



Cluster Management using ThinkSystem Storage Manager for DM Series



Version 9.9.1

Second Edition (March 2023)

© Copyright Lenovo 2021, 2023.

LIMITED AND RESTRICTED RIGHTS NOTICE: If data or software is delivered pursuant to a General Services Administration (GSA) contract, use, reproduction, or disclosure is subject to restrictions set forth in Contract No. GS-35F-05925

Contents

Contents	i
---------------------------	----------

Chapter 1. Understanding Storage Manager	1
---	----------

Chapter 2. Icons used in the application interface	3
---	----------

Chapter 3. Storage Manager window layout	5
---	----------

Chapter 4. Enhancements to ThinkSystem Storage Manager for DM Series	7
---	----------

Chapter 5. Setting up your cluster environment	11
---	-----------

Setting up the cluster by using ThinkSystem Storage Manager for DM Series	11
Setting up the cluster manually	13
Configuring Storage Manager options	16
Viewing log files of ThinkSystem Storage Manager for DM Series	16
How system logging works	17
Configuring a cluster by using Storage Manager	17
Setting up the cluster	17
Setting up the network	19
Setting up physical storage	20
Setting up logical storage	21

Chapter 6. Setting up SAML authentication	27
--	-----------

Enabling SAML authentication	27
Disabling SAML authentication	28

Chapter 7. Setting up peering	29
--	-----------

Prerequisites for cluster peering	29
Creating intercluster LIFs	30
Creating cluster peer relationships	30
Creating SVM peers	31
What passphrases are	31
Enabling SnapMirror Business Continuity	32

Chapter 8. Managing clusters	33
---	-----------

Understanding quorum and epsilon	33
Dashboard window	34
Monitoring a cluster using the dashboard	35
Applications	35

Storage service definitions	35
Configuration update	36
Configuring the administration details of an SVM	36
Service Processors	36
Isolating management network traffic	36
Assigning IP addresses to Service Processors	37
Editing Service Processor settings	38
Understanding the Service Processor	38
Cluster peers	38
Generating a peering passphrase	39
Modifying the cluster peer passphrase	39
Modifying LIFs that are configured for the remote cluster	39
Deleting cluster peer relationships	39
Cluster Peers window	40
Licenses	41
License types and entitlement risk	41
Updating clusters	42
Updating clusters in a non MetroCluster configuration	42
Updating clusters in a MetroCluster configuration	43
Obtaining ONTAP software images	44
Updating a cluster nondisruptively	45
How to update a cluster nondisruptively	46
SNMP	47
Enabling or disabling SNMP	47
Editing SNMP information	47
Enabling or disabling SNMP traps	47
SNMP window	48
LDAP	48
Viewing the LDAP client configuration	48
Using LDAP services	49
LDAP window	49
Users and Roles	50
Adding a cluster user account	50
Editing a cluster user account	50
Changing passwords for cluster user accounts	50
Locking or unlocking cluster user accounts	51
User accounts (cluster administrators only)	51
Users and Roles window	51

Chapter 9. Managing the network	53
--	-----------

IPspaces	53
--------------------	----

Editing IPspaces.	53
Deleting IPspaces	53
Broadcast domains	53
Editing broadcast domain settings.	53
Deleting broadcast domains	54
Network interfaces	54
Creating network interfaces	54
Editing network interface settings	55
Deleting network interfaces	55
Migrating a LIF	56
Ethernet ports	56
Creating interface groups	56
Creating VLAN interfaces	56
Editing interface group settings	57
Deleting VLANs	57
Ports and adapters.	57

Chapter 10. Managing physical storage 59

Storage tiers	59
Renaming a Storage Tier	59
Deleting aggregates	59
Moving FlexVol volumes	59
Viewing aggregate information	60
How moving a FlexVol volume works.	60
How you can use effective ONTAP disk type for mixing HDDs.	61
What compatible spare disks are	61
How Storage Manager works with hot spares.	61
Rules for displaying disk types and disk RPM	62
What a FabricPool is	62
Storage recommendations for creating aggregates	62
External capacity tier.	63
Configuring and managing cloud tiers	63
Aggregates	67
Disks	67
Viewing disk information	67
How ONTAP reports disk types	67
Minimum number of hot spares required for disks	68
Considerations for sizing RAID groups	69
About disks	69
Events	69
Events window	69
System alerts.	71
Acknowledging system health alerts	71
Suppressing system health alerts	71
Available cluster health monitors	71

