the efficiency policy is disabled. You can enable the policy by using the status drop-down menu and assign it later to a deduplication operation.

**Run By**

Specifies whether the storage efficiency policy is run based on a schedule or based on a threshold value (change log threshold).

**QoS Policy**

Specifies the QoS type for the storage efficiency policy. The QoS type can be one of the following:

- Background
  Specifies that the QoS policy is running in the background, which reduces potential performance impact on the client operations.
- Best-effort
  Specifies that the QoS policy is running on a best-effort basis, which enables you to maximize the utilization of system resources.

**Maximum Runtime**

Specifies the maximum run-time duration of an efficiency policy. If this value is not specified, the efficiency policy is run till the operation is complete.

### Details area

The area below the efficiency policy list displays additional information about the selected efficiency policy, including the schedule name and the schedule details for a schedule-based policy, and the threshold value for a threshold-based policy.

# Protection policies

You can use System Manager to create, edit, and delete protection policies.

## Creating protection policies

You can use System Manager to create asynchronous mirror policies, vault policies, or mirror and vault policies, and to apply these policies to a data protection relationship.

### Steps

1. Click **Storage** > **SVMs**.
2. Select the storage virtual machine (SVM) for which you want to create a protection policy, and then click **SVM Settings**.
3. In the **Policies** pane, click **Protection Policies**.
4. Click **Create**.
5. In the **Create Policy** dialog box, select the policy type that you want to create.
6. Specify the policy name and transfer priority.

   Low indicates that the transfer has the least priority, and the transfer is usually scheduled after normal priority transfers. By default, the priority is set to Normal.
7. Optional: For a policy of type asynchronous mirror, select the **Transfer All Source Snapshot Copies** check box to include the "all_source_snapshots" rule to the mirror policy, which backs up all of the Snapshot copies from the source volume.
8. Optional: Select the **Enable Network Compression** check box to compress the data that is being transferred during a data transfer.
9. Optional: Click **Add Comments** to add additional comments for the policy.
10. For a policy of type vault or mirror vault, specify a SnapMirror label and a destination retention count.
11. Click **Create**.

## Deleting protection policies

You can use System Manager to delete a protection policy if you no longer want to use the policy.

### About this task

The cluster-level mirror policies or vault policies are not displayed.

### Steps

1. Click **Storage** > **SVMs**.
2. Select the storage virtual machine (SVM), and then click **SVM Settings**.
3. In the **Protection Policies** window, select the policy that you want to delete, and then click **Delete**.
4. In the **Delete Policy** dialog box, click **Delete**.

## Editing protection policies

You can use System Manager to modify a protection policy and to apply the policy to a data protection relationship.

### About this task

The protection policies are not displayed at the cluster level.

### Steps

1. Click **Storage** > **SVMs**.
2. Select the storage virtual machine (SVM), and then click **SVM Settings**.
3. In the **Policies** pane, click **Protection Policies**.
4. Select the protection policy that you want to edit, and then click **Edit**.
5. Modify the transfer priority, and then enable or disable network compression.
6. For an asynchronous mirror policy, back up all of the source Snapshot copies.
7. For a vault policy or mirror vault policy, modify the SnapMirror label and retention count.

   You cannot remove the sm_created label for a mirror vault policy.
8. Click **Save**.

## Protection Policies window

You can use the Protection Policies window to create, manage, and display information about mirror, vault, and mirror vault policies.

- *Command buttons* on page 277
- *Protection policies list*
- *Details area*

### Command buttons

**Create**

Opens the Create Policy dialog box, which enables you to create a mirror, vault, or mirror vault policy.

**Edit**

Opens the Edit Policy dialog box, which enables you to edit a policy.

**Delete**

Opens the Delete Policy dialog box, which enables you to delete a policy.

**Refresh**

Updates the information in the window.

### Protection policies list

**Name**

Displays the name of the protection policy.

**Type**

Displays the policy type, which can be Vault, Mirror Vault, or Asynchronous Mirror.

**Comment**

Displays the description specified for the policy.

**Transfer Priority**

Displays the data transfer priority, such as Normal or Low.

### Details area

**Policy Details tab**

Displays details of the protection policy, such as the user who created the policy, number of rules, retention count, and status of network compression.

**Policy Rules tab**

Displays details of the rules that are applied to the policy. The Policy Rules tab is displayed only if the selected policy contains rules.

# QoS policy groups

You can use System Manager to create, edit, and delete QoS policy groups.

## Creating QoS policy groups

You can use System Manager to create storage Quality of Service (QoS) policy groups to limit the throughput of workloads and to monitor workload performance.

### Steps

1. Click **Storage** > **SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **Policies** pane, click **QoS Policy Groups**.
4. In the **QoS Policy Groups** window, click **Create**.
5. In the **Create Policy Group** dialog box, specify a group name for the policy.
6. Specify the minimum throughput limit.
   - You can set the minimum throughput limit only on an AFF platform.
   - You cannot set the minimum throughput limit for volumes on a FabricPool-enabled aggregate.
   - If you do not specify the minimum throughput value or if the minimum throughput value is set to 0, the system automatically displays "None" as the value.

     This value is case-sensitive.
7. Specify the maximum throughput limit.
   - The minimum throughput limit and the maximum throughput limit must be of the same unit type.
   - If you do not specify the minimum throughput limit, you can set the maximum throughput limit in IOPS and B/s, KB/s, MB/s, and so on.
   - If you do not specify the maximum throughput limit, the system automatically displays "Unlimited" as the value.

     This value is case-sensitive. The unit that you specify does not affect the maximum throughput.
8. Click **OK**.

## Deleting QoS policy groups

You can use System Manager to delete a Storage Quality of Service (QoS) policy group that is no longer required.

### Before you begin

You must have unassigned all of the storage objects that are assigned to the policy group.

### Steps

1. Click **Storage** > **SVMs**.

2. Select the SVM, and then click **SVM Settings**.
3. In the **Policies** pane, click **QoS Policy Groups**.
4. In the **QoS Policy Groups** window, select the policy group that you want to delete, and then click **Delete**.
5. In the confirmation dialog box, click **Delete**.

## Editing QoS policy groups

You can use the Edit Policy Group dialog box in System Manager to modify the name and maximum throughput of an existing storage Quality of Service (QoS) policy group.

### About this task

- You can set the minimum throughput limit only on an AFF platform.
- You cannot set the minimum throughput limit for volumes on a FabricPool-enabled aggregate.

### Steps

1. Click **Storage** > **SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **Policies** pane, click **QoS Policy Groups**.
4. Select the QoS policy group that you want to edit, and then click **Edit**.
   - The minimum throughput limit and the maximum throughput limit must be of the same unit type.
   - If you do not specify the minimum throughput limit, you can set the maximum throughput limit in IOPS and B/s, KB/s, MB/s, and so on.
   - If you do not specify the maximum throughput limit, the value is set to unlimited, and the unit that you specify does not affect the maximum throughput.
5. In the **Edit Policy Group** dialog box, edit the QoS policy group details, and then click **Save**.

## Managing workload performance by using Storage QoS

Storage Quality of Service (QoS) can help you manage risks around meeting your performance objectives. You can use Storage QoS to limit the throughput to workloads and to monitor workload performance. You can reactively limit workloads to address performance problems, and you can proactively limit workloads to prevent performance problems.

A workload represents the input/output (I/O) operations to one of the following kinds of storage objects:

- FlexVol volumes
- LUNs
- FlexGroup volumes

You can assign a storage object to a policy group to control and monitor a workload. You can monitor workloads without controlling them.

The following illustration shows a sample environment before and after using Storage QoS. On the left, the workloads compete for cluster resources to transmit I/O. These workloads get "best effort" performance, which means that you have less performance predictability (for example, a workload might get such good performance that it negatively impacts other workloads). On the right, the same workloads are assigned to policy groups. The policy groups enforce a maximum throughput limit.

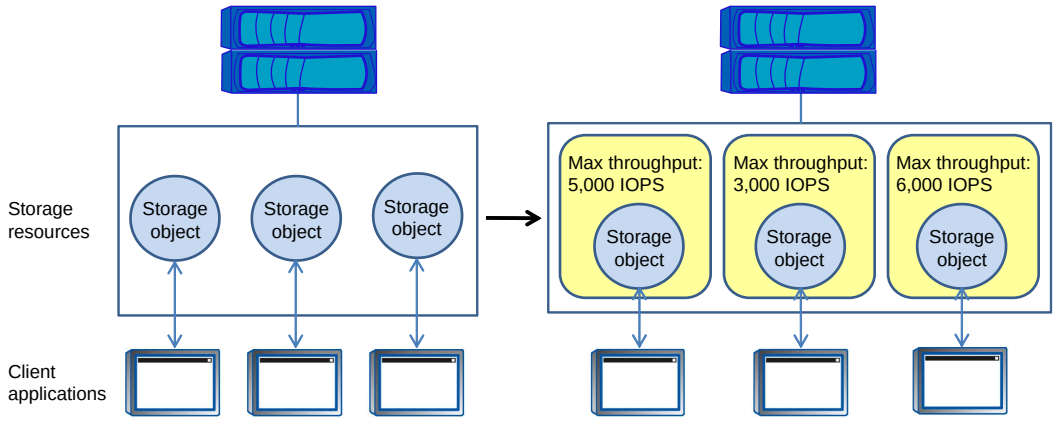


## How Storage QoS works

Storage QoS controls workloads that are assigned to policy groups by throttling and prioritizing client operations (SAN and NAS data requests) and system operations.

The following illustration shows a sample environment before and after using Storage QoS. On the left, workloads compete for cluster resources to transmit I/O. These workloads get "best effort" performance, which means that you have less performance predictability (for example, a workload might get such good performance that it negatively impacts other workloads). On the right, the same workloads are assigned to policy groups that enforce maximum throughput limits.

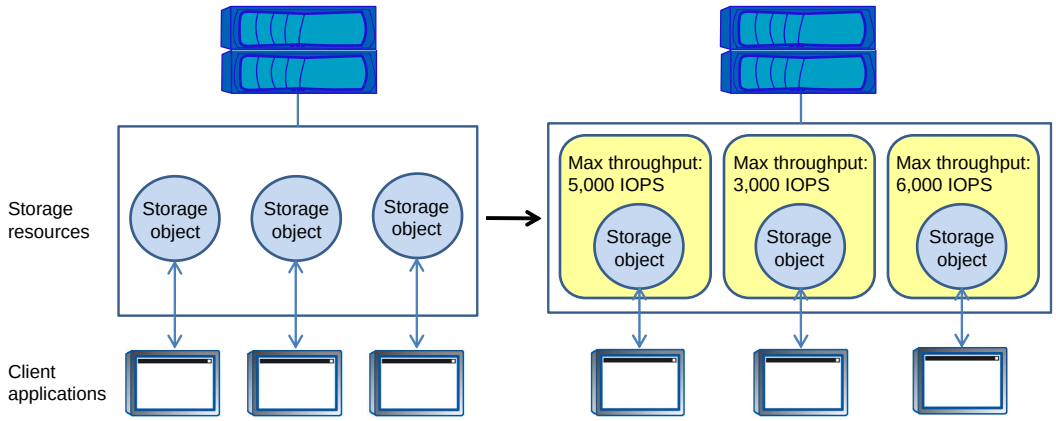


The `-max-throughput` parameter specifies the maximum throughput limit for the policy group that the policy group must not exceed. The value of this parameter is specified in terms of IOPS or MB/s, or a combination of comma-separated IOPS and MB/s values, and the range is zero to infinity.

The units are base 10. There should be no space between the number and the unit. The default value for the `-max-throughput` parameter is **`infinity`**, which is specified by the special value INF.

> **Note:** There is no default unit for the `-max-throughput` parameter. For all values except zero and infinity, you must specify the unit.

The keyword "none" is available for a situation that requires the removal of a value. The keyword "INF" is available for a situation that requires the maximum available value to be specified. Examples of valid throughput specifications are: ""100B/s"", "10KB/s", "1gb/s", "500MB/s", "1tb/s", "100iops", "100iops,400KB/s", and "800KB/s,100iops".

## How the maximum throughput limit works

You can specify one service-level objective for a Storage QoS policy group: a maximum throughput limit. A maximum throughput limit, which you define in terms of IOPS, MBps, or both, specifies the throughput that the workloads in the policy group cannot collectively exceed.

When you specify a maximum throughput for a policy group, Storage QoS controls client operations to ensure that the combined throughput for all workloads in the policy group does not exceed the specified maximum throughput.

For example, assume that you create the policy group "untested_apps" and specify a maximum throughput of 300 MBps. You assign three volumes to the policy group. The combined throughput to those three volumes cannot exceed 300 MBps.

**Note:** The combined throughput to the workloads in a policy group might exceed the specified limit by up to 10 percent. A deviation might occur if you have a workload that experiences rapid changes in throughput (sometimes called a *bursty workload*).

Note the following about specifying a maximum throughput:

- You must not set the limit too low because you might underutilize the cluster.
- You must consider the minimum amount of throughput that you want to reserve for workloads that do not have limits.

  For example, you can ensure that your critical workloads get the throughput that they need by limiting noncritical workloads.
- You might want to provide room for growth.

  For example, if you see an average utilization of 500 IOPS, you might specify a limit of 1,000 IOPS.

## Rules for assigning storage objects to policy groups

You should be aware of rules that dictate how you can assign storage objects to Storage QoS policy groups.

### Storage objects and policy groups must belong to the same SVM

A storage object must be contained by the SVM to which the policy group belongs. You specify the SVM to which the policy group belongs when you create the policy group. Multiple policy groups can belong to the same SVM.

In the following illustration, the policy group pg1 belongs to SVM vs1. You cannot assign volumes vol2 or vol3 to policy group pg1 because those volumes are contained by a different SVM.

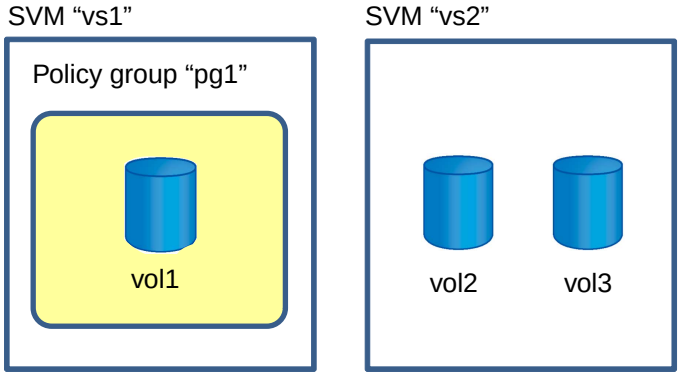

### Nested storage objects cannot belong to policy groups

You cannot assign a storage object to a policy group if its containing object or its child objects belong to a policy group. The following table lists the restrictions.

| If you assign the... | Then you cannot assign... |
|---|---|
| SVM to a policy group | Any storage objects contained by the SVM to a policy group |
| Volume to a policy group | The volume's containing SVM or any child LUNs to a policy group |
| LUN to a policy group | The LUN's containing volume or SVM to a policy group |
| File to a policy group | The file's containing volume or SVM to a policy group |

In the following illustration, the SVM vs3 is assigned to policy group pg2. You cannot assign volumes vol4 or vol5 to a policy group because an object in the storage hierarchy (SVM vs3) is assigned to a policy group.

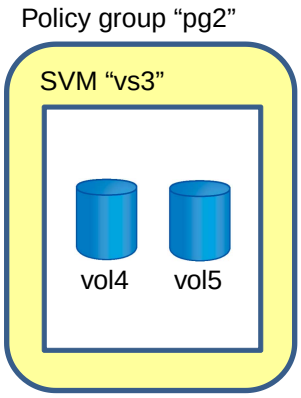


## QoS Policy Groups window

Storage QoS (Quality of Service) can help you manage risks related to meeting your performance objectives. Storage QoS enables you to limit the throughput of workloads and to monitor workload performance. You can use the QoS Policy groups window to manage your policy groups and view information about them.