Ways to respond to system health alerts	72
System Alerts window	72
AutoSupport notifications	73
Enabling or disabling AutoSupport settings	73
Adding AutoSupport email recipients	73
Testing AutoSupport settings.	73
Generating AutoSupport data	74
AutoSupport severity types	74
AutoSupport window	74
Jobs	75
Jobs	75

Chapter 11. Managing logical storage 77

Storage Virtual Machines	77
Editing SVM settings	77
Deleting SVMs	77
Starting SVMs	78
Stopping SVMs	78
Managing SVMs	79
Types of SVMs	79
Why you use SVMs	79
How ONTAP name service switch configuration works	80
Trace File Access window	81
Volumes.	82
Editing volume properties	83
Editing data protection volumes.	83
Deleting volumes	84
Creating FlexClone volumes	84
Splitting a FlexClone volume from its parent volume	85
Setting the Snapshot copy reserve	85
Scheduling automatic creation of Snapshot copies.	86
Extending the expiry date of Snapshot copies.	86
Resizing volumes	86
Moving FlexVol volumes between aggregates or nodes.	87
Assigning volumes to Storage QoS	87
Creating a mirror relationship from a source SVM	88
Creating a vault relationship from a source SVM	89
Creating a mirror and vault relationship from a source SVM	90
Changing the tiering policy of a volume.	91
Creating FlexGroup volumes	91
Viewing FlexGroup volume information	91
Editing FlexGroup volumes	92
Resizing FlexGroup volumes	92

Deleting FlexGroup volumes	92	Quotas	109
What Volume Encryption is	93	Creating quotas	109
Snapshot configuration	93	Deleting quotas	109
How volume guarantees work for FlexVol volumes	93	Editing quota limits.	109
FlexClone volumes and space guarantees	94	Types of quotas	110
Thin provisioning for greater efficiencies using FlexVol volumes	94	Quota limits	110
Using space reservations with FlexVol volumes	95	Quota management	111
Benefits of storage efficiency.	95	CIFS protocol	111
Data compression and deduplication	96	Setting up CIFS	111
Guidelines for using deduplication.	96	Editing the name for a CIFS SVM	112
Considerations when moving volumes	97	Adding home directory paths.	112
Shares	97	Deleting home directory paths	112
Creating a CIFS share	97	NFS protocol	113
Creating home directory shares.	98	Editing NFS settings	113
Editing share settings.	98	NVMe protocol	113
How ONTAP enables dynamic home directories	98	Setting up NVMe	113
LUNs	100	Creating an NVMe namespace	114
Creating LUNs	100	Editing an NVMe namespace.	114
Deleting LUNs	100	Cloning an NVMe namespace	115
Creating initiator groups.	101	What NVMe is.	115
Deleting initiator groups.	101	What an NVMe subsystem is	116
Adding initiators.	101	NVMe namespaces	116
Deleting initiators from an initiator group	101	iSCSI protocol	117
Cloning LUNs.	102	Starting or stopping the iSCSI service	117
Editing LUNs	102	FC/FCoE protocol	117
Moving LUNs	103	Starting or stopping the FC or FCoE service	118
Editing initiator groups	103	Export policies	118
Viewing LUN information	103	Creating an export policy	118
Viewing initiator groups	104	Renaming export policies	118
Guidelines for working with FlexVol volumes that contain LUNs	104	Deleting export policies	119
Understanding space reservations for LUNs	105	Adding rules to an export policy.	119
Guidelines for using LUN multiprotocol type	105	Modifying export policy rules	119
Understanding LUN clones	106	Deleting export policy rules	120
Initiator hosts	106	How export policies control client access to volumes or qtrees	120
igroup name	106	QoS policy groups.	120
igroup type.	107	Creating QoS policy groups	120
igroup ostype.	107	Managing workload performance by using Storage QoS	121
Qtrees	107	How Storage QoS works	121
Creating qtrees	107	How the maximum throughput limit works.	122
Deleting qtrees	107	Rules for assigning storage objects to policy groups	123
Editing qtrees.	107	LDAP client services	124
Assigning export policies to qtrees	108	Adding an LDAP client configuration	124
Viewing qtree information	108	Deleting an LDAP client configuration	124
Qtree options	108	Editing an LDAP client configuration	125
		LDAP configuration services	125
		Editing active LDAP clients	125
		Deleting active LDAP clients	125

LDAP Configuration window	126
DNS Services	126
Enabling or disabling DNS	126

Chapter 12. Managing data protection127

Mirror relationships	127
Creating a mirror relationship from a destination SVM	127
Deleting mirror relationships	128
Editing mirror relationships	128
Updating mirror relationships.	129
Quiescing mirror relationships	130
Resuming mirror relationships	130
Breaking SnapMirror relationships.	130
Resynchronizing mirror relationships.	131
Reverse resynchronizing mirror relationships	132

Aborting a mirror transfer	132
Restoring a volume in a mirror relationship	133
How SnapMirror relationships work	133
What lag time is.	134
Types of data protection relationships	134
Understanding workloads supported by StrictSync and Sync policies	135
SnapMirror licensing	136
SVM Relationships	136
Creating SVM relationships	136
Editing SVM relationships	137
Managing SVM relationships	137

Appendix A. Contacting Support . . .139

Appendix B. Notices.141

Trademarks	142
----------------------	-----

Chapter 1. Understanding Storage Manager

Storage 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 Storage Manager to administer the entire cluster and its resources.

Storage 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.
- Create S3 Object Store bucket

Chapter 2. 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.

Chapter 3. Storage Manager window layout

Understanding the typical window layout helps you to navigate and use Storage Manager effectively.

Typical layout of Storage Manager windows



Chapter 4. Enhancements to ThinkSystem Storage Manager for DM Series

You should be aware of the features that have been added or changed in this release of Storage Manager.

Features and enhancements added in ONTAP 9.9.1

- Full Support for SnapMirror Synchronous relationships
- Enhanced Error and Warning messages in Dashboard
- Hardware visualization support in System Manager
- Fast directory delete with System Manager
- SVM DR support for FlexGroup volumes
- Single LUN performance enhancement

Features and enhancements added in ONTAP 9.8

- Elimination of classic GUI
- S3 Object Store may now be configured from Storage Manager
- SnapMirror Business Continuity added under Data Protection tab
- Support for ONTAP Mediator from ThinkSystem Storage Manager
- Support for NFS 4.2

Features and enhancements added in ONTAP 9.7

- New simplified Management screen and processes
- Added tier options in place of aggregate view

Features and enhancements added in ONTAP 9.6

- MetroCluster switchover and switchback operations

Starting with Storage Manager 9.6, you can use MetroCluster switchover and switchback operations to allow one cluster site to take over the tasks of another cluster site. This capability allows you to facilitate maintenance or recovery from disasters.

- Trace file access

A new Trace File Access window allows you to diagnose issues when users have problems accessing files and shares on an SVM using the CIFS or NFS protocol.

- Encrypted SnapMirror

Starting with Storage Manager 9.6, you can generate a passphrase for the local cluster IPspace and use the same passphrase in the remote cluster when creating peering relationships. For security reasons, the passphrase can be modified.

- FlexCache capabilities

Storage Manager now provides the capability to create, edit, view, and delete FlexCache volumes.

- Fabric Pool enhancements

You can now use Alibaba Cloud and Google Cloud as object stores for use as a cloud tier. Also, Storage Manager now supports the “All” tiering policy, which allows you to specify that all data should be tiered to the cloud. Enhancements were made to the Volume Performance tab of the Volume 360 page to show cloud latency of the volume.

- FlexGroup enhancements

Starting with Storage Manager 9.6, you can edit the properties of an existing FlexGroup volume, such as renaming or resizing the volume.

- SnapLock volume enhancements

When creating or modifying a storage QoS policy group, you can set the minimum throughput limit for an ONTAP Select Premium system in addition to a performance-based All Flash Optimized personality.

Features and enhancements added in ONTAP 9.5

- Volume encryption

You can now enable volume encryption while editing a FlexVol volume or a FlexGroup volume. Also, this feature is enhanced to support the Rekey option to change the data encryption key of the volume.

- Cluster update

Beginning with Storage Manager 9.5, you can update a cluster in MetroCluster configurations. You must perform each operation on both the clusters except for updating the cluster.

- Volume replication policies

Two new policies, StrictSync and Sync, are added in Storage Manager 9.5. You can use these policies to provide zero RPO replication with and without primary IO restriction during replication failures. You can also enable volume protection using the protection tab.

- Cloud Registration

You can use Storage Manager to register the ONTAP cluster with Lenovo Data Availability Services to save data in the cloud.