- *Command buttons* on page 282
- *QoS Policy Groups list* on page 283
- *Details area* on page 283

### Command buttons

**Create**

Opens the Create QoS Policy Group dialog box, which enables you to create new policy groups.

**Edit**

Opens the Edit QoS Policy Group dialog box, which enables you to modify the selected policy group.

**Delete**

Deletes the selected policy groups.

**Refresh**

Updates the information in the window.

### QoS Policy Groups list

The QoS Policy Groups list displays the policy group name and the maximum throughput for each policy group.

**Name**

Displays the name of the QoS policy group.

**Minimum Throughput**

Displays the minimum throughput limit specified for the policy group.

If you have not specified any minimum throughput value, the system automatically displays "None" as the value and this value is case-sensitive.

**Maximum Throughput**

Displays the maximum throughput limit specified for the policy group.

If you have not specified any maximum throughput value, the system automatically displays "Unlimited" as the value and this value is case-sensitive.

**Storage Objects Count**

Displays the number of storage objects assigned to the policy group.

### Details area

The area below the QoS Policy Groups list displays detailed information about the selected policy group.

**Assigned Storage Objects tab**

Displays the name and type of the storage object that is assigned to the selected policy group.

# NIS services

You can use System Manager to add, edit, and manage Network Information Service (NIS) domains.

### Related information

*NFS configuration*

## Adding NIS domains

You can maintain host information centrally by using NIS. You can use System Manager to add the NIS domain name of your storage system. Only one NIS domain can be active on a storage virtual machine (SVM) at any given time.

### Steps

1. Click **Storage** > **SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **Services** pane, click **NIS**.
4. Click **Create**.
5. Type the NIS domain name, and then add one or more NIS servers.
6. Click **Create**.

## Editing NIS domains

You can use System Manager to modify NIS domains based on the requirement for storage virtual machine (SVM) authentication and authorization.

### Steps

1. Click **Storage** > **SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **Services** pane, click **NIS**.
4. Select the NIS domain, and then click **Edit**.
5. Make the required changes, and then click **Edit**.

## NIS window

The NIS window enables you to view the current NIS settings for your storage system.

### Command buttons

**Create**

Opens the Create NIS Domain dialog box, which enables you to create NIS domains.

**Edit**

Opens the Edit NIS Domain dialog box, which enables you to add, delete, or modify NIS servers.

**Delete**

Deletes the selected NIS domain.

**Refresh**

Updates the information in the window.

# LDAP client services

You can use System Manager to add, edit, and delete LDAP client configurations.

## Adding an LDAP client configuration

You can use System Manager to add an LDAP client configuration at the cluster level or the storage virtual machine (SVM) level if you want to use LDAP services. You must first set up an LDAP client to use LDAP services.

### About this task

At the SVM level, you can add an LDAP client only for a selected SVM.

### Steps

1. Add an LDAP client configuration by using one of the following methods:
   - Cluster level: click ⚙ > **LDAP**.
   - SVM level: click **SVM** > **SVM Settings** > **LDAP Client**.
2. Click **Add**.
3. Type the name of the LDAP client.
4. Add either the Active Directory domain or the LDAP server.
5. Click ⚙ (advanced options), select the **Schema**, and click **Apply**.
6. Specify the **Base DN** and **TCP Port**.
7. Click **Binding**, and then specify the authentication details.

8. Click **Save and Close**.
9. Verify that the LDAP client that you added is displayed.

### Related concepts

*LDAP* on page 94
You can use System Manager to configure an LDAP server that centrally maintains user information.

## Deleting an LDAP client configuration

You can use System Manager to delete an LDAP client configuration at the cluster level or the storage virtual machine (SVM) level.

### About this task

At the SVM level, you can delete an LDAP client only for a selected SVM.

### Steps

1. To delete an LDAP client configuration:
   - Cluster level: Click ⚙ > **LDAP**.
   - SVM level: Click **SVM** > **SVM Settings** > **LDAP Client**.
2. Select the LDAP client that you want to delete, and then click **Delete**.
3. Select the confirmation check box, and then click **Delete**.
4. Verify that the LDAP client that you deleted is no longer displayed.

### Related concepts

*LDAP* on page 94
You can use System Manager to configure an LDAP server that centrally maintains user information.

## Editing an LDAP client configuration

You can use System Manager to edit an LDAP client configuration at the cluster level or the storage virtual machine (SVM) level.

### About this task

At the SVM level, you can edit an LDAP client only for a selected SVM.

### Steps

1. To edit an LDAP client configuration:
   - Cluster level: Click ⚙ > **LDAP**.
   - SVM level: Click **SVM** > **SVM Settings** > **LDAP Client**.
2. Select the LDAP client that you want to modify, and then click **Edit**.
3. In the **Edit LDAP Client** dialog box, edit the LDAP client configuration as required.
4. Click **Save and Close**.
5. Verify that the changes that you made to the LDAP client configuration are displayed.

### Related concepts

*LDAP* on page 94

You can use System Manager to configure an LDAP server that centrally maintains user information.

## LDAP Client window

You can use the LDAP Client window to create LDAP clients for user authentication, file access authorization, user search, and mapping services between NFS and CIFS at the storage virtual machine (SVM) level.

### Command buttons

**Add**

Opens the Create LDAP Client dialog box, which enables you to create and configure LDAP clients.

**Edit**

Opens the Edit LDAP Client dialog box, which enables you to edit LDAP client configurations. You can also edit active LDAP clients.

**Delete**

Opens the Delete LDAP Client(s) dialog box, which enables you to delete LDAP client configurations. You can also delete an active LDAP client.

**Refresh**

Updates the information in the window.

### LDAP client list

Displays (in tabular format) details about LDAP clients.

**LDAP Client Configuration**

Displays the name of the LDAP client configuration that you specified.

**Storage Virtual Machine**

Displays the name of the SVM for each LDAP client configuration.

**Schema**

Displays the schema for each LDAP client.

**Minimum Bind Level**

Displays the minimum bind level for each LDAP client.

**Active Directory Domain**

Displays the Active Directory domain for each LDAP client configuration.

**LDAP Servers**

Displays the LDAP server for each LDAP client configuration.

**Preferred Active Directory Servers**

Displays the preferred Active Directory server for each LDAP client configuration.

# LDAP configuration services

You can use System Manager to manage LDAP configurations.

## Editing active LDAP clients

You can use System Manager to associate an active LDAP client with a storage virtual machine (SVM), which enables you to use LDAP as a name service or for name mapping.

### Steps

1. Click **Storage** > **SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **Services** pane, click **LDAP Configuration**.
4. In the **LDAP Configuration** window, click **Edit**.
5. In the **Active LDAP Client** dialog box, select the LDAP client that you want to edit, and perform the following actions:
   - Modify the Active Directory domain servers.
   - Modify the preferred Active Directory servers.
6. Click **OK**.
7. Verify that the changes that you made are updated in the **LDAP Configuration** window.

## Deleting active LDAP clients

You can use System Manager to delete an active LDAP client when you do not want a storage virtual machine (SVM) to be associated with it.

### Steps

1. Click **Storage** > **SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. Click the **SVM Settings** tab.
4. In the **Services** pane, click **LDAP Configuration**.
5. Click **Delete**.
6. Select the confirmation check box, and then click **Delete**.

## LDAP Configuration window

You can use the LDAP Configuration window to edit or delete active LDAP clients at the storage virtual machine (SVM) level.

### Command buttons

**Edit**

Opens the Active LDAP Client dialog box, which enables you to edit the properties of the active LDAP client, such as Active Directory domain servers and preferred Active Directory servers.

**Delete**

Opens the Delete Active LDAP Client dialog box, which enables you to delete the active LDAP client.

**Refresh**

Updates the information in the window.

### LDAP Configuration area

Displays the details about the active LDAP client.

**LDAP client name**

Displays the name of the active LDAP client.

**Active Directory Domain Servers**

Displays the Active Directory domain for the active LDAP client.

**Preferred Active Directory Servers**

Displays the preferred Active Directory server for the active LDAP client.

# Kerberos realm services

You can use System Manager to create and manage Kerberos realm services.

### Related information

*NFS management*

## Creating a Kerberos realm configuration

If you want to use Kerberos authentication for client access, you must configure the storage virtual machine (SVM) to use an existing Kerberos realm. You can use System Manager to create a Kerberos realm configuration, which enables SVMs to use Kerberos security services for NFS.

### Before you begin

- The CIFS license must be installed if CIFS shares are used, and the NFS license must be installed if an LDAP server is used.
- Active Directory (Windows 2003 or Windows 2008) with DES MD5 encryption capability must be available.
- You must have set the time zone and synchronized the time across the cluster by configuring NTP.

  This prevents authentication errors, and ensures that the timestamps in log files are consistent across the cluster.

### About this task

While creating a Kerberos realm, you must set the following attributes in the Create Kerberos Realm wizard:

- Kerberos realm
- KDC IP address and port number

  The default port number is 88.
- Kerberos Key Distribution Center (KDC) vendor
- Administrative server IP address if the KDC vendor is not Microsoft
- Password server IP address
- Active Directory server name and IP address if the KDC vendor is Microsoft

### Steps

1. Click **Storage** > **SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **Services** pane, click **Kerberos Realm**.
4. In the **Kerberos Realm** window, click **Create**.
5. Type or select information as prompted by the wizard.
6. Confirm the details, and then click **Finish** to complete the wizard.

### Related tasks

*Setting the time zone for a cluster* on page 35

### Related information

*NetApp Technical Report 4067: NFS in NetApp ONTAP*

*NetApp Technical Report 4616: NFS Kerberos in ONTAP with Microsoft Active Directory*

*NetApp Technical Report 4835: How to Configure LDAP in ONTAP*

## Editing a Kerberos realm configuration

You can use System Manager to edit a Kerberos realm configuration at the storage virtual machine (SVM) level.

### About this task

You can modify the following attributes by using the Kerberos Realm Edit wizard:

- The KDC IP address and port number
- The IP address of the administrative server if the KDC vendor is not Microsoft
- The IP address of the password server
- The Active Directory server name and IP address if the KDC vendor is Microsoft

### Steps

1. Click **Storage** > **SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **Services** pane, click **Kerberos Realm**.
4. In the **Kerberos Realm** window, select the Kerberos realm configuration that you want to modify, and then click **Edit**.
5. Type or select information as prompted by the wizard.
6. Confirm the details, and then click **Finish** to complete the wizard.

## Deleting Kerberos realm configurations

You can use System Manager to delete a Kerberos realm configuration.

### Steps

1. Click **Storage** > **SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **Services** pane, click **Kerberos Realm**.
4. In the **Kerberos Realm** window, select one or more Kerberos realm configurations that you want to delete, and then click **Delete**.
5. Select the confirmation check box, and then click **Delete**.

## Using Kerberos with NFS for strong security

You can use Kerberos to provide strong authentication between SVMs and NFS clients to provide secure NFS communication. Configuring NFS with Kerberos increases the integrity and security of NFS client communications with the storage system.

## Kerberos authentication for CIFS

With Kerberos authentication, upon connection to your CIFS server, the client negotiates the highest possible security level. However, if the client cannot use Kerberos authentication, Microsoft NTLM or NTLM V2 is used to authenticate with the CIFS server.

## Kerberos Realm window

You can use the Kerberos Realm window to provide authentication between storage virtual machines (SVMs) and NFS clients to ensure secure NFS communication.

### Command buttons

**Create**

Opens the Kerberos Realm Create wizard, which enables you to configure a Kerberos realm to retrieve user information.

**Edit**

Opens the Kerberos Realm Edit wizard, which enables you to edit a Kerberos realm configuration based on the requirement for SVM authentication and authorization.

**Delete**

Opens the Delete Kerberos Realm(s) dialog box, which enables you to delete Kerberos realm configuration.

**Refresh**

Updates the information in the window.

### Kerberos Realm list

Provides details about the Kerberos realms, in tabular format.

**Realm**

Specifies the name of the Kerberos realm.

**KDC Vendor**

Specifies the name of the Kerberos Distribution Center (KDC) vendor.

**KDC IP Address**

Specifies the KDC IP address used by the configuration.

### Details area

The details area displays information such as the KDC IP address and port number, KDC vendor, administrative server IP address and port number, Active Directory server and server IP address of the selected Kerberos realm configuration.

# Kerberos interface services

You can use System Manager to manage Kerberos interface services.

## Editing Kerberos configuration

You can use System Manager to enable Kerberos and to edit a Kerberos configuration that is associated with a storage virtual machine (SVM), which enables the SVM to use Kerberos security services for NFS.

### Before you begin

- You must have at least one Kerberos realm configured at the SVM level.
- You must have a minimum of two data LIFs on the SVM.

  One data LIF is used by the Service Principal Name (SPN) for both the UNIX and CIFS-related Kerberos traffic. The other data LIF is used for accessing non-Kerberos traffic.

  **Note:** A CIFS server is not required for basic NFS Kerberos access. A CIFS server is required for multiprotocol access or when using Active Directory as an LDAP server for name mapping purposes.

### About this task

If you are using Microsoft Active Directory Kerberos, the first 15 characters of any SPNs that are used in the domain must be unique. Microsoft Active Directory has a limitation for SPNs of 15 characters maximum and does not allow duplicate SPNs.

### Steps

1. Click **Storage** > **SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **Services** pane, click **Kerberos Interface**.
4. In the **Kerberos Interface** window, select the interface, and then click **Edit**.
5. In the **Edit Kerberos Configuration** dialog box, make the required changes, and then click **OK**.

### Related information

*NetApp Technical Report 4067: NFS in NetApp ONTAP*
*NetApp Technical Report 4616: NFS Kerberos in ONTAP with Microsoft Active Directory*
*NetApp Technical Report 4835: How to Configure LDAP in ONTAP*

## Kerberos Interface window

You can use the Kerberos Interface window to enable Kerberos and to edit the Kerberos configuration for storage virtual machines (SVMs).

### Command buttons

**Edit**

Opens the Edit Kerberos Configuration dialog box, which you can use to enable Kerberos and to edit the Kerberos configuration associated with the SVM.

**Refresh**

Updates the information in the window.

### Kerberos Interface list

Provides details about the Kerberos configuration.

**Interface Name**

Specifies the logical interfaces associated with the Kerberos configuration for SVMs.

**Service Principal Name**

Specifies the Service Principal Name (SPN) that matches the Kerberos configuration.

**Realm**

Specifies the name of the Kerberos realm associated with the Kerberos configuration.

**Kerberos Status**

Specifies whether Kerberos is enabled.

# DNS/DDNS Services

You can use System Manager to manage DNS/DDNS services.

## Enabling or disabling DDNS

You can use System Manager to enable or disable DDNS on a storage system.

### About this task

- DNS is enabled by default.
- DDNS is disabled by default.
- System Manager does not perform any validation checks for the DNS and DDNS settings.

### Steps

1. Click **Storage** > **SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **Services** pane, click **DNS/DDNS**.
4. In the **DNS/DDNS Services** window, click **Edit**.
5. In the **Edit DNS/DDNS Settings** dialog box, enable DDNS by selecting the **DDNS service** check box.

   You can disable DDNS by clearing the **DDNS service** check box.
6. Click **OK**.

### Related reference

*DNS/DDNS Services window* on page 293

The DNS/DDNS Services window enables you to view and edit the current DNS and DDNS settings for your system.

## Editing DNS and DDNS settings

You can maintain host information centrally by using DNS. You can use System Manager to add or modify the DNS domain name of your storage system. You can also enable DDNS on your storage system to update the name server automatically in the DNS server.

### Before you begin

You must have set up a CIFS server or an Active Directory account for the storage virtual machine (SVM) for secure DDNS to work.

### About this task

System Manager does not perform any validation checks for the DNS and DDNS settings.