- SVM DR

Storage virtual machine (SVM) disaster recovery (DR) provides disaster recovery capability at the SVM level by enabling the recovery of the data that is present in the constituent volumes of the SVM and the recovery of the SVM configuration. You can use Storage Manager to create and manage mirror relationships and mirror and vault relationships between SVMs.

- L2/L3 applications displayed

Starting with ONTAP 9.5, Storage Manager lists L2/L3 applications on the Applications page under different host names. Clicking on the host name opens a new window in the L2 Cockpit interface. For each application, Storage Manager also lists IOPs and latency measurements.

- Virtual IP support

Starting with ONTAP 9.5, Storage Manager displays information about Virtual IP (VIP) LIFs; however, you cannot create, delete, or manage VIP LIFs from Storage Manager.

- NVMe subsystems licensing requirement

Starting with ONTAP 9.5, NVMe is licensed. Storage Manager supports the licensing requirement.

- Support for NVMeoF subsystems

Storage Manager supports the use of an NVMe over Fabric (NVMeoF) subsystem, which is a separate kernel object that resides in the FreeBSD kernel. 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. An NVMeoF subsystem interfaces with SAN components, WAFL, and RAS components.

- NVMe multipath support

Starting with ONTAP 9.5, at least one LIF must be configured for each node in an HA pair using the NVMe protocol. You can also define two LIFs for a node. When you upgrade to ONTAP 9.5, you must ensure that a minimum of one LIF is defined for each node in an HA pair using the NVMe protocol.

- FlexGroup eligible aggregates

When you create a FlexGroup, aggregates are selected by default according to best practices. For All-Flash Optimized storage systems, thin provisioning is enabled by default, and for other storage systems, thick provisioning is enabled by default. You can override the best practices defaults and select your choices from a list of eligible FabricPool aggregates.

- Public SSL Certificate authentication

Starting with Storage Manager 9.5, you can view a public SSL certificate associated with an SVM. You can view the certificate details, the serial number, the start date, and the expiration date. You can also copy the certificate to the clipboard, and email the certificate details. Additionally, when you add the vsadmin user account to an SVM, a login method is automatically included that uses HTTP as the application and is authenticated with a certificate.

- Qtrees appearing as directories on a FlexVol

If a FlexVol contains both qtrees and volumes, the qtrees appear as directories.

- FlexCache volumes

FlexCache volumes are displayed in Storage Manager as a FlexGroup. The parent volume details are shown in the 360 page.

- Deprecation of infinite volumes

Starting with ONTAP 9.5, infinite volumes are deprecated and no longer supported.

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 Storage 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. Storage 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 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 Storage Manager Volumes window.
- Removed partial support for infinite volumes

You cannot create infinite volumes and protect infinite volumes by using Storage Manager.
- 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.

Chapter 5. Setting up your cluster environment

You can create a cluster by using Storage Manager or the command-line interface (CLI). To create a cluster by using Storage 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 Storage Manager.



Setting up the cluster by using ThinkSystem Storage Manager for DM Series

Beginning with ONTAP 9.5, you can use ThinkSystem Storage Manager for DM Series to set up a cluster by creating a cluster, setting up the node management network and cluster management network, and then setting up 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 9.5 or later versions.
- 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 ThinkSystem Storage Manager for DM Series.

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



Setting up the cluster manually

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

Creating a cluster

You can use ThinkSystem Storage Manager for DM Series to create and set up a cluster in your data center.

About this task

If the cluster supports ONTAP 9.5 or later versions, you can add only those storage systems that are running ONTAP 9.5 or later versions.

- Step 1. Open the web browser, and then enter the node management IP address that you have configured:
`https://node-management-IP`
- Step 2. Enter a cluster name and administrator password.
- Step 3. Assign your cluster IP address and subnet mask with the second node IP address in the networking section.
- Step 4. Choose if you want to enable DNS, sending of AutoSupport data and if you want to use NTP.
- Step 5. Click **Submit**.

After you finish

You are ready to launch Storage Manager for DM Series.

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 Storage Manager.

To create data aggregates as per the storage recommendations, click **Submit and Continue**.

To create data aggregates at a later time using Storage Manager, click **Skip this step**.

After you finish

If you opted to create aggregates per the storage recommendations, you should 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.

Step 1. Enter a name for the SVM.

Step 2. Select data protocols for the SVM:

If you want to...	Then...
Enable CIFS protocol by configuring the CIFS server using an Active Directory	1. Select the Active Directory box. 2. Enter the Active Directory administrator name. 3. Enter the Active Directory administrator password. 4. Enter a name for the CIFS server. 5. Enter a name for the Active Directory domain. 6. Depending on your requirements, select the One data LIF on this SVM or One data LIF per node on this SVM box. 7. Provide data LIF details such as IP address, netmask, gateway, and port. 8. Provide DNS details.
Enable CIFS protocol by configuring the CIFS server using a workgroup	1. Select the Workgroup box. 2. Enter a name for the workgroup. 3. Enter a name for the CIFS server. 4. Depending on your requirements, select the One data LIF on this SVM or One data LIF per node on this SVM check box. 5. Provide data LIF details such as IP address, netmask, gateway, and port.
Enable NFS protocol	1. Select the NFS box. 2. Depending on your requirements, select the One data LIF on this SVM or One data LIF per node on this SVM check box. 3. Provide data LIF details such as IP address, netmask, gateway, and port.
Enable iSCSI protocol	1. Select the iSCSI box. 2. Provide data LIF details such as IP address, netmask, gateway, and port.

If you want to...	Then...
Enable FC/FCoE protocol	1. Select the FC/FCoE box. 2. Select the FC/FCoE ports for FC or FCoE protocols. Note: Each node must have at least one correctly configured port for each protocol (FC and FCoE).
Enable NVMe protocol	1. Select the NVMe box. 2. Select the NVMe ports for NVMe protocols. Note: At least one NVMe capable adapter must be available in each of the nodes configured for NVMe. Also, starting with ONTAP 9.5, at least one NVMe LIF must be configured for each node of a HA pair associated with the SVM. You can create a maximum of two NVMe LIFs for each node in the pair.

Step 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.

Step 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 Storage Manager, or click **Provision an Application** to provision storage applications, or click **Export Configuration** to download the configuration file.

Configuring Storage Manager options

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

About this task

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

Step 1. Click **Cluster → Settings**.

Step 2. In the **UI Settings** card, click the pencil icon .

Step 3. Specify the log level, inactivity timeout, auto refresh interval and global search interval.

Step 4. Click **Update**.

Viewing log files of ThinkSystem Storage Manager for DM Series

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

- Step 1. Identify the node that hosts the cluster management LIF.
- Step 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.
- Step 3. Type your cluster administrator credentials, and then click **OK**.
- Step 4. In the Data ONTAP - Root Volume File Access window, click the **logs** link for the node that hosts the cluster management LIF.
- Step 5. Navigate to the `mlog` directory to access the Storage 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 Storage 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 Storage Manager logging at runtime without modifying the application binary.

Log output can be voluminous and therefore can become difficult to manage. Storage 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 Storage Manager

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

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

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 Storage Manager to modify the name of a cluster when required.

- Step 1. Click **Cluster → Overview**.

- Step 2. In the **Overview** card, click the more icon  .

- Step 3. Select **Rename**.
- Step 4. Type the value in the **CLUSTER NAME** field.
- Step 5. Click **Save**.

Changing the cluster password

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

- Step 1. Click **Cluster → Settings**.
- Step 2. In the **Users and Roles** card, click the arrow icon .
- Step 3. Choose the user account whose password that you wish to change and then click the more icon  next to it.
- Step 4. Select **Change Password**.
- Step 5. Type the value in the **NEW PASSWORD** field and confirm it in the **CONFIRM PASSWORD** field.
- Step 6. Click **Save**.

Editing DNS configurations

You can use Storage 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.

- Step 1. Click **Cluster → Overview**.
- Step 2. In the **Overview** card, click the more icon .
- Step 3. Select **Edit**.
- Step 4. In the **DNS DOMAINS** field, add or modify the DNS domain name.
- Step 5. In the **NAME SERVERS** field, add or modify the IP addresses of the name servers to assign address to a node.
- Step 6. Click **Save**.

Creating a cluster management logical interface

You can use Storage Manager to create a cluster management logical interface (LIF) to provide a single management interface for a cluster. You can use this LIF to manage all of the activities of the cluster.

- Step 1. Click **Cluster → Overview**.
- Step 2. In the **Overview** card, click the more icon .
- Step 3. Select **Edit**.
- Step 4. Select **Add cluster management interface**.
- Step 5. Select a node to host the cluster management address.
- Step 6. Type values in the **IP ADDRESS** and **SUBNET MASK** fields.
- Step 7. Click **Save**.