### Steps

1. Click **Storage** > **SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **Services** pane, click **DNS/DDNS**.
4. Click **Edit**.
5. In the **DNS Domains and Name Servers** area, add or modify the DNS domain names and the IP addresses.
6. Select the **DDNS service** check box to enable DDNS.
   a. Select the **Enable Secure DDNS** check box to enable secure DDNS.
   b. Specify the fully qualified domain name (FQDN) and the time to live value for the DDNS service.

      By default, time to live is set to 24 hours and FQDN is set to `SVM name. domain name`.
7. Click **OK** to save the changes that you made.

### Related reference

*DNS/DDNS Services window* on page 293

The DNS/DDNS Services window enables you to view and edit the current DNS and DDNS settings for your system.

## DNS/DDNS Services window

The DNS/DDNS Services window enables you to view and edit the current DNS and DDNS settings for your system.

### Command buttons

**Edit**

Opens the Edit DNS/DDNS Settings dialog box, which you can use to add or modify DNS or DDNS details. You can also enable or disable DDNS.

**Refresh**

Updates the information in the window.

### Related tasks

*Enabling or disabling DDNS* on page 292

You can use System Manager to enable or disable DDNS on a storage system.

*Editing DNS and DDNS settings* on page 292

You can maintain host information centrally by using DNS. You can use System Manager to add or modify the DNS domain name of your storage system. You can also enable DDNS on your storage system to update the name server automatically in the DNS server.

# Users

You can use System Manager to create and manage storage virtual machine (SVM) user accounts.

## Adding SVM user accounts

You can use System Manager to add a storage virtual machine (SVM) user account and to specify a user login method for accessing the storage system.

### Steps

1. Click **Storage** > **SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **SVM User Details** pane, click **Users**.
4. Click **Add**.
5. Type the user name and password that the user uses to connect to the storage system, and confirm the password.
6. Add one or more user login methods, and then click **Add**.

## Changing the password for SVM user accounts

You can use System Manager to reset the password for a storage virtual machine (SVM) user account.

### Steps

1. Click **Storage** > **SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **SVM User Details** pane, click **Users**.

4. Select the user account for which you want to modify the password, and then click **Reset Password**.
5. In the **Reset Password** dialog box, type the new password, confirm the new password, and then click **Change**.

## Editing SVM user accounts

You can use System Manager to edit a storage virtual machine (SVM) user account by modifying the user login methods for accessing the storage system.

### Steps

1. Click **Storage** > **SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **SVM User Details** pane, click **Users**.
4. Select the user account that you want to edit, and then click **Edit**.
5. Modify one or more user login methods, and then click **Modify**.

## Locking or unlocking SVM user accounts

You can use System Manager to lock or unlock storage virtual machine (SVM) user accounts.

### Steps

1. Click **Storage** > **SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **SVM User Details** pane, click **Users**.
4. In the **Users** window, select the user account for which you want to modify the account status, and then click either **Lock** or **Unlock**, as required.

## Users window

You can use the Users window to manage user accounts, to reset the password of a user, and to view information about all of the user accounts.

### Command buttons

**Add**

Opens the Add User dialog box, which enables you to add user accounts.

**Edit**

Opens the Modify User dialog box, which enables you to modify user login methods.

> **Note:** It is a best practice to use a single role for all of the access and authentication methods of a user account.

**Delete**

Enables you to delete a selected user account.

**Change Password**

Opens the Change Password dialog box, which enables you to reset a selected user's password.

**Lock**

Locks the user account.

**Refresh**

Updates the information in the window.

### Users list

The area below the users list displays detailed information about the selected user.

**User**

Displays the name of the user account.

**Account Locked**

Displays whether the user account is locked.

### User Login Methods area

**Application**

Displays the access method that a user can use to access the storage system. The supported access methods include the following:

- System console (console)
- HTTP(S) (http)
- ONTAP API (ontapi)
- Service Processor (service-processor)
- SSH (ssh)

**Authentication**

Displays the default supported authentication method, which is "password".

**Role**

Displays the role of a selected user.

# Roles

You can use System Manager to create and manage roles.

### Related information

*Administrator authentication and RBAC*

## Adding roles

You can use System Manager to add an access-control role and to specify the command or command directory that the users of the role can access. You can also control the level of access the role has to the command or command directory, and you can specify a query that applies to the command or command directory.

### Steps

1. Click **Storage** > **SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **SVM User Details** pane, click **Roles**.
4. Click **Add**.
5. In the **Add Role** dialog box, specify the role name, and then add the role attributes.
6. Click **Add**.

## Editing roles

You can use System Manager to modify the access of an access-control role to a command or command directory and to restrict a user's access to only a specified set of commands.

### Steps

1. Click **Storage** > **SVMs**.

2. Select the SVM, and then click **SVM Settings**.
3. In the **SVM User Details** pane, click **Roles**.
4. Select the role that you want to modify, and then click **Edit**.
5. Modify the role attributes, and then click **Modify**.

## Roles window

You can use the Roles window to manage the roles that are associated with user accounts.

### Command buttons

**Add**

Opens the Add Role dialog box, which enables you to create an access-control role and specify the command or command directory that the role's users can access.

**Edit**

Opens the Edit Role dialog box, which enables you to add or modify role attributes.

**Refresh**

Updates the information in the window.

### Roles list

The roles list provides a list of roles that are available to be assigned to users.

### Role Attributes area

The details area displays the role attributes, such as the command or command directory that the selected role can access, the access level, and the query that applies to the command or command directory.

# UNIX

You can use System Manager to maintain a list of local UNIX users and groups for each storage virtual machine (SVM).

## UNIX window

You can use the UNIX window to maintain a list of local UNIX users and groups for each storage virtual machine (SVM). You can use local UNIX users and groups for authentication and name mappings.

### Groups tab

You can use the Groups tab to add, edit, or delete UNIX groups that are local to an SVM.

### Command buttons

**Add Group**

Opens the Add Group dialog box, which enables you to create UNIX groups that are local to SVMs. Local UNIX groups are used with local UNIX users.

**Edit**

Opens the Edit Group dialog box, which enables you to edit a group ID.

**Delete**

Deletes the selected group.

**Refresh**

Updates the information in the window.

### Groups list

**Group Name**

Displays the name of the group.

**Group ID**

Displays the ID of the group.

### Users tab

You can use the **Users** tab to add, edit, and delete UNIX users that are local to SVMs.

### Command buttons

**Add User**

Opens the Add User dialog box, which enables you to create UNIX users that are local to SVMs.

**Edit**

Opens the Edit User dialog box, which enables you to edit the User ID, UNIX group to which the user belongs, and the full name of the user.

**Delete**

Deletes the selected user.

**Refresh**

Updates the information in the window.

### Users list

**User Name**

Displays the name of the user.

**User ID**

Displays the ID of the user.

**Full Name**

Displays the full name of the user.

**Primary Group ID**

Displays the ID of the group to which the user belongs.

**Primary Group Name**

Displays the name of the group to which the user belongs.

# Windows

You can use System Manager to create and manage Windows groups and user accounts.

### Related information

*SMB/CIFS management*

## Creating a local Windows group

You can use System Manager to create local Windows groups that can be used for authorizing access to the data contained in the storage virtual machine (SVM) over an SMB connection. You can also assign the privileges that define the user rights or capabilities that a member of the group has when performing administrative activities.

### Before you begin

CIFS server must be configured for the SVM.

### About this task

- You can specify a group name with or without the local domain name.
  The local domain is the name of the CIFS server for the SVM. For example, if the CIFS server name of the SVM is "CIFS_SERVER" and you want to create an "engineering" group, you can specify either "engineering" or "CIFS_SERVER\engineering" as the group name.
  The following rules apply when using a local domain as part of the group name:
  - You can specify only the local domain name for the SVM to which the group is applied.
    For example, if the local CIFS server name is "CIFS_SERVER", you cannot specify "CORP_SERVER\group1" as the group name.
  - You cannot use "BUILTIN" as a local domain in the group name.
    For example, you cannot create a group with "BUILTIN\group1" as the name.
  - You cannot use an Active Directory domain as a local domain in the group name.
    For example, you cannot create a group named "AD_DOM\group1", where "AD_DOM" is the name of an Active Directory domain.
- You cannot use a group name that already exists.
- The group name that you specify must meet the following requirements:
  - Must not exceed 256 characters
  - Must not end in a period
  - Must not include commas
  - Must not include any of the following printable characters: " / \ [ ] : | < > + = ; ? * @
  - Must not include characters in the ASCII range 1 through 31, which are non-printable

### Steps

1. Click **Storage** > **SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **Host Users and Groups** pane, click **Windows**.
4. In the **Groups** tab, click **Create**.
5. In the **Create Group** dialog box, specify a name for the group and a description that helps you to identify the new group.
6. Assign a set of privileges to the group.
   You can select the privileges from the predefined set of supported privileges.
7. Click **Add** to add users to the group.
8. In the **Add Members to Group** dialog box, perform one of the following actions:
   - Specify the Active Directory user or Active Directory group to be added to a particular local group.
   - Select the users from the list of available local users in the SVM.
   - Click **OK**.
9. Click **Create**.

### Result

The local Windows group is created and is listed in the Groups window.

#### Related reference

*Windows window* on page 305

You can use the Windows window to maintain a list of local Windows users and groups for each storage virtual machine (SVM) on the cluster. You can use the local Windows users and groups for authentication and name mappings.

## Editing local Windows group properties

You can manage local group memberships by adding and removing a local user, an Active Directory user, or an Active Directory group by using System Manager. You can modify the privileges that are assigned to a group and the description of a group to easily identify the group.

### About this task

You must keep the following in mind when adding members to or removing members from a local Windows group:

- You cannot add users to or remove users from the special *Everyone* group.
- You cannot add a local Windows group to another local Windows group.

### Steps

1. Click **Storage** > **SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **Host Users and Groups** pane, click **Windows**.
4. In the **Groups** tab, click **Edit**.
5. Specify a name for the group and a description to identify the new group.
6. Assign a set of privileges to the group.

   You can select the privileges from the predefined set of supported privileges.
7. Click **Add** to add users to the group.
8. In the **Add Members** window, perform one of the following actions:
   - Specify the Active Directory user or Active Directory group to be added to a particular local group.
   - Select the users from the list of available local users in the storage virtual machine (SVM).
9. Click **Edit**.

### Result

The local Windows group settings are modified, and the changes are displayed in the **Groups** tab.

### Related reference

*Windows window* on page 305
You can use the Windows window to maintain a list of local Windows users and groups for each storage virtual machine (SVM) on the cluster. You can use the local Windows users and groups for authentication and name mappings.

## Adding user accounts to a Windows local group

You can add a local user, an Active Directory user, or an Active Directory group (if you want users to have the privileges that are associated with that group) to a Windows local group by using System Manager.

### Before you begin

- The group must exist before you can add a user to the group.
- The user must exist before you can add the user to a group.

### About this task

You must keep the following in mind when adding members to a local Windows group:

- You cannot add users to the special *Everyone* group.
- You cannot add a local Windows group to another local Windows group.
- You cannot add a user account that contains a space in the user name by using System Manager.

  You can either rename the user account or add the user account by using the command-line interface (CLI).

### Steps

1. Click **Storage** > **SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **Host Users and Groups** pane, click **Windows**.
4. In the **Groups** tab, select the group to which you want to add a user, and then click **Add Members**.
5. In the **Add Members** window, perform one of the following actions:
   - Specify the Active Directory user or Active Directory group to be added to a particular local group.
   - Select the users from the list of available local users in the storage virtual machine (SVM).
6. Click **OK**.

### Result

The user that you added is listed in the Users tab of the **Groups** tab.

#### Related reference

*Windows window* on page 305

You can use the Windows window to maintain a list of local Windows users and groups for each storage virtual machine (SVM) on the cluster. You can use the local Windows users and groups for authentication and name mappings.

## Renaming a local Windows group

You can use System Manager to rename a local Windows group to identify the group more easily.

### About this task

- The new group name must be created in the same domain as the old group name.
- The group name must meet the following requirements:
  - Must not exceed 256 characters
  - Must not end in a period
  - Must not include commas
  - Must not include any of the following printable characters: " / \ [ ] : | < > + = ; ? * @
  - Must not include characters in the ASCII range 1 through 31, which are non-printable

### Steps

1. Click **Storage** > **SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **Host Users and Groups** pane, click **Windows**.
4. In the **Groups** tab, select the group that you want to rename, and then click **Rename**.
5. In the **Rename Group** window, specify a new name for the group.

### Result

The local group name is changed, and the group is listed with the new name in the Groups window.

### Related reference

*Windows window* on page 305

You can use the Windows window to maintain a list of local Windows users and groups for each storage virtual machine (SVM) on the cluster. You can use the local Windows users and groups for authentication and name mappings.

## Deleting a local Windows group

You can use System Manager to delete a local Windows group from a storage virtual machine (SVM) if the group is no longer required for determining access rights to the data contained on the SVM or for assigning SVM user rights (privileges) to group members.

### About this task

- Removing a local group removes the membership records of the group.
- The file system is not altered.
  Windows Security Descriptors on files and directories that refer to this group are not adjusted.
- The special "Everyone" group cannot be deleted.
- Built-in groups such as BUILTIN\Administrators and BUILTIN\Users cannot be deleted.

### Steps

1. Click **Storage** > **SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **Host Users and Groups** pane, click **Windows**.
4. In the **Groups** tab, select the group that you want to delete, and then click **Delete**.
5. Click **Delete**.

### Result

The local group is deleted along with its membership records.

### Related reference

*Windows window* on page 305

You can use the Windows window to maintain a list of local Windows users and groups for each storage virtual machine (SVM) on the cluster. You can use the local Windows users and groups for authentication and name mappings.

## Creating a local Windows user account

You can use System Manager to create a local Windows user account that can be used to authorize access to the data contained in the storage virtual machine (SVM) over an SMB connection. You can also use local Windows user accounts for authentication when creating a CIFS session.

### Before you begin

- The CIFS server must be configured for the SVM.

### About this task

A local Windows user name must meet the following requirements:

- Must not exceed 20 characters
- Must not end in a period
- Must not include commas
- Must not include any of the following printable characters: " / \ [ ] : | < > + = ; ? * @
- Must not include characters in the ASCII range 1 through 31, which are non-printable

The password must meet the following criteria:

- Must be at least six characters in length
- Must not contain the user account name
- Must contain characters from at least three of the following four categories:
  - English uppercase characters (A through Z)
  - English lowercase characters (a through z)
  - Base 10 digits (0 through 9)
  - Special characters: ~ ! @ # 0 ^ & * _ - + = ` \ | ( ) [ ] : ; " ' < > , . ? /

### Steps

1. Click **Storage** > **SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **Host Users and Groups** pane, click **Windows**.
4. In the **Users** tab, click **Create**.
5. Specify a name for the local user.
6. Specify the full name of the local user and a description that helps you to identify this new user.
7. Enter a password for the local user, and confirm the password.

   The password must meet the password requirements.
8. Click **Add** to assign group memberships to the user.
9. In the **Add Groups** window, select the groups from the list of available groups in the SVM.
10. Select **Disable this account** to disable this account after the user is created.
11. Click **Create**.

### Result

The local Windows user account is created and is assigned membership to the selected groups. The user account is listed in the **Users** tab.

### Related reference

*Windows window* on page 305

You can use the Windows window to maintain a list of local Windows users and groups for each storage virtual machine (SVM) on the cluster. You can use the local Windows users and groups for authentication and name mappings.

## Editing the local Windows user properties

You can use System Manager to modify a local Windows user account if you want to change an existing user's full name or description, or if you want to enable or disable the user account. You can also modify the group memberships that are assigned to the user account.

### Steps

1. Click **Storage** > **SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **Host Users and Groups** pane, click **Windows**.
4. In the **Users** tab, click **Edit**.
5. In the **Modify User** window, make the required changes.
6. Click **Modify**.

### Result

The attributes of the local Windows user account are modified and are displayed in the **Users** tab.

### Related reference

*Windows window* on page 305

You can use the Windows window to maintain a list of local Windows users and groups for each storage virtual machine (SVM) on the cluster. You can use the local Windows users and groups for authentication and name mappings.

## Assigning group memberships to a user account

You can use System Manager to assign group membership to a user account if you want a user to have the privileges that are associated with a particular group.

### Before you begin

- The group must exist before you can add a user to the group.
- The user must exist before you can add the user to a group.

### About this task

You cannot add users to the special *Everyone* group.

### Steps

1. Click **Storage** > **SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **Host Users and Groups** pane, click **Windows**.
4. In the **Users** tab, select the user account to which you want to assign group memberships, and then click **Add to Group**.
5. In the **Add Groups** window, select the groups to which you want to add the user account.
6. Click **OK**.

### Result

The user account is assigned membership to all of the selected groups, and the user has the privileges that are associated with these groups.

### Related reference

*Windows window* on page 305

You can use the Windows window to maintain a list of local Windows users and groups for each storage virtual machine (SVM) on the cluster. You can use the local Windows users and groups for authentication and name mappings.

## Renaming a local Windows user

You can use System Manager to rename a local Windows user account to identify the local user more easily.

### About this task

- The new user name must be created in the same domain as the previous user name.
- The user name that you specify must meet the following requirements:
  - Must not exceed 20 characters
  - Must not end in a period
  - Must not include commas
  - Must not include any of the following printable characters: " / \ [ ] : | < > + = ; ? * @
  - Must not include characters in the ASCII range 1 through 31, which are non-printable

### Steps

1. Click **Storage** > **SVMs**.

2. Select the SVM, and then click **SVM Settings**.
3. In the **Host Users and Groups** pane, click **Windows**.
4. In the **Users** tab, select the user that you want to rename, and then click **Rename**.
5. In the **Rename User** window, specify a new name for the user.
6. Confirm the new name, and then click **Rename**.

### Result

The user name is changed, and the new name is listed in the **Users** tab.

### Related reference

*Windows window* on page 305

You can use the Windows window to maintain a list of local Windows users and groups for each storage virtual machine (SVM) on the cluster. You can use the local Windows users and groups for authentication and name mappings.

## Resetting the password of a Windows local user

You can use System Manager to reset the password of a Windows local user. For example, you might want to reset the password if the current password is compromised or if the user has forgotten the password.

### About this task

The password that you set must meet the following criteria:

- Must be at least six characters in length
- Must not contain the user account name
- Must contain characters from at least three of the following four categories:
  - English uppercase characters (A through Z)
  - English lowercase characters (a through z)
  - Base 10 digits (0 through 9)
  - Special characters: ~ ! @ # 0 ^ & * _ - + = ` \ | ( ) [ ] : ; " ' < > , . ? /

### Steps

1. Click **Storage** > **SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **Host Users and Groups** pane, click **Windows**.
4. In the **Users** tab, select the user whose password you want to reset, and then click **Set Password**.
5. In the **Reset Password** dialog box, set a new password for the user.
6. Confirm the new password, and then click **Reset**.

### Related reference

*Windows window* on page 305

You can use the Windows window to maintain a list of local Windows users and groups for each storage virtual machine (SVM) on the cluster. You can use the local Windows users and groups for authentication and name mappings.

## Deleting a local Windows user account

You can use System Manager to delete a local Windows user account from a storage virtual machine (SVM) if the user account is no longer required for local CIFS authentication to the CIFS server of the SVM or for determining access rights to the data contained in the SVM.

### About this task

- Standard users such as Administrator cannot be deleted.
- ONTAP removes references to the deleted local user from the local-group database, from the local-user-membership, and from the user-rights database.

### Steps

1. Click **Storage** > **SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **Host Users and Groups** pane, click **Windows**.
4. In the **Users** tab, select the user account that you want to delete, and then click **Delete**.
5. Click **Delete**.

### Result

The local user account is deleted along with its group membership entries.

### Related reference

*Windows window* on page 305

You can use the Windows window to maintain a list of local Windows users and groups for each storage virtual machine (SVM) on the cluster. You can use the local Windows users and groups for authentication and name mappings.

## Windows window

You can use the Windows window to maintain a list of local Windows users and groups for each storage virtual machine (SVM) on the cluster. You can use the local Windows users and groups for authentication and name mappings.

- *Users tab*
- *Groups tab*

### Users tab

You can use the Users tab to view the Windows users that are local to an SVM.

### Command buttons

**Create**

Opens the Create User dialog box, which enables you to create a local Windows user account that can be used to authorize access to data contained in the SVM over an SMB connection.

**Edit**

Opens the Edit User dialog box, which enables you to edit local Windows user properties, such as group memberships and the full name. You can also enable or disable the user account.

**Delete**

Opens the Delete User dialog box, which enables you to delete a local Windows user account from an SVM if it is no longer required.

**Add to Group**

Opens the Add Groups dialog box, which enables you to assign group membership to a user account if you want the user to have privileges associated with that group.

**Set Password**

Opens the Reset Password dialog box, which enables you to reset the password of a Windows local user. For example, you might want to reset the password if the password is compromised or if the user has forgotten the password.

**Rename**

Opens the Rename User dialog box, which enables you to rename a local Windows user account to more easily identify it.

**Refresh**

Updates the information in the window.

### Users list

**Name**

Displays the name of the local user.

**Full Name**

Displays the full name of the local user.

**Account Disabled**

Displays whether the local user account is enabled or disabled.

**Description**

Displays the description for this local user.

### Users Details Area

**Group**

Displays the list of groups in which the user is a member.

### Groups tab

You can use the Groups tab to add, edit, or delete Windows groups that are local to an SVM.

### Command buttons

**Create**

Opens the Create Group dialog box, which enables you to create local Windows groups that can be used for authorizing access to data contained in SVMs over an SMB connection.

**Edit**

Opens the Edit Group dialog box, which enables you to edit the local Windows group properties, such as privileges assigned to the group and the description of the group.

**Delete**

Opens the Delete Group dialog box, which enables you to delete a local Windows group from an SVM if it is no longer required.

**Add Members**

Opens the Add Members dialog box, which enables you to add local or Active Directory users, or Active Directory groups to the local Windows group.

**Rename**

Opens the Rename Group dialog box, which enables you to rename a local Windows group to more easily identify it.

**Refresh**

Updates the information in the window.

### Groups list

**Name**

Displays the name of the local group.

**Description**

Displays the description for this local group.

### Groups Details Area

**Privileges**

Displays the list of privileges associated with the selected group.

**Users**

Displays the list of local users associated with the selected group.

### Related tasks

*Creating a local Windows group* on page 297
You can use System Manager to create local Windows groups that can be used for authorizing access to the data contained in the storage virtual machine (SVM) over an SMB connection. You can also assign the privileges that define the user rights or capabilities that a member of the group has when performing administrative activities.

*Editing local Windows group properties* on page 299
You can manage local group memberships by adding and removing a local user, an Active Directory user, or an Active Directory group by using System Manager. You can modify the privileges that are assigned to a group and the description of a group to easily identify the group.

*Adding user accounts to a Windows local group* on page 299
You can add a local user, an Active Directory user, or an Active Directory group (if you want users to have the privileges that are associated with that group) to a Windows local group by using System Manager.

*Renaming a local Windows group* on page 300
You can use System Manager to rename a local Windows group to identify the group more easily.

*Deleting a local Windows group* on page 301
You can use System Manager to delete a local Windows group from a storage virtual machine (SVM) if the group is no longer required for determining access rights to the data contained on the SVM or for assigning SVM user rights (privileges) to group members.

*Creating a local Windows user account* on page 301
You can use System Manager to create a local Windows user account that can be used to authorize access to the data contained in the storage virtual machine (SVM) over an SMB connection. You can also use local Windows user accounts for authentication when creating a CIFS session.

*Editing the local Windows user properties* on page 302
You can use System Manager to modify a local Windows user account if you want to change an existing user's full name or description, or if you want to enable or disable the user account. You can also modify the group memberships that are assigned to the user account.

*Assigning group memberships to a user account* on page 303

You can use System Manager to assign group membership to a user account if you want a user to have the privileges that are associated with a particular group.

*Renaming a local Windows user* on page 303
You can use System Manager to rename a local Windows user account to identify the local user more easily.

*Resetting the password of a Windows local user* on page 304
You can use System Manager to reset the password of a Windows local user. For example, you might want to reset the password if the current password is compromised or if the user has forgotten the password.

*Deleting a local Windows user account* on page 305
You can use System Manager to delete a local Windows user account from a storage virtual machine (SVM) if the user account is no longer required for local CIFS authentication to the CIFS server of the SVM or for determining access rights to the data contained in the SVM.

# Name mapping

You can use System Manager to specify name mapping entries to map users from different platforms.

**Related information**

*SMB/CIFS management*

## Name mapping conversion rules

An ONTAP system keeps a set of conversion rules for each SVM. Each rule consists of two pieces: a *pattern* and a *replacement*. Conversions start at the beginning of the appropriate list and perform a substitution based on the first matching rule. The pattern is a UNIX-style regular expression. The replacement is a string containing escape sequences representing subexpressions from the pattern, as in the UNIX `sed` program.

## Name Mapping window

You can use the Name Mapping window to specify the name mapping entries to map users from different platforms.

### Name Mappings

You can create and use name mappings to map your UNIX users to Windows users, Windows users to UNIX users, or Kerberos users to UNIX users.

### Command buttons

**Add**

Opens the Add Name Mapping Entry dialog box, which enables you to create a name mapping on storage virtual machines (SVMs).

**Edit**

Opens the Edit Name Mapping Entry dialog box, which enables you to edit a name mapping on SVMs.

**Delete**

Opens the Delete Name Mapping Entries dialog box, which enables you to delete a name mapping entry.

**Swap**

Opens the Swap Name Mapping Entries dialog box, which enables you to interchange positions of the two selected name mapping entries.

**Refresh**

Updates the information in the window.

### Name mappings list

**Position**

Specifies the name mapping's position in the priority list. Name mappings are applied in the order in which they occur in the priority list.

**Pattern**

Specifies the user name pattern that must be matched.

**Replacement**

Specifies the replacement pattern for the user name.

**Direction**

Specifies the direction of the name mapping. Possible values are krb_unix for a Kerberos-to-UNIX name mapping, win_unix for a Windows-to-UNIX name mapping, and unix_win for a UNIX-to-Windows name mapping.

### Command buttons

**Add**

Opens the Add Group Mapping Entry dialog box, which enables you to create a group mapping on SVMs.

**Edit**

Opens the Edit Group Mapping Entry dialog box, which enables you to edit the group mapping on SVMs.

**Delete**

Opens the Delete Group Mapping Entries dialog box, which enables you to delete a group mapping entry.

**Swap**

Opens the Swap Group Mapping Entries dialog box, which enables you to interchange positions of the two selected group mapping entries.

**Refresh**

Updates the information in the window.

# Managing data protection

You can use System Manager to protect your data by creating and managing mirror relationships, vault relationships, and mirror and vault relationships. You can also create and manage the Snapshot policies and schedules.

## Mirror relationships

You can use System Manager to create and manage mirror relationships by using the mirror policy.

### Creating a mirror relationship from a destination SVM

You can use System Manager to create a mirror relationship from the destination storage virtual machine (SVM), and to assign a policy and schedule to the mirror relationship. The mirror copy enables quick availability of data if the data on the source volume is corrupted or lost.

#### Before you begin

- The source cluster must be running ONTAP 8.2.2 or later.
- The SnapMirror license must be enabled on the source cluster and destination cluster.

  **Note:** For some platforms, it is not mandatory for the source cluster to have the SnapMirror license enabled if the destination cluster has the SnapMirror license and Data Protection Optimization (DPO) license enabled.
- While mirroring a volume, if you select a SnapLock volume as the source, then the SnapMirror license and SnapLock license must be installed on the destination cluster.
- The source cluster and destination cluster must be in a healthy peer relationship.
- The destination SVM must have space available.
- A source volume of type read/write (rw) must exist.
- The FlexVol volumes must be online and must be of type read/write.
- The SnapLock aggregate type must be of the same type.
- If you are connecting from a cluster running ONTAP 9.2 or earlier to a remote cluster on which Security Assertion Markup Language (SAML) authentication is enabled, password-based authentication must be enabled on the remote cluster.

#### About this task

- System Manager does not support a cascade relationship.

  For example, a destination volume in a relationship cannot be the source volume in another relationship.
- You cannot create a mirror relationship between a sync-source SVM and a sync-destination SVM in a MetroCluster configuration.
- You can create a mirror relationship between sync-source SVMs in a MetroCluster configuration.
- You can create a mirror relationship from a volume on a sync-source SVM to a volume on a data-serving SVM.
- You can create a mirror relationship from a volume on a data-serving SVM to a data protection (DP) volume on a sync-source SVM.
- You can create a mirror relationship between SnapLock volumes of the same type only.

  For example, if the source volume is a SnapLock Enterprise volume, then the destination volume must also be a SnapLock Enterprise volume. You must ensure that the destination SVM has aggregates of the same SnapLock type available.
- The destination volume that is created for a mirror relationship is not thin provisioned.
- A maximum of 25 volumes can be protected in one selection.

### Steps

1. Click **Protection** > **Relationships**.
2. In the **Relationships** window, click **Create**.
3. In the **Browse SVM** dialog box, select an SVM for the destination volume.
4. In the **Create Protection Relationship** dialog box, select **Mirror** from the **Relationship Type** drop-down list.
5. Specify the cluster, the SVM, and the source volume.

   If the specified cluster is running a version of ONTAP software earlier than ONTAP 9.3, then only peered SVMs are listed. If the specified cluster is running ONTAP 9.3 or later, peered SVMs and permitted SVMs are listed.
6. For FlexVol volumes, specify a volume name suffix.

   The volume name suffix is appended to the source volume names to generate the destination volume names.
7. Optional: Click **Browse**, and then change the mirror policy.
8. Select a schedule for the relationship from the list of existing schedules.
9. Optional: Select **Initialize Relationship** to initialize the mirror relationship.
10. Enable FabricPool-enabled aggregates, and then select an appropriate tiering policy.
11. Click **Create**.

### Result

If you chose to create a destination volume, a destination volume of type *dp* is created, with the language attribute set to match the language attribute of the source volume.

A mirror relationship is created between the source volume and the destination volume. The base Snapshot copy is transferred to the destination volume if you have opted to initialize the relationship.

### Related reference

*Protection window* on page 336
You can use the Protection window to create and manage mirror relationships, vault relationships, and mirror vault relationships and to display details about these relationships. The Protection window does not display load-sharing (LS) relationships and transition (TDP) relationships.

## Deleting mirror relationships

You can delete a mirror relationship and permanently end the mirror relationship between the source and destination volumes. When a mirror relationship is deleted, the base Snapshot copy on the source volume is deleted.

### About this task

It is a best practice to break the mirror relationship before deleting the relationship.

### Steps

1. Click **Protection** > **Relationships**.
2. Select the mirror relationship that you want to delete and click **Delete**.
3. Select the confirmation check boxes to delete the mirror relationship and to release the base Snapshot copies, and then click **Delete**.

### Result

The relationship is deleted and the base Snapshot copy on the source volume is deleted.

### Related reference

*Protection window* on page 336

You can use the Protection window to create and manage mirror relationships, vault relationships, and mirror vault relationships and to display details about these relationships. The Protection window does not display load-sharing (LS) relationships and transition (TDP) relationships.

## Editing mirror relationships

You can use System Manager to edit a mirror relationship either by selecting an existing policy or schedule in the cluster, or by creating a policy or schedule.