Editing the node name

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

- Step 1. Click **Cluster → Overview**.
- Step 2. In the **Nodes** card, click the name of the node that you want to modify.

Step 3. Type the new name and press **Enter**.

Editing AutoSupport settings

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

Step 1. Click **Cluster → Settings**.

Step 2. In the **AutoSupport** card, click the more icon .

Step 3. Select **More options**.

Step 4. Move the toggle button to enable or disable the sending of AutoSupport data.

- If you choose to disable it, the Disable AutoSupport dialog box is displayed and you need to click **Disable**.
- If you choose to enable it, you can also perform the following operations:
 - In the **Connections** card, click **Edit** to select either SMTP, HTTPS, or HTTP for sending the data, and then click **Save**.
 - In the **Email** card, click **Edit** to choose to send the data to another location that can be configured, and then click **Save**.

Adding licenses

Your storage system software was installed at the factory. Storage Manager automatically adds the software to its list of licenses.

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 Storage Manager to add the Cloud Volumes ONTAP license. The Cloud Volumes ONTAP license is not listed in the license page. Storage Manager does not raise any alert about the entitlement risk status of the Cloud Volumes ONTAP license.

Step 1. Click **Cluster → Settings**.

Step 2. In the **Licenses** card, click the arrow icon .

Step 3. Click **+Add**.

Step 4. Either select **Use 28 character license keys** and type the value in the **LICENSE KEYS** field, or select the key file generated from the LKMS by clicking **Browse**.

Step 5. Click **Save**.

Result

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

Step 1. Click **Network → Overview**.

Step 2. In the IPspaces card, click the plus icon .

Step 3. Specify the name of the new IPspace.

Step 4. Click **Save**.

Creating broadcast domains

You can create a broadcast domain by using Storage 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.

Step 1. Click **Network → Overview**.

Step 2. In the **Broadcast Domains** card, click the plus icon .

Step 3. Select the name for the new broadcast domain and the IPspace that it will be placed in.

Step 4. Specify the MTU size to use and the ports assigned to the broadcast domain.

Step 5. Click **Save**.

Setting up physical storage

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

Creating a local storage tier

You can create a local tier based on storage recommendations.

Before you begin

You must have enough spare disks to create an aggregate.

About this task

You cannot perform the following actions by using Storage 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.

Provisioning storage by creating a local tier based on storage recommendations

You can use Storage Manager to create a local tier based on storage recommendations. Storage 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 MetroCluster configurations.
- Errors, if any, are displayed on the screen.

Step 1. Click **Storage**.

Step 2. Select tiers.

Step 3. Click **+Add Local Tier**.

Step 4. Click **Save**.

Setting up logical storage

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

Creating SVMs

You can use Storage 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.

- Step 1. Click **Storage → Storage VMs**.
- Step 2. Click **+Add**.
- Step 3. Specify **Storage VM Name** and the protocol that the SVM will support.
- Step 4. Click **Save**.

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 Storage 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 or a URI keytab file.

About this task

SnapLock aggregates are not considered for automatically creating volumes.

- Step 1. Click **Storage → Storage VMs**.
- Step 2. Click **+Add**.
- Step 3. Specify **Storage VM Name** and choose either CIFS or NFS protocol that the SVM will support.
- Step 4. When prompted, provide an account to add the system to a domain or workgroup.
- Step 5. Provide the password, System name, and Active Directory domain to add the new SVM into.
- Step 6. Specify the DNS domain name and name servers to use.
- Step 7. Configure the LIFs that will be used to service the CIFS data.
- Step 8. Click **Save**.

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 Storage 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.

Step 1. Click **Storage → Storage VMs**.

Step 2. Click **+Add**.

Step 3. Specify **Storage VM Name**.

Step 4. Select iSCSI from the protocol list and select **Enable iSCSI**.

Step 5. Create at least one LIF on each node.

Step 6. Click **Save**.

Result

The data LIFs and port sets are created with the specified configuration if you select the option to configure the iSCSI protocol. 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 Storage 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.

About this task

- You can configure the FC protocol while creating the SVM or you can configure the protocol 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.

Step 1. Click **Storage → Storage VMs**.

Step 2. Click **+Add**.

- Step 3. Specify **Storage VM Name** and select **enable FC** for the SVM.
- Step 4. When prompted, provide the FibreChannel Ports to use.
- Step 5. Click **Save**.

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.