### About this task

- You cannot edit a mirror relationship that is created between a volume in Data ONTAP 8.2.1 and a volume in ONTAP 8.3 or later.
- You cannot edit the parameters of an existing policy or schedule.
- You can modify the relationship type of a version-flexible mirror relationship, vault relationship, or mirror and vault relationship by modifying the policy type.

### Steps

1. Click **Protection** > **Relationships**.
2. Select the mirror relationship for which you want to modify the policy or schedule, and then click **Edit**.
3. In the **Edit Relationship** dialog box, select an existing policy or create a policy:

| If you want to… | Do the following… |
|---|---|
| Select an existing policy | Click **Browse**, and then select an existing policy. |
| Create a policy | **a.** Click **Create Policy**.<br>**b.** Specify a name for the policy.<br>**c.** Set the priority for scheduled transfers.<br>Low indicates that the transfer has the least priority and is usually scheduled after normal priority transfers. By default, the priority is set to Normal.<br>**d.** Select the **Transfer All Source Snapshot Copies** check box to include the "all_source_snapshots" rule to the mirror policy, which enables you to back up all of the Snapshot copies from the source volume.<br>**e.** Select the **Enable Network Compression** check box to compress the data that is being transferred.<br>**f.** Click **Create**. |

4. Specify a schedule for the relationship:

| If… | Do the following… |
|---|---|
| You want to assign an existing schedule | From the list of schedules, select an existing schedule. |
| You want to create a schedule | **a.** Click **Create Schedule**.<br>**b.** Specify a name for the schedule.<br>**c.** Select either **Basic** or **Advanced**.<br>• Basic specifies only the day of the week, time, and the transfer interval.<br>• Advanced creates a cron-style schedule.<br>**d.** Click **Create**. |

| If… | Do the following… |
|---|---|
| You do not want to assign a schedule | Select **None**. |

5. Click **OK** to save the changes.

### Related reference

*Protection window* on page 336
You can use the Protection window to create and manage mirror relationships, vault relationships, and mirror vault relationships and to display details about these relationships. The Protection window does not display load-sharing (LS) relationships and transition (TDP) relationships.

## Initializing mirror relationships

When you start a mirror relationship for the first time, you have to initialize the relationship. Initializing a relationship consists of a complete baseline transfer of data from the source volume to the destination. You can use System Manager to initialize a mirror relationship if you have not already initialized the relationship while creating it.

### Steps

1. Click **Protection** > **Relationships**.
2. Select the mirror relationship that you want to initialize.
3. Click **Operations** > **Initialize**.
4. Select the confirmation check box and click **Initialize**.
5. Verify the status of the mirror relationship in the **Protection** window.

### Result

A Snapshot copy is created and transferred to the destination. This Snapshot copy is used as a baseline for subsequent incremental Snapshot copies.

### Related reference

*Protection window* on page 336
You can use the Protection window to create and manage mirror relationships, vault relationships, and mirror vault relationships and to display details about these relationships. The Protection window does not display load-sharing (LS) relationships and transition (TDP) relationships.

## Updating mirror relationships

You can initiate an unscheduled mirror update of the destination. You might have to perform a manual update to prevent data loss due to an upcoming power outage, scheduled maintenance, or data migration.

### Before you begin

The mirror relationship must be in a Snapmirrored state.

### Steps

1. Click **Protection** > **Relationships**.
2. Select the mirror relationship for which you want to update the data, and click **Operations** > **Update**.
3. Choose one of the following options:
   - Select **On demand** to perform an incremental transfer from the recent common Snapshot copy between the source and destination volumes.
   - Select **Select Snapshot copy** and specify the Snapshot copy that you want to transfer.

4. Optional: Select **Limit transfer bandwidth to** to limit the network bandwidth used for transfers and specify the maximum transfer speed.
5. Click **Update**.
6. Verify the transfer status in the **Details** tab.

### Related reference

*Protection window* on page 336

You can use the Protection window to create and manage mirror relationships, vault relationships, and mirror vault relationships and to display details about these relationships. The Protection window does not display load-sharing (LS) relationships and transition (TDP) relationships.

## Quiescing mirror relationships

Use System Manager to quiesce a mirror destination to stabilize the destination before creating a Snapshot copy. The quiesce operation enables active mirror transfers to finish and disables future transfers for the mirroring relationship.

### About this task

You can quiesce only mirror relationships that are in the Snapmirrored state.

### Steps

1. Click **Protection** > **Relationships**.
2. Select the mirror relationship that you want to quiesce.
3. Click **Operations** > **Quiesce**.
4. Select the confirmation check box and click **Quiesce**.

### Related reference

*Protection window* on page 336

You can use the Protection window to create and manage mirror relationships, vault relationships, and mirror vault relationships and to display details about these relationships. The Protection window does not display load-sharing (LS) relationships and transition (TDP) relationships.

## Resuming mirror relationships

You can resume a quiesced mirror relationship. When you resume the relationship, normal data transfer to the mirror destination is resumed and all the mirror activities are restarted.

### About this task

If you have quiesced a broken mirror relationship from the command-line interface (CLI), you cannot resume the relationship from System Manager. You must use the CLI to resume the relationship.

### Steps

1. Click **Protection** > **Relationships**.
2. Select the mirror relationship that you want to resume.
3. Click **Operations** > **Resume**.
4. Select the confirmation check box and click **Resume**.

### Result

Data transfer to the mirror destination is resumed for the selected mirror relationship.

### Related reference

*Protection window* on page 336

You can use the Protection window to create and manage mirror relationships, vault relationships, and mirror vault relationships and to display details about these relationships. The Protection window does not display load-sharing (LS) relationships and transition (TDP) relationships.

## Breaking SnapMirror relationships

You must break a mirror relationship if a mirror source becomes unavailable and you want client applications to be able to access the data from the mirror destination. After the mirror relationship is broken, the destination volume type changes from DP to RW.

### Before you begin

- The SnapMirror destination must be in the quiesced state or idle state.
- The destination volume must be mounted on the destination storage virtual machine (SVM) namespace.

### About this task

- You can use the destination volume to serve data while you repair or replace the source, update the source, and reestablish the original configuration of the systems.
- You can break SnapMirror relationships between ONTAP systems and SolidFire storage systems.
- If you are break a FlexGroup volume relationship, you must refresh the page to view the updated status of the relationship.

### Steps

1. Click **Protection** > **Relationships**.
2. Select the mirror relationship that you want to break.
3. Click **Operations** > **Break**.
4. Select the confirmation check box, and then click **Break**.

### Result

The data protection SnapMirror relationship is broken. The destination volume type changes from data protection (DP) read-only to read/write. The system stores the base Snapshot copy for the data protection mirror relationship for later use.

### Related reference

*Protection window* on page 336

You can use the Protection window to create and manage mirror relationships, vault relationships, and mirror vault relationships and to display details about these relationships. The Protection window does not display load-sharing (LS) relationships and transition (TDP) relationships.

## Resynchronizing mirror relationships

You can reestablish a mirror relationship that was broken earlier. You can perform a resynchronization operation to recover from a disaster that disabled the source volume.

### Before you begin

The source and destination clusters and the source and destination storage virtual machines (SVMs) must be in peer relationships.

### About this task

- When you perform a resynchronization operation, the contents on the mirror destination are overwritten by the contents on the source.

**Attention:** The resynchronization operation can cause loss of newer data written to the destination volume after the base Snapshot copy was created.

- If the Last Transfer Error field in the Protection window recommends a resynchronization operation, you must first break the relationship and then perform the resynchronization operation.

### Steps

1. Click **Protection** > **Relationships**.
2. Select the mirror relationship that you want to resynchronize.
3. Click **Operations** > **Resync**.
4. Select the confirmation check box and click **Resync**.

### Related reference

*Protection window* on page 336

You can use the Protection window to create and manage mirror relationships, vault relationships, and mirror vault relationships and to display details about these relationships. The Protection window does not display load-sharing (LS) relationships and transition (TDP) relationships.

## Reverse resynchronizing mirror relationships

You can use System Manager to reestablish a mirror relationship that was previously broken. In a reverse resynchronization operation, you reverse the functions of the source and destination.

### Before you begin

The source volume must be online.

### About this task

- You can use the destination volume to serve data while you repair or replace the source, update the source, and reestablish the original configuration of the systems.
- When you perform reverse resynchronization, the contents on the mirror source are overwritten by the contents on the destination.

**Attention:** This operation can cause data loss on the source.

- When you perform reverse resynchronization, the mirror policy of the relationship is set to DPDefault and the mirror schedule is set to None.

### Steps

1. Click **Protection** > **Relationships**.
2. Select the mirror relationship that you want to reverse.
3. Click **Operations** > **Reverse Resync**.
4. Select the confirmation check box, and click **Reverse Resync**.

### Related reference

*Protection window* on page 336

You can use the Protection window to create and manage mirror relationships, vault relationships, and mirror vault relationships and to display details about these relationships. The Protection window does not display load-sharing (LS) relationships and transition (TDP) relationships.

## Aborting a mirror transfer

You can abort a volume replication operation before the data transfer is complete. You can abort a scheduled update, a manual update, or an initial data transfer.

### Steps

1. Click **Protection** > **Relationships**.
2. Select the relationship for which you want to stop the data transfer, and click **Operations** > **Abort**.
3. Select the **Yes, I want to abort the transfer** check box to confirm the operation.
4. Optional: Select the **Keep any partially transferred data** check box to retain the data that is already transferred to the destination volume.
5. Click **Abort**.

### Related reference

*Protection window* on page 336

You can use the Protection window to create and manage mirror relationships, vault relationships, and mirror vault relationships and to display details about these relationships. The Protection window does not display load-sharing (LS) relationships and transition (TDP) relationships.

## Restoring a volume in a mirror relationship

For a version-independent mirror relationship, you can use System Manager to restore Snapshot copies to a source volume or to other volumes if the source data is corrupted and is no longer usable. You can replace the original data with the Snapshot copies in the destination volume.

### Before you begin

- The SnapMirror license must be enabled on both the source cluster and the destination cluster or on the nodes that contain the source volume and destination volume.
- The source cluster and destination cluster must be in a healthy peer relationship.
- The source aggregate or any other aggregate that you select for the restore operation must be a 64-bit aggregate.
- If you are connecting from a cluster running ONTAP 9.2 or earlier to a remote cluster on which SAML authentication is enabled, password-based authentication must also be enabled on the remote cluster.

### About this task

- You cannot restore a volume that is in a mirror relationship between a source storage virtual machine (SVM) and a destination SVM in a MetroCluster configuration.
- You can restore a mirror relationship between sync-source SVMs in a MetroCluster configuration.
- You can restore a mirror relationship from a volume on a sync-source SVM to a default SVM.
- You can restore a mirror relationship from a volume on a default SVM to a DP volume on a sync-source SVM.

### Steps

1. Click **Protection** > **Relationships**.
2. Select the mirror relationship, and then click **Operations** > **Restore**.

3. In the **Restore** dialog box, restore the data to the source volume in the mirror relationship or select any other volume:

| If you want to restore the data to... | Do this... |
|---|---|
| The source volume | **a.** Select **Source volume**.<br>**b.** Go to Step 7. |
| Any other volume | Select **Other volume**, and then select the cluster and SVM from the list. |

4. Restore the data to a new volume or to an existing volume:

| If you want to restore the data to... | Do this... |
|---|---|
| A new volume | If you want to change the default name, displayed in the format *`destination_SVM_name_destination_volume_name_`* restore, specify a new name, and then select the containing aggregate for the volume. |
| An existing volume | Select the **Select Volume** option.<br>You must select a volume other than the source volume, or a read/write volume with some data in it and with a common Snapshot copy.<br>Only those volumes with the same language attribute as the source volume are listed. |

5. Select either the latest Snapshot copy or the specific Snapshot copy that you want to restore.
6. Select the confirmation check box to restore the volume from the Snapshot copy.
7. Optional: Select the **Enable Network Compression** check box to compress the data that is being transferred during the restore operation.
8. Click **Restore**.

## How SnapMirror relationships work

You can create a data protection mirror relationship to a destination within a cluster to protect your data. For greater disaster protection, you can also create a mirror relationship to a destination in a different cluster in a different location.

A data protection mirror configuration consists of a source volume that can be replicated to one or more destination volumes. Each data protection mirror relationship is independent from the other data protection mirror relationships.

**Note:** The destination volume must be running either the same ONTAP version as that of the source volume or a later version of ONTAP than that of the source volume.

Snapshot copies are used to update destination volumes. Snapshot copies are transferred from the source volume to the destination volume by using an automated schedule or manually; therefore, mirrors copies are updated asynchronously.

You can create data protection mirror relationships to destinations that are on the same aggregate as the source volume as well as to destinations that are on the same storage virtual machine (SVM) or on a different SVM. For greater protection, you can create the relationships to destinations on a different aggregate, which enables you to recover from any failure of the source volume's aggregate. However, these two configurations do not protect against a cluster failure.

To protect against a cluster failure, you can create a data protection mirror relationship in which the source volume is on one cluster and the destination volume is on a different cluster. If the

cluster on which the source volume resides experiences a disaster, you can direct user clients to the destination volume on the cluster peer until the source volume is available again.

# Vault relationships

You can use System Manager to create and manage vault relationships by using the vault policy.

## Creating a vault relationship from a destination SVM

You can use System Manager to create a vault relationship from the destination storage virtual machine (SVM), and to assign a vault policy to create a backup vault. In the event of data loss or corruption on a system, backed-up data can be restored from the backup vault destination.

### Before you begin

- The source cluster must be running ONTAP 8.2.2 or later.
- SnapVault license or SnapMirror license must be enabled on both the source cluster and the destination cluster.

  **Note:** For some platforms, it is not mandatory for the source cluster to have the SnapVault license or the SnapMirror license enabled if the destination cluster has the SnapVault license or the SnapMirror license, and DPO license enabled.
- The source cluster and destination cluster must be in a healthy peer relationship.
- The destination SVM must have space available.
- The source aggregate and the destination aggregate must be 64-bit aggregates.
- A source volume of type read/write (rw) must exist.
- A vault (XDP) policy must exist.

  If a vault policy does not exist, you must create a vault policy or accept the default vault policy (XDPDefault) that is automatically assigned.
- FlexVol volumes must be online and read/write.
- The SnapLock aggregate type must be the same.
- If you are connecting from a cluster running ONTAP 9.2 or earlier to a remote cluster on which SAML authentication is enabled, password-based authentication must be enabled on the remote cluster.

### About this task

- System Manager does not support a cascade relationship.

  For example, a destination volume in a relationship cannot be the source volume in another relationship.
- You cannot create a vault relationship between a sync-source SVM and a sync-destination SVM in a MetroCluster configuration.
- You can create a vault relationship between sync-source SVMs in a MetroCluster configuration.
- You can create a vault relationship from a volume on a sync-source SVM to a volume on a data-serving SVM.
- You can create a vault relationship from a volume on a data-serving SVM to a data protection (DP) volume on a sync-source SVM.
- You can create a vault relationship only between a non-SnapLock (primary) volume and a Snaplock destination (secondary) volume.
- A maximum of 25 volumes can be protected in one selection.

### Steps

1. Click **Protection** > **Relationships**.
2. In the **Relationships** window, click **Create**.

3. In the **Browse SVM** dialog box, select an SVM for the destination volume.
4. In the **Create Protection Relationship** dialog box, select **Vault** from the **Relationship Type** drop-down list.
5. Specify the cluster, the SVM, and the source volume.

   If the specified cluster is running a version of ONTAP software earlier than ONTAP 9.3, then only peered SVMs are listed. If the specified cluster is running ONTAP 9.3 or later, peered SVMs and permitted SVMs are listed.
6. Enter a volume name suffix.

   The volume name suffix is appended to the source volume names to generate the destination volume names.
7. If you are creating a SnapLock volume, specify the default retention period.

   The default retention period can be set to any value between 1 day through 70 years or Infinite.
8. Optional: Click **Browse**, and then change the vault policy.
9. Select a schedule for the relationship from the list of existing schedules.
10. Optional: Select **Initialize Relationship** to initialize the vault relationship.
11. Enable SnapLock aggregates, and then select a SnapLock Compliance aggregate or a SnapLock Enterprise aggregate.
12. Enable FabricPool-enabled aggregates, and then select an appropriate tiering policy.
13. Click **Validate** to verify whether the selected volumes have matching labels.
14. Click **Create**.

### Result

If you chose to create a destination volume, a volume of type *dp* is created with the following default settings:

- Autogrow is enabled.
- Deduplication is enabled or disabled according to the user preference or the source volume deduplication setting.
- Compression is disabled.
- The language attribute is set to match the language attribute of the source volume.

A vault relationship is created between the destination volume and the source volume. The base Snapshot copy is transferred to the destination volume if you have opted to initialize the relationship.

### Related reference

*Protection window* on page 336
You can use the Protection window to create and manage mirror relationships, vault relationships, and mirror vault relationships and to display details about these relationships. The Protection window does not display load-sharing (LS) relationships and transition (TDP) relationships.

## Deleting vault relationships

You can use System Manager to end a vault relationship between a source and destination volume, and release the Snapshot copies from the source.

### About this task

Releasing the relationship permanently removes the base Snapshot copies used by the vault relationship on the source volume. To re-create the vault relationship, you must run the resynchronization operation from the source volume by using the command-line interface (CLI).

### Steps

1. Click **Protection** > **Relationships**.
2. Select the volume for which you want to delete the vault relationship, and click **Delete**.
3. Select the confirmation check box, and then click **Delete**.

   You can also select the release base Snapshot copies check box to delete the base Snapshot copies used by the vault relationship on the source volume.

   If the relationship is not released, then you must use the CLI to run the release operation on the source cluster to delete the base Snapshot copies that were created for the vault relationship from the source volume.

### Related reference

*Protection window* on page 336

You can use the Protection window to create and manage mirror relationships, vault relationships, and mirror vault relationships and to display details about these relationships. The Protection window does not display load-sharing (LS) relationships and transition (TDP) relationships.

## Editing vault relationships

You can use System Manager to edit a vault relationship either by selecting an existing policy or schedule in the cluster, or by creating a new policy or schedule. However, you cannot edit the parameters of an existing policy or schedule.

### Before you begin

The source and destination clusters must be in a healthy peer relationship.

### Steps

1. Click **Protection** > **Relationships**.
2. Select the vault relationship for which you want to modify the policy or schedule, and then click **Edit**.
3. In the **Edit Relationship** dialog box, select the appropriate action:

| If you want to… | Do the following… |
|---|---|
| Select an existing policy | Click **Browse**, and then select an existing policy.<br>You can select a policy that has the maximum number of matching labels with the Snapshot policy that is attached to the source volume. |
| Create a new policy | **a.** Click **Create Policy**.<br>**b.** Specify a name for the policy.<br>**c.** Set the priority for scheduled transfers.<br>Low indicates that the transfer has the least priority and is usually scheduled after normal priority transfers. By default, the priority is set to Normal.<br>**d.** Select the **Enable Network Compression** check box to compress the data that is being transferred.<br>**e.** Specify a SnapMirror label and destination retention count for the vault policy.<br>You must ensure that a Snapshot copy with the same label is created on the source volume for the new SnapMirror label to be effective.<br>**f.** Click **Create**. |

4. Specify a schedule for the relationship:

| If… | Do the following… |
|---|---|
| You want to assign an existing schedule | Select an existing schedule from the list. |
| You want to create a new schedule | **a.** Click **Create Schedule**.<br>**b.** Specify a name for the schedule.<br>**c.** Select one of the following options:<br>• **Basic**<br>You can select this option to specify only the day of the week, time, and the transfer interval.<br>• **Advanced**<br>You can select this option to specify a cron-style schedule.<br>**d.** Click **Create**. |
| You do not want to assign a schedule | Select **None**. |

5. Click **OK**.

### Related reference

*Protection window* on page 336

You can use the Protection window to create and manage mirror relationships, vault relationships, and mirror vault relationships and to display details about these relationships. The Protection window does not display load-sharing (LS) relationships and transition (TDP) relationships.

## Initializing a vault relationship

You can use System Manager to initialize a vault relationship if you have not already initialized it while creating the relationship. A baseline transfer of data is initiated from the source FlexVol volume to the destination FlexVol volume.

### Before you begin

The source and destination clusters must be in a healthy peer relationship.

### Steps

1. Click **Protection** > **Relationships**.
2. Select the relationship you want to initialize, and click **Operations** > **Initialize**.
3. In the **Initialize** window, click **Initialize**.

### Result

A Snapshot copy is created and transferred to the destination.

This Snapshot copy is used as a baseline for subsequent incremental Snapshot copies.

### Related reference

*Protection window* on page 336

You can use the Protection window to create and manage mirror relationships, vault relationships, and mirror vault relationships and to display details about these relationships. The Protection window does not display load-sharing (LS) relationships and transition (TDP) relationships.

## Updating a vault relationship

You can use System Manager to manually initiate an unscheduled incremental update. You might require a manual update to prevent data loss due to an upcoming power outage, scheduled maintenance, or data migration.

### Before you begin

The vault relationship must be initialized.

### Steps

1. Click **Protection** > **Relationships**.
2. Select the relationship for which you want to update the data, and click **Operations** > **Update**.
3. Choose one of the following options:
   - Select **As Per Policy** to perform an incremental transfer from the recent common Snapshot copy between the source and destination volumes.
   - Select **Select Snapshot copy** and specify the Snapshot copy that you want to transfer.
4. Optional: Select **Limit transfer bandwidth to** to limit the network bandwidth that is used for transfers and specify the maximum transfer speed.
5. Click **Update**.
6. Verify the transfer status in the **Details** tab.

### Related reference

*Protection window* on page 336

You can use the Protection window to create and manage mirror relationships, vault relationships, and mirror vault relationships and to display details about these relationships. The Protection window does not display load-sharing (LS) relationships and transition (TDP) relationships.

## Quiescing a vault relationship

You can use System Manager to disable data transfers to the destination FlexVol volume by quiescing the vault relationship.

### Steps

1. Click **Protection** > **Relationships**.
2. Select the relationship for which you want to stop the scheduled data transfers, and click **Operations** > **Quiesce**.
3. In the **Quiesce** window, click **Quiesce**.

### Result

If there is no transfer in progress, the transfer status is displayed as Quiesced. If a transfer is in progress, the transfer is not affected, and the transfer status is displayed as Quiescing until the transfer is complete.

### Related reference

*Protection window* on page 336

You can use the Protection window to create and manage mirror relationships, vault relationships, and mirror vault relationships and to display details about these relationships. The Protection window does not display load-sharing (LS) relationships and transition (TDP) relationships.

## Resuming a vault relationship

You can resume a quiesced vault relationship by using System Manager. When you resume the relationship, normal data transfer to the destination FlexVol volume is resumed and all vault activities are restarted.

### Steps

1. Click **Protection** > **Relationships**.
2. Select the relationship for which you want to resume the data transfer, and click **Operations** > **Resume**.
3. In the **Resume** window, click **Resume**.

### Result

Normal data transfers are resumed. If there is a scheduled transfer for the relationship, the transfer is started from the next schedule.

### Related reference

*Protection window* on page 336

You can use the Protection window to create and manage mirror relationships, vault relationships, and mirror vault relationships and to display details about these relationships. The Protection window does not display load-sharing (LS) relationships and transition (TDP) relationships.

## Aborting a Snapshot copy transfer

You can use System Manager to abort or stop a data transfer that is currently in progress.

### Steps

1. Click **Protection** > **Relationships**.
2. Select the relationship for which you want to stop the data transfer, and click **Operations** > **Abort**.
3. Select the **Yes, I want to abort the transfer** check box to confirm the operation.
4. Optional: Select the **Keep any partially transferred data** check box to retain the data that is already transferred to the destination volume.
5. Click **Abort**.

### Result

The transfer status is displayed as "Aborting" until the operation is complete and displayed as "Idle" after the operation is complete.

### Related reference

*Protection window* on page 336

You can use the Protection window to create and manage mirror relationships, vault relationships, and mirror vault relationships and to display details about these relationships. The Protection window does not display load-sharing (LS) relationships and transition (TDP) relationships.

## Restoring a volume in a vault relationship

You can use System Manager to restore Snapshot copies to a source volume or to other volumes if the source data is corrupted and is no longer usable. You can replace the original data with the Snapshot copies in the destination volume.

### Before you begin

- The SnapMirror license must be enabled on both the source storage system and the destination storage system or on the nodes that contain the source volume and destination volume.
- The source cluster and destination cluster must be in a healthy peer relationship.
- The source aggregate or any other aggregate that you select for the restore operation must be a 64-bit aggregate.
- If you are connecting from a cluster running ONTAP 9.2 or earlier to a remote cluster on which SAML authentication is enabled, password-based authentication must also be enabled on the remote cluster.

### About this task

- You cannot restore a volume that is in a vault relationship between a source storage virtual machine (SVM) and a destination SVM in a MetroCluster configuration.
- You can restore a vault relationship between sync-source SVMs in a MetroCluster configuration.
- You can restore a vault relationship from a volume on a sync-source SVM to a default SVM.
- You can restore a vault relationship from a volume on a default SVM to a DP volume on a sync-source SVM.

### Steps

1. Click **Protection** > **Relationships**.
2. Select the vault relationship, and then click **Operations** > **Restore**.
3. In the **Restore** dialog box, restore the data to the source volume in the vault relationship or select any other volume:

| If you want to restore the data to... | Do this... |
|---|---|
| The source volume | **a.** Select **Source volume**.<br>**b.** Go to Step 6. |
| Any other volume | Select **Other volume**, and then select the cluster and SVM from the list. |

4. Restore the data to a new volume or select any existing volume:

| If you want to restore the data to... | Do this... |
|---|---|
| A new volume | If you want to change the default name, displayed in the format *destination_SVM_name_destination_volume_name_* restore, specify a new name, and then select the containing aggregate for the volume. |

| If you want to restore the data to... | Do this... |
|---|---|
| An existing volume | Select the **Select Volume** option.<br>You must select a volume other than the source volume, or a read/write volume with some data in it and with a common Snapshot copy.<br>Only those volumes with the same language attribute as the source volume are listed. |

5. Select either the latest Snapshot copy or the specific Snapshot copy that you want to restore.
6. Select the confirmation check box to restore the volume from the Snapshot copy.
7. Optional: Select the **Enable Network Compression** check box to compress the data that is being transferred during the restore operation.
8. Click **Restore**.

### Related reference

*Protection window* on page 336

You can use the Protection window to create and manage mirror relationships, vault relationships, and mirror vault relationships and to display details about these relationships. The Protection window does not display load-sharing (LS) relationships and transition (TDP) relationships.

## What a SnapVault backup is

A SnapVault backup is a collection of Snapshot copies on a FlexVol volume that you can restore data from if the primary data is not usable. Snapshot copies are created based on a Snapshot policy. The SnapVault backup backs up Snapshot copies based on its schedule and SnapVault policy rules.

A SnapVault backup is a disk-to-disk backup solution that you can also use to offload tape backups. In the event of data loss or corruption on a system, backed-up data can be restored from the SnapVault secondary volume with less downtime and uncertainty than is associated with conventional tape backup and restore operations.

The following terms are used to describe SnapVault backups:

**baseline transfer**

An initial complete backup of a primary storage volume to a corresponding volume on the secondary system.

**secondary volume**

A volume to which data is backed up from a primary volume. Such a volume can be a secondary or tertiary (and onward) destination in a cascade or fanout backup configuration. The SnapVault secondary system maintains Snapshot copies for long-term storage and possible restore operations.

**incremental transfer**

A follow-up backup to the secondary system that contains only the changes to the primary data since the last transfer action.

**SnapMirror label**

An attribute that identifies Snapshot copies for the purpose of selection and retention in SnapVault backups. Each SnapVault policy configures the rules for selecting Snapshot copies on the primary volume and transferring the Snapshot copies that match a given SnapMirror label.

**Snapshot copy**

The backup images on the source volume that are created manually or automatically as scheduled by an assigned policy. Baseline Snapshot copies contain a copy of the entire source data being protected; subsequent Snapshot copies contain differential copies of the source data. Snapshot copies can be stored on the source volume or on a different destination volume in a different storage virtual machine (SVM) or cluster.

Snapshot copies capture the state of volume data on each source system. For SnapVault and mirror relationships, this data is transferred to destination volumes.

**primary volume**

A volume that contains data that is to be backed up. In cascade or fanout backup deployments, the primary volume is the volume that is backed up to a SnapVault backup, regardless of where in the chain the SnapVault source is. In a cascade chain configuration in which A has a mirror relationship to B and B has a SnapVault relationship to C, B serves as the source for the SnapVault backup even though it is a secondary destination in the chain.

**SnapVault relationship**

A backup relationship, configured as a SnapVault relationship, between a primary volume and a secondary volume.

### Related reference

*Protection window* on page 336

You can use the Protection window to create and manage mirror relationships, vault relationships, and mirror vault relationships and to display details about these relationships. The Protection window does not display load-sharing (LS) relationships and transition (TDP) relationships.

# Mirror and vault relationships

You can use System Manager to create and manage mirror and vault relationships by using the mirror and vault policy.

## Creating a mirror and vault relationship from a destination SVM

You can use System Manager to create a mirror and vault relationship from the destination storage virtual machine (SVM). Creating this relationship enables you to better protect your data by periodically transferring data from the source volume to the destination volume. It also enables you to retain data for long periods by creating backups of the source volume.

### Before you begin

- The destination cluster must be running ONTAP 8.3.2 or later.
- SnapMirror license must be enabled on the source cluster and destination cluster.

  **Note:** For some platforms, it is not mandatory for the source cluster to have the SnapMirror license enabled if the destination cluster has the SnapMirror license and Data Protection Optimization (DPO) license enabled.
- The source cluster and destination cluster must be in a healthy peer relationship.
- The destination SVM must have space available.
- The source aggregate and destination aggregate must be 64-bit aggregates.
- A source volume of type read/write (rw) must already exist.
- The SnapLock aggregate type must be the same.
- If you are connecting from a cluster running ONTAP 9.2 or earlier to a remote cluster on which SAML authentication is enabled, password-based authentication must be enabled on the remote cluster.

### About this task

- System Manager does not support a cascade relationship.

  For example, a destination volume in a relationship cannot be the source volume in another relationship.
- You cannot create a mirror and vault relationship between a sync-source SVM and a sync-destination SVM in a MetroCluster configuration.
- You can create a mirror and vault relationship between sync-source SVMs in a MetroCluster configuration.
- You can create a mirror and vault relationship from a volume on a sync-source SVM to a volume of a data-serving SVM.
- You can create a mirror and vault relationship from a volume on a data-serving SVM to a DP volume on a sync-source SVM.
- A maximum of 25 volumes can be protected in one selection.

### Steps

1. Click **Protection** > **Relationships**.
2. In the **Relationships** window, click **Create**.
3. In the **Browse SVM** dialog box, select an SVM for the destination volume.
4. In the **Create Protection Relationship** dialog box, select **Mirror and Vault** from the **Relationship Type** drop-down list.
5. Specify the cluster, the SVM, and the source volume.

   If the specified cluster is running a version of ONTAP software earlier than ONTAP 9.3, then only peered SVMs are listed. If the specified cluster is running ONTAP 9.3 or later, peered SVMs and permitted SVMs are listed.
6. Enter a volume name suffix.

   The volume name suffix is appended to the source volume names to generate the destination volume names.
7. Optional: Click **Browse**, and then change the mirror and vault policy.