Configuring NVMe protocol on SVMs

You can configure the NVMe protocol on a storage virtual machine (SVM) using Storage 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.

- Step 1. Click **Storage → Storage VMs**.
- Step 2. Click **+Add**.
- Step 3. Select a name for the new virtual machine and click **NVMe\FC**.
- Step 4. Enable **NVMe\FC**.
- Step 5. Select the FC ports to use.
- Step 6. Click **Save**.

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.

- Step 1. Click **Storage → Storage VMs**.
- Step 2. Click the storage VM to modify.
- Step 3. Click the **Settings** tab on the displayed right pane.
- Step 4. Scroll down to locate the **Users and Roles** card and click the arrow icon .
- Step 5. Move the cursor onto the **vsadmin** account and then click the more icon .
- Step 6. Select **Change Password**.
- Step 7. Set the new password and confirm it.
- Step 8. Move the cursor onto the **vsadmin** account and then click the more icon .
- Step 9. Select **Unlock**.
The user is unlocked.
- Step 10. Choose **Network → Overview** on the left navigation pane.
- Step 11. In the Network Interfaces card, click the plus icon .

Step 12. Specify the storage virtual machine and assign a new IP address, subnet mask, and gateway to the virtual machine.

Step 13. Click **Save**.

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 volumes

You can create a volume for your data by using the Volumes dialog box in Storage Manager. You must 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 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 Storage Manager, and you must have enabled “key-manager setup” by using the command-line interface (CLI).

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

- Volumes can only be created for CIFS or NFS enabled Storage virtual machines.

About this task

- You can enable storage Quality of Service (QoS) only for a read/write (rw) volume.

Step 1. Click **Storage → Volumes**.

Step 2. Select **+Add**.

Step 3. Select a name for the new volume, Virtual machine to assign it to and a capacity.


Notes:

- If you select more options you can also set the volume to be distributed across the cluster (FlexGroup).
- You can also enable access permissions for the share or export and Protection options including Snapshot local copies and SnapMirror.

Step 4. Click **Save**.

Chapter 6. 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 Storage Manager.



Enabling SAML authentication

You can use Storage 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 Storage Manager are Shibboleth and Active Directory Federation Services.

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

Step 1. Click **Cluster → Settings**.

Step 2. In the **SAML Authenticaion** card, click the gear icon .

Step 3. Select **Enable SAML Authentication**.

Step 4. Enter the IdP URI and IP address of host system.

Step 5. Click **Save**.

Step 6. Log in to Storage 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, the system automatically changes the IP address to the IP address of the host system that was specified during the IdP configuration.

Disabling SAML authentication

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

About this task

Disabling SAML authentication does not delete SAML configuration.

Step 1. Click **Cluster → Settings**.

Step 2. In the **SAML Authenticaion** card, click the gear icon .


Step 3. Clear the **Enable SAML authentication** check box.

Step 4. Click **Save**.

Step 5. Log in to the Storage Manager by using cluster credentials.

Chapter 7. 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

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 belong to the broadcast domain that contains the ports that 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.

Creating intercluster LIFs

Creating intercluster logical interfaces (LIFs) enables the cluster network to communicate with a node. You must create an intercluster LIF within each IPspace that will be used for peering, on each node in each cluster for which you want to create a peer relationship.

Step 1. Click **Network → Overview**.

Step 2. In the **Network Interfaces** card, click the plus icon .

Step 3. Select the required Intercluster for the interface role.

Step 4. Select an IPspace for the new LIF.

Step 5. Assign a name, home node, IP address and subnet mask.

Step 6. Click **Save**.

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” on page 29](#)

- You must have created intercluster logical interfaces (LIFs).
- You should be aware of which version of ONTAP each cluster is running.

About this task

- You can create a peer relationship between a cluster running ONTAP 9.5 and a cluster running ONTAP 9.6. However, encryption is not supported in ONTAP 9.5, so the peer relationship cannot be encrypted.
- 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. However, the passphrases of both clusters must match.

Step 1. Click **Protection → Overview**.

Step 2. Check whether there are network interfaces listed under Intercluster Settings.

If there are no network interfaces listed under Intercluster Settings, go to setting up Intercluster LIFs.

Step 3. Go to **Cluster Peers**.

Step 4. Select **Peer Cluster**.

Step 5. Specify a user-defined passphrase or a system-generated passphrase and at least one intercluster LIF.

Step 6. Repeat steps 1 to 5 for the second cluster.