   You can select the policy that has the maximum number of matching labels with the Snapshot policy that is attached to the source volume.
8. Select a schedule for the relationship from the list of existing schedules.
9. Optional: Select **Initialize Relationship** to initialize the relationship.
10. Enable FabricPool-enabled aggregates, and then select an appropriate tiering policy.
11. Click **Validate** to verify whether the selected volumes have matching labels.
12. Click **Create**.

## Deleting mirror and vault relationships

You can use System Manager to end a mirror and vault relationship between a source and destination volume, and release the Snapshot copies from the source volume.

### About this task

- It is a best practice to break the mirror and vault relationship before deleting the relationship.
- To re-create the relationship, you must run the resynchronization operation from the source volume by using the command-line interface (CLI).

### Steps

1. Click **Protection** > **Relationships**.
2. Select the mirror and vault relationship that you want to delete and click **Delete**.

3. Select the confirmation check box, and then click **Delete**.

   You can also select the release base Snapshot copies check box to delete the base Snapshot copies used by the mirror and vault relationship on the source volume.

   If the relationship is not released, then you must use the CLI to run the release operation on the source cluster to delete the base Snapshot copies that were created for the mirror and vault relationship from the source volume.

### Result

The relationship is deleted and the base Snapshot copies on the source volume are permanently deleted.

## Editing mirror and vault relationships

You can use System Manager to edit a mirror and vault relationship by modifying the selected policy or schedule. However, you cannot edit the parameters of an existing policy or schedule.

### Before you begin

The source and destination clusters must be in a healthy peer relationship.

### About this task

You can modify the relationship type of a version-flexible mirror relationship, vault relationship, or mirror and vault relationship by modifying the policy type.

### Steps

1. Click **Protection** > **Relationships**.
2. Select the mirror and vault relationship that you want to modify, and then click **Edit**.
3. In the **Edit Relationship** dialog box, select the appropriate action:

| If you want to… | Do the following… |
|---|---|
| Select an existing policy | Click **Browse**, and then select an existing policy.<br>You can select a policy that has the maximum number of matching labels with the Snapshot policy that is attached to the source volume. |
| Create a new policy | **a.** Click **Create Policy**.<br>**b.** Specify a name for the policy.<br>**c.** Set the priority for scheduled transfers.<br>Low indicates that the transfer has the least priority and is usually scheduled after normal priority transfers. By default, the priority is set to Normal.<br>**d.** Select the **Enable Network Compression** check box to compress the data that is being transferred.<br>**e.** Specify a SnapMirror label and destination retention count for the vault policy.<br>You must ensure that a Snapshot copy with the same label is created on the source volume for the new SnapMirror label to be effective.<br>**f.** Click **Create**. |

4. Specify a schedule for the relationship:

| If… | Do the following… |
|---|---|
| You want to assign an existing schedule | Click **Browse**, and then select an existing schedule. |

| If… | Do the following… |
|---|---|
| You want to create a new schedule | **a.** Click **Create Schedule**.<br>**b.** Specify a name for the schedule.<br>**c.** Select one of the following options:<br>• **Basic**<br>You can select this option to specify only the day of the week, time, and the transfer interval.<br>• **Advanced**<br>You can select this option to specify a cron style schedule.<br>**d.** Click **Create**. |
| You do not want to assign a schedule | Select **None**. |

5. Click **OK**.

## Initializing mirror and vault relationships

You can use System Manager to initialize a mirror and vault relationship if you have not already initialized the relationship while creating it. When you initialize a relationship, a complete baseline transfer of data is performed from the source volume to the destination.

### Before you begin

The source and destination clusters must be in a healthy peer relationship.

### Steps

1. Click **Protection** > **Relationships**.
2. Select the mirror and vault relationship that you want to initialize, and then click **Operations** > **Initialize**.
3. Select the confirmation check box, and then click **Initialize**.
4. Verify the status of the relationship in the **Protection** window.

### Result

A Snapshot copy is created and transferred to the destination.

This Snapshot copy is used as a baseline for subsequent incremental Snapshot copies.

## Updating mirror and vault relationships

You can use System Manager to manually initiate an unscheduled incremental update. You might require a manual update to prevent data loss due to an upcoming power outage, scheduled maintenance, or data migration.

### Before you begin

The mirror and vault relationship must be initialized and in a Snapmirrored state.

### Steps

1. Click **Protection** > **Relationships**.
2. Select the mirror relationship for which you want to update the data, and then click **Operations** > **Update**.
3. Choose one of the following options:

- Select **As Per Policy** to perform an incremental transfer from the recent common Snapshot copy between the source and destination volumes.
- Select **Select Snapshot copy** and specify the Snapshot copy that you want to transfer.

4. Optional: Select **Limit transfer bandwidth to** to limit the network bandwidth that is used for transfers, and then specify the maximum transfer speed.
5. Click **Update**.
6. Verify the transfer status in the **Details** tab.

## Quiescing mirror and vault relationships

You can use System Manager to quiesce a destination volume to stabilize the destination before creating a Snapshot copy. The quiesce operation enables active data transfers to finish and disables future transfers for the mirror and vault relationship.

### Before you begin

The mirror and vault relationship must be in a Snapmirrored state.

### Steps

1. Click **Protection** > **Relationships**.
2. Select the mirror and vault relationship that you want to quiesce, and then click **Operations** > **Quiesce**.
3. Select the confirmation check box, and then click **Quiesce**.

### Result

If there is no transfer in progress, the transfer status is displayed as Quiesced. If a transfer is in progress, the transfer is not affected, and the transfer status is displayed as Quiescing until the transfer is complete.

## Resuming mirror and vault relationships

If you have a quiesced mirror and vault relationship, you can resume the relationship by using System Manager. When you resume the relationship, normal data transfer to the destination volume is resumed and all the protection activities are restarted.

### About this task

If you have quiesced a broken mirror and vault relationship from the command-line interface (CLI), you cannot resume the relationship from System Manager. You must use the CLI to resume the relationship.

### Steps

1. Click **Protection** > **Relationships**.
2. Select the mirror and vault relationship that you want to resume, and then click **Operations** > **Resume**.
3. Select the confirmation check box, and then click **Resume**.

### Result

Normal data transfers are resumed. If there is a scheduled transfer for the relationship, the transfer is started from the next schedule.

## Breaking mirror and vault relationships

You can use System Manager to break a mirror and vault relationship if a source volume becomes unavailable and you want client applications to access the data from the destination volume. You

can use the destination volume to serve data while you repair or replace the source volume, update the source volume, and reestablish the original configuration of the systems.

### Before you begin

- The mirror and vault relationship must be in the quiesced state or idle state.
- The destination volume must be mounted on the destination storage virtual machine (SVM) namespace.

### About this task

You can break mirror relationships between ONTAP systems and SolidFire storage systems.

### Steps

1. Click **Protection** > **Relationships**.
2. Select the mirror and vault relationship that you want to break, and then click **Operations** > **Break**.
3. Select the confirmation check box, and then click **Break**.

### Result

The mirror and vault relationship is broken. The destination volume type changes from data protection (DP) read-only to read/write. The system stores the base Snapshot copy for the mirror and vault relationship for later use.

## Resynchronizing mirror and vault relationships

You can use System Manager to reestablish a mirror and vault relationship that was broken earlier. You can perform a resynchronization operation to recover from a disaster that disabled the source volume.

### Before you begin

The source and destination clusters and the source and destination storage virtual machines (SVMs) must be in peer relationships.

### About this task

- When you perform a resynchronization operation, the contents on the destination volume are overwritten by the contents on the source.

  **Attention:** The resynchronization operation can cause loss of newer data written to the destination volume after the base Snapshot copy was created.

- If the Last Transfer Error field in the Protection window recommends a resynchronization operation, you must first break the relationship and then perform the resynchronization operation.

### Steps

1. Click **Protection** > **Relationships**.
2. Select the mirror and vault relationship that you want to resynchronize, and then click **Operations** > **Resync**.
3. Select the confirmation check box, and then click **Resync**.

## Reverse resynchronizing mirror and vault relationships

You can use System Manager to reestablish a mirror and vault relationship that was previously broken. In a reverse resynchronization operation, the functions of the source and destination

volumes are reversed. You can use the destination volume to serve data while you repair or replace the source, update the source, and reestablish the original configuration of the systems.

### Before you begin

The source volume must be online.

### About this task

- When you perform reverse resynchronization, the contents on the source volume are overwritten by the contents on the destination volume.

  **Attention:** The reverse resynchronization operation can cause data loss on the source volume.

- When you perform reverse resynchronization, the policy of the relationship is set to MirrorAndVault and the schedule is set to None.

### Steps

1. Click **Protection** > **Relationships**.
2. Select the mirror and vault relationship that you want to reverse, and then click **Operations** > **Reverse Resync**.
3. Select the confirmation check box, and then click **Reverse Resync**.

## Aborting mirror and vault relationships

You can abort a volume replication operation if you want to stop the data transfer. You can abort a scheduled update, a manual update, or an initial data transfer.

### Steps

1. Click **Protection** > **Relationships**.
2. Select the mirror and vault relationship for which you want to stop the data transfer, and then click **Operations** > **Abort**.
3. Select the **Yes, I want to abort the transfer** check box to confirm the operation.
4. Optional: Select the **Keep any partially transferred data** check box to retain the data that is already transferred to the destination volume.
5. Click **Abort**.

### Result

The transfer status is displayed as "Aborting" until the operation is complete and displayed as "Idle" after the operation is complete.

## Restoring a volume in a mirror and vault relationship

You can use System Manager to restore Snapshot copies to a source volume or to other volumes if the source data is corrupted and is no longer usable. You can replace the original data with the Snapshot copies in the destination volume.

### Before you begin

- The SnapMirror license and SnapVault license must be enabled on both the source cluster and the destination cluster or on the nodes that contain the source volume and destination volume.
- The source cluster and destination cluster must be in a healthy peer relationship.
- The source aggregate or any other aggregate that you select for the restore operation must be a 64-bit aggregate.

- If you are connecting from a cluster running ONTAP 9.2 or earlier to a remote cluster on which SAML authentication is enabled, password-based authentication must also be enabled on the remote cluster.

### About this task

- You cannot restore a volume that is in a mirror and vault relationship between a source storage virtual machine (SVM) and a destination SVM in a MetroCluster configuration.
- You can restore a mirror and vault relationship for the following configurations:
  - Between sync-source SVMs in a MetroCluster configuration
  - From a volume on a sync-source SVM to a default SVM
  - From a volume on a default SVM to a DP volume on a sync-source SVM

### Steps

1. Click **Protection** > **Relationships**.
2. Select the mirror and vault relationship that you want to restore, and then click **Operations** > **Restore**.
3. In the **Restore** dialog box, restore the data to the source volume in the relationship or select any other volume:

| If you want to restore the data to… | Do this… |
|---|---|
| The source volume | **a.** Select **Source volume**.<br>**b.** *Go to step 6*. |
| Any other volume | Select **Other volume**, and then select the cluster and the SVM. |

4. Restore the data to a new volume or to an existing volume:

| If you want to restore the data to… | Do this… |
|---|---|
| A new volume | If you want to change the default name, displayed in the format "destination_SVM_name_destination_volume_name_restore", specify a new name, and then select the containing aggregate for the volume. |
| An existing volume | Select the **Select Volume** option.<br>You must select a volume other than the source volume, or a read/write volume with some data in it and with a common Snapshot copy.<br>Only those volumes with the same language attribute as the source volume are listed. |

5. Select either the latest Snapshot copy or the specific Snapshot copy that you want to restore.
6. Select the confirmation check box to restore the volume from the Snapshot copy.
7. Optional: Select the **Enable Network Compression** check box to compress the data that is being transferred during the restore operation.
8. Click **Restore**.

## What lag time is

Lag time is the amount of time by which the destination system lags behind the source system.

The lag time is the difference between the current time and the timestamp of the Snapshot copy that was last successfully transferred to the destination system. The lag time will always be at least

as much as the duration of the last successful transfer, unless the clocks on the source and destination systems are not synchronized. The time zone difference is automatically calculated into the lag time.

## Types of data protection relationships

Depending on your data protection and backup requirements, OnCommand System Manager provides different types of protection relationships that enable you to protect data against accidental, malicious, or disaster-induced loss of data.

### Mirror relationship (SnapMirror license required)

A mirror relationship provides asynchronous disaster recovery. Data protection mirror relationships enable you to periodically create Snapshot copies of the data on one volume, to copy those Snapshot copies to a partner volume (the destination volume), which is usually on another cluster, and then to retain those Snapshot copies. If the data on the source volume is corrupted or lost, the mirror copy on the destination volume ensures quick availability and restoration of data from the time of the latest Snapshot copy.

For mirror relationships, the version of ONTAP that is running on the destination cluster must be the same version as or a later version than the ONTAP version running on the source cluster. However, version-flexible mirror relationships are not dependent on the ONTAP version. Therefore, you can create a version-flexible mirror relationship with a destination cluster that is running either a later ONTAP version or an earlier ONTAP version than the ONTAP version of the source cluster or an earlier version of ONTAP than the ONTAP version of the source cluster.

> **Note:** The version-flexible mirror relationship feature is available only from ONTAP 8.3 onward. You cannot have a version-flexible mirror relationship with a volume in Data ONTAP 8.3 or earlier.

Mirror relationships are valid for FlexVol volumes.

### Vault relationship (SnapMirror or SnapVault license required)

A vault relationship provides storage-efficient and long-term retention of backups. Vault relationships enable you to back up selected Snapshot copies of volumes to a destination volume and to retain the backups.

Vault relationships are valid only for FlexVol volumes.

### Mirror and vault relationship (SnapMirror licenses required)

A mirror and vault relationship provides data protection by periodically transferring data from the source volume to the destination volume and also facilitates long-term retention of data by creating backups of the source volume.

> **Note:** The mirror and vault relationship feature is available only from ONTAP 8.3.2 onward. You cannot have a mirror and vault relationship with a volume in Data ONTAP 8.3.2 or earlier.

A mirror and vault relationship is valid for FlexVol volumes and FlexGroup volumes.

## SnapMirror licensing

A SnapMirror license is required on both the source and destination clusters. A SnapVault license is not required if a SnapMirror license is already installed.

### DP_Optimized (DPO) license

Starting with ONTAP 9.3, a new DP_Optimized (DPO) license is available that supports an increased number of volumes and peer relationships. A SnapMirror license is still required on both the source and destination.

On the following platforms, a DPO license is required only on the destination cluster. Otherwise, it is required on both the source and destination:

- FAS22xx
- FAS25xx
- FAS26xx
- FAS62xx
- FAS80xx
- FAS82xx
- FAS9000
- V32xx
- V62xx

# Protection window

You can use the Protection window to create and manage mirror relationships, vault relationships, and mirror vault relationships and to display details about these relationships. The Protection window does not display load-sharing (LS) relationships and transition (TDP) relationships.

- *Command buttons* on page 336
- *Protection relationships list* on page 336
- *Details area* on page 337

## Command buttons

**Create**

Opens the Create Protection Relationship dialog box, which you can use to create a mirror relationship, vault relationship, or mirror and vault relationship from a destination volume.

System Manager does not display any storage virtual machine (SVM) configured for disaster recovery (DR) in the Create Protection Relationship dialog box.

**Edit**

Opens the Edit Protection Relationship dialog box, which you can use to edit the schedule and policy of a relationship.

For a vault relationship, mirror and vault relationship, or version-flexible mirror relationship, you can modify the relationship type by modifying the policy type.

**Delete**

Opens the Delete Protection Relationship dialog box, which you can use to delete a relationship.

**Operations**

Displays the operations that can be performed on a protection relationship.

**Refresh**

Updates the information in the window.

## Protection relationships list

**Source Storage Virtual Machine**

Displays the SVM that contains the volume from which data is mirrored or vaulted in a relationship.

**Source Volume**

Displays the volume from which data is mirrored or vaulted in a relationship.

**Destination Volume**

Displays the volume to which data is mirrored or vaulted in a relationship.

**Is Healthy**

Displays whether the relationship is healthy or not.

**Relationship State**

Displays the state of the relationship, such as Snapmirrored, Uninitialized, or Broken Off.

**Transfer Status**

Displays the relationship status, such as Idle, Transferring, or Aborting.

**Relationship Type**

Displays the type of relationship, such as Mirror, Vault, or Mirror and Vault.

**Lag Time**

Displays the difference between the current time and the timestamp of the Snapshot copy that was last transferred successfully to the destination storage system. It indicates the time difference between the data that is currently on the source system and the latest data that is stored on the destination system. The value that is displayed can be positive or negative. The value is negative if the time zone of the destination system is behind the time zone of the source storage system.

**Policy Name**

Displays the name of the policy that is assigned to the relationship.

**Policy Type**

Displays the type of policy that is assigned to the relationship. The policy type can be Vault, Mirror Vault, or Asynchronous Mirror.

### Details area

**Details tab**

Displays general information about the selected relationship, such as the source cluster and destination cluster, data transfer rate, state of the relationship, details about the network compression ratio, data transfer status, type of current data transfer, type of last data transfer, latest Snapshot copy, and timestamp of the latest Snapshot copy.

**Policy Details tab**

Displays details about the policy that is assigned to the selected protection relationship. This tab also displays the SnapMirror label and the Snapshot copy schedules in the source volume that match the specified label.

**Snapshot Copies tab**

Displays the count of Snapshot copies with the SnapMirror label attribute for the selected protection relationship and the timestamp of the latest Snapshot copy.

### Related concepts

*What a SnapVault backup is* on page 326

A SnapVault backup is a collection of Snapshot copies on a FlexVol volume that you can restore data from if the primary data is not usable. Snapshot copies are created based on a Snapshot policy. The SnapVault backup backs up Snapshot copies based on its schedule and SnapVault policy rules.

### Related tasks

*Creating a mirror relationship from a source SVM* on page 189

You can use System Manager to create a mirror relationship from the source storage virtual machine (SVM), and to assign a mirror policy and schedule to the mirror relationship. The mirror copy enables quick availability of data if the data on the source volume is corrupted or lost.

*Creating a mirror relationship from a destination SVM* on page 310
You can use System Manager to create a mirror relationship from the destination storage virtual machine (SVM), and to assign a policy and schedule to the mirror relationship. The mirror copy enables quick availability of data if the data on the source volume is corrupted or lost.

*Deleting mirror relationships* on page 311
You can delete a mirror relationship and permanently end the mirror relationship between the source and destination volumes. When a mirror relationship is deleted, the base Snapshot copy on the source volume is deleted.

*Editing mirror relationships* on page 312
You can use System Manager to edit a mirror relationship either by selecting an existing policy or schedule in the cluster, or by creating a policy or schedule.

*Initializing mirror relationships* on page 313
When you start a mirror relationship for the first time, you have to initialize the relationship. Initializing a relationship consists of a complete baseline transfer of data from the source volume to the destination. You can use System Manager to initialize a mirror relationship if you have not already initialized the relationship while creating it.

*Updating mirror relationships* on page 313
You can initiate an unscheduled mirror update of the destination. You might have to perform a manual update to prevent data loss due to an upcoming power outage, scheduled maintenance, or data migration.

*Quiescing mirror relationships* on page 314
Use System Manager to quiesce a mirror destination to stabilize the destination before creating a Snapshot copy. The quiesce operation enables active mirror transfers to finish and disables future transfers for the mirroring relationship.

*Resuming mirror relationships* on page 314
You can resume a quiesced mirror relationship. When you resume the relationship, normal data transfer to the mirror destination is resumed and all the mirror activities are restarted.

*Breaking SnapMirror relationships* on page 315
You must break a mirror relationship if a mirror source becomes unavailable and you want client applications to be able to access the data from the mirror destination. After the mirror relationship is broken, the destination volume type changes from DP to RW.

*Resynchronizing mirror relationships* on page 315
You can reestablish a mirror relationship that was broken earlier. You can perform a resynchronization operation to recover from a disaster that disabled the source volume.

*Reverse resynchronizing mirror relationships* on page 316
You can use System Manager to reestablish a mirror relationship that was previously broken. In a reverse resynchronization operation, you reverse the functions of the source and destination.

*Aborting a mirror transfer* on page 317
You can abort a volume replication operation before the data transfer is complete. You can abort a scheduled update, a manual update, or an initial data transfer.

*Creating a vault relationship from a source SVM* on page 191
You can use System Manager to create a vault relationship from the source storage virtual machine (SVM), and to assign a vault policy to the vault relationship to create a backup vault. In the event of data loss or corruption on a system, backed-up data can be restored from the backup vault destination.

*Creating a vault relationship from a destination SVM* on page 319

You can use System Manager to create a vault relationship from the destination storage virtual machine (SVM), and to assign a vault policy to create a backup vault. In the event of data loss or corruption on a system, backed-up data can be restored from the backup vault destination.

*Deleting vault relationships* on page 320
You can use System Manager to end a vault relationship between a source and destination volume, and release the Snapshot copies from the source.

*Editing vault relationships* on page 321
You can use System Manager to edit a vault relationship either by selecting an existing policy or schedule in the cluster, or by creating a new policy or schedule. However, you cannot edit the parameters of an existing policy or schedule.

*Initializing a vault relationship* on page 322
You can use System Manager to initialize a vault relationship if you have not already initialized it while creating the relationship. A baseline transfer of data is initiated from the source FlexVol volume to the destination FlexVol volume.

*Updating a vault relationship* on page 323
You can use System Manager to manually initiate an unscheduled incremental update. You might require a manual update to prevent data loss due to an upcoming power outage, scheduled maintenance, or data migration.

*Quiescing a vault relationship* on page 323
You can use System Manager to disable data transfers to the destination FlexVol volume by quiescing the vault relationship.

*Resuming a vault relationship* on page 324
You can resume a quiesced vault relationship by using System Manager. When you resume the relationship, normal data transfer to the destination FlexVol volume is resumed and all vault activities are restarted.

*Aborting a Snapshot copy transfer* on page 324
You can use System Manager to abort or stop a data transfer that is currently in progress.

*Restoring a volume in a vault relationship* on page 325
You can use System Manager to restore Snapshot copies to a source volume or to other volumes if the source data is corrupted and is no longer usable. You can replace the original data with the Snapshot copies in the destination volume.

# Protection policies

You can use System Manager to create, edit, and delete protection policies.

## Creating protection policies

You can use System Manager to create cluster-level asynchronous mirror policies, vault policies, or mirror and vault policies, and to apply these policies to a cluster-level data protection relationship.

### Steps

1. Click **Protection** > **Protection Policies**.
2. Click **Create**.
3. In the **Create Policy** dialog box, select the type of policy that you want to create.
4. Specify the policy name and transfer priority.

   Low indicates that the transfer has the lowest priority. Low priority transfers are usually scheduled after normal priority transfers. By default, the transfer priority is set to Normal.

5. Optional: For an asynchronous mirror policy, select the **Transfer All Source Snapshot Copies** check box to include the "all_source_snapshots" rule to the mirror policy, which backs up all of the Snapshot copies from the source volume.
6. Optional: Select the **Enable Network Compression** check box to compress the data that is being transferred during a data transfer.
7. Optional: Click **Add Comments** to add additional comments for the policy.
8. For a vault policy or mirror vault policy, specify a SnapMirror label and a destination retention count.
9. Click **Create**.

## Protection Policies window

You can use the Protection Policies window to create, manage, and display information about mirror, vault, and mirror vault policies.

- *Command buttons* on page 340
- *Protection policies list*
- *Details area*

### Command buttons

**Create**

Opens the Create Policy dialog box, which enables you to create a mirror, vault, or mirror vault policy.

**Edit**

Opens the Edit Policy dialog box, which enables you to edit a policy.

**Delete**

Opens the Delete Policy dialog box, which enables you to delete a policy.

**Refresh**

Updates the information in the window.

### Protection policies list

**Name**

Displays the name of the protection policy.

**Type**

Displays the policy type, which can be Vault, Mirror Vault, or Asynchronous Mirror.

**Comment**

Displays the description specified for the policy.

**Transfer Priority**

Displays the data transfer priority, such as Normal or Low.

### Details area

**Policy Details tab**

Displays details of the protection policy, such as the user who created the policy, number of rules, retention count, and status of network compression.

**Policy Rules tab**

Displays details of the rules that are applied to the policy. The Policy Rules tab is displayed only if the selected policy contains rules.

# Snapshot policies

You can use System Manager to create and manage Snapshot policies in your storage system.

## Creating Snapshot policies

You can create a Snapshot policy in System Manager to specify the maximum number of Snapshot copies that can be automatically created and the frequency of creating them.

### Steps

1. Click **Protection** > **Snapshot Policies**.
2. Click **Create**.
3. In the **Create Snapshot Policy** dialog box, specify the policy name.
4. Click **Add**, and then specify the schedule name, the maximum number of Snapshot copies that you want to retain, and the SnapMirror label name.

   The maximum number of Snapshot copies that can be retained by the specified schedules must not exceed 254.
5. Click **OK**, and then click **Create**.

## Editing Snapshot policies

You can modify the details of an existing Snapshot policy, such as the schedule name, SnapMirror label, or the maximum number of Snapshot copies that are created by using the Edit Snapshot Policy dialog box in System Manager.

### Steps

1. Click **Protection** > **Snapshot Policies**.
2. In the **Snapshot Policies** window, select the Snapshot policy that you want to modify and click **Edit**.
3. In the **Edit Snapshot Policy** dialog box, select the schedule that you want to modify and click **Edit**.
4. Click **OK**.
5. Verify the changes you made to the selected Snapshot policy in the **Edit Snapshot Policy** dialog box and click **Save**.

## Deleting Snapshot policies

You can use System Manager to delete Snapshot policies. If you delete a Snapshot policy that is being used by one or more volumes, Snapshot copies of the volume or volumes are no longer created according to the deleted policy.

### Before you begin

You must have dissociated the Snapshot policy from each volume that uses it.

### Steps

1. Click **Protection** > **Snapshot Policies**.
2. Select the Snapshot policy and click **Delete**.
3. Select the confirmation check box, and then click **Delete**.

## About Snapshot policies

When applied to a volume, a Snapshot policy specifies a schedule or schedules according to which Snapshot copies are created and specifies the maximum number of Snapshot copies that each schedule can create. A Snapshot policy can include up to five schedules.

For vault relationships, the SnapMirror Label attribute is used to select Snapshot copies on the source volumes. Only Snapshot copies with the labels configured in the vault policy rules are replicated in backup vault operations. The Snapshot policy assigned to the source volume must include the SnapMirror Label attribute.

## Snapshot Policies window

You can use the Snapshot Policies window to manage Snapshot policy tasks, such as adding, editing, and deleting Snapshot policies.

### Command buttons

**Create**

Opens the Create Snapshot Policy dialog box, which enables you to add backup schedules and specify the maximum number of Snapshot copies to be retained in a policy.

**Edit**

Opens the Edit Snapshot Policy dialog box, which enables you to modify the frequency at which Snapshot copies should be created and the maximum number of Snapshot copies to be retained.

**Delete**

Opens the Delete dialog box, which enables you to delete the selected Snapshot policy.

**View as**

Enables you to view the Snapshot policies either as a list or as a tree.

**Status**

Opens the menu, which you can use to either enable or disable the selected Snapshot policy.

**Refresh**

Updates the information in the window.

### Snapshot policy list

**Policy/Schedule Name**

Specifies the name of the Snapshot policy and the schedules in the policy.

**Storage Virtual Machine**

Specifies the name of the storage virtual machine (SVM) to which the Snapshot copies belong.

**Status**

Specifies the status of the Snapshot policy, which can be Enabled or Disabled.

**Maximum Snapshots to be retained**

Specifies the maximum number of Snapshot copies to be retained.

**SnapMirror Label**

Specifies the name of the SnapMirror label attribute of the Snapshot copy generated by the backup schedule.

# Schedules

You can use System Manager to create and manage schedules in your storage system.

## Creating schedules

You can create schedules to run a job at a specific time or at regular periods by using System Manager.

### About this task

When you create a schedule in a MetroCluster configuration, it is a best practice to create an equivalent schedule on the cluster in the surviving site as well.

### Steps

1. Click **Protection** > **Schedules**.
2. Click **Create**.
3. In the **Create Schedule** dialog box, specify the schedule name.
4. Create a schedule based on your requirements:

| If you want to create... | Do this... |
|---|---|
| A daily or a specific schedule on certain days | Select **Basic**, and specify the schedule and recurrence details (in hours and minutes). |
| A schedule that runs at a specific interval | Select **Interval**, and specify the schedule and recurrence details (in days, hours, and minutes). |
| A schedule that runs at a specific period | Select **Advanced**, and specify the schedule and recurrence details (in months, days, weekdays, hours, and minutes). |

5. Click **Create**.

## Editing schedules

You can make changes to a previously created cron schedule or an interval schedule if it does not meet your requirements by using System Manager. You can modify schedule details such as recurring days and hours, interval options, and advanced cron options.

### About this task

When you edit a schedule in a MetroCluster configuration, it is a best practice to edit the equivalent schedule on the surviving site cluster as well.

### Steps

1. Click **Protection** > **Schedules**.
2. Select the schedule that you want to modify and click **Edit**.
3. In the **Edit Schedule** dialog box, modify the schedule by performing the appropriate action:

| If you select the schedule option as... | Do this.. |
|---|---|
| Basic | Specify the recurring days and recurring schedule details. |
| Interval | Specify the interval options in days, hours, and minutes. |
| Advanced | Specify the advanced cron options in months, days, week days (if applicable), hours, and minutes. |

4. Click **OK**.

## Deleting schedules

You can use System Manager to delete the schedules that run specific storage management tasks.

### Steps

1. Click **Protection** > **Schedules**.
2. Select the schedule that you want to delete and click **Delete**.
3. Select the confirmation check box, and then click **Delete**.

## Schedules

You can configure many tasks (for instance, volume Snapshot copies and mirror replications) to run on specified schedules. Schedules that are run at specified schedules are known as *cron* schedules because of their similarity to UNIX `cron` schedules. Schedules that are run at intervals are known as *interval* schedules.

You can manage schedules in the following ways:

- Creating a cron schedule or an interval schedule
- Displaying information about all the schedules
- Modifying a cron schedule or an interval schedule
- Deleting a cron schedule or an interval schedule
  You cannot delete a schedule that is currently in use by a running job.

The cluster administrator can perform all the schedule management tasks.

## Schedules window

You can use the Schedules window to manage scheduled tasks, such as creating, displaying information about, modifying, and deleting schedules.

### Command buttons

**Create**

Opens the Create Schedule dialog box, which enables you to create time-based and interval schedules.

**Edit**

Opens the Edit Schedule dialog box, which enables you to edit the selected schedules.

**Delete**

Opens the Delete Schedule dialog box, which enables you to delete the selected schedules.

**Refresh**

Updates the information in the window.

### Schedules list

**Name**

Specifies the name of the schedule.

**Type**

Specifies the type of the schedule—time-based or interval-based.

### Details area

The details area displays information about when a selected schedule is run.

# Copyright, trademark, and machine translation

## Copyright

Copyright © 2021 NetApp, Inc. All rights reserved. Printed in the U.S.

No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

Data contained herein pertains to a commercial item (as defined in FAR 2.101) and is proprietary to NetApp, Inc. The U.S. Government has a non-exclusive, non-transferrable, non-sublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b).

## Trademark

NETAPP, the NETAPP logo, and the marks listed on the NetApp Trademarks page are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.

*https://www.netapp.com/company/legal/trademarks/*

## Machine translation

See important information about localized content at netapp.com.

*https://www.netapp.com/company/legal/machine-translation/*