Step 7. Click **Initiate Cluster 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 **Protection → Overview** 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 Storage Manager.

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

Step 1. Click **Protection → Overview**.

Step 2. Select **Storage VM Peers** and click the more icon .

Step 3. Select **Peer Storage VMs**.

Step 4. Select the peer cluster from the list.

Step 5. Select the target storage VM.

Step 6. Click **Peer Storage VMs**.

After you finish

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

When you use Storage Manager to create the peer relationship, the encryption status is “Enabled” by default.

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.
- 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.

Enabling SnapMirror Business Continuity

Before you begin

You must first setup cluster and SVM peering first. You will also need to add the ONTAP Mediator before starting the following procedure.

- Step 1. Click **Protection → Overview**.
- Step 2. Select **Protect for Business Continuity**.
- Step 3. Select the LUNs that will be part of the new consistency group by checking each one from the **LUNs** field.
- Step 4. Specify a name for the new consistency group.
- Step 5. Select the destination cluster from the **CLUSTER** field.
- Step 6. Verify that the peer SVM is listed in the **Storage VM** field.
- Step 7. If you need to make any changes to the destination settings, click that option and alter the required fields.
- Step 8. Click **Save**.

Chapter 8. Managing clusters

You can use Storage Manager to manage clusters.


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 system health, performance, overall capacity and network configuration.

Cluster Health

Presents a graphical display of the overall system configuration.

It will display the model of the DM system that is installed including both nodes and all drives that are attached.

By highlighting any of the nodes, you can verify that they are online by looking for the green check box next to the name.

You can also see the overall health at the top of the panel where it says all systems are healthy. Any errors or warnings will also display in that text field.

If you click on the arrow at the top of the window, it will jump to the **Cluster → Hardware** status screen where you can see a graphical display of the system and the status for any devices.

Capacity

Displays the status of all configured tiers for the system. You will see the currently consumed capacity and the available capacity that has been allocated to a local tier.

If no tiers have been created, it will display the number of spare drives and the option to prepare storage.

If you select the arrow at the top of the window, it will launch the Storage -> Tiers screen and give you the ability to add a local or cloud tier.

This window will also display the overall data reduction for the system.

You can also see if a cloud tier is configured on the system.

A new feature in 9.9.1 is that you will also see the logical capacity consumed by the system.

Network

Displays all of the devices configured.

This includes:

- Storage VMs
- Volumes
- LUNs
- Hosts
- Ethernet\FC\NVMe ports
- FC\NFS\CIFS\iSCSI Interfaces

Clicking on any of these links will launch the corresponding panel

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, Storage 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 Storage 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 1. Click the **Dashboard** tab to view the health and performance dashboard panels.

Applications

Currently, you can not provision applications in Storage Manager 9.7 but it will display existing applications that were provisioned previously. It will also display applications provisioned using the CLI.

Storage service definitions

ONTAP includes predefined storage services that are mapped to corresponding minimum performance factors.

The actual set of storage services available in a cluster or SVM is determined by the type of storage that makes up an aggregate in the SVM.

The following table shows how the minimum performance factors are mapped to the predefined storage services:

Storage service	Expected IOPS (SLA)	Peak IOPS (SLO)	Minimum volume IOPS	Estimated latency	Are expected IOPS enforced?
value	128 TB	512 TB	75	17 ms	On AFA: Yes Otherwise: No
performance	2048 TB	4096 TB	500	2 ms	Yes
extreme	6144 TB	12288 TB	1000	1 ms	Yes

The following table defines the available storage service level for each type of media or node:

Media or node	Available storage service level
Disk	value
Virtual machine disk	value
FlexArray LUN	value
Hybrid	value
Capacity-optimized Flash	value

Media or node	Available storage service level
Solid-state drive (SSD) - non-AFA	value
Performance-optimized Flash - SSD (AFA)	extreme, performance, value

Configuration update

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

Configuring the administration details of an SVM

You can use Storage 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 Storage Manager to manage delegated SVMs. You can manage the SVMs only by using the command-line interface (CLI).

- Step 1. Click **Storage → Storage VMs**.
- Step 2. Click the required storage VM.
- Step 3. Click the **Settings** tab on the displayed right pane.
- Step 4. Scroll down to locate the **Users and Roles** card and click the arrow icon .
- Step 5. Move the cursor onto the **vsadmin** account and then click the more icon .
- Step 6. Select **Change Password**.
- Step 7. Set the new password and confirm it.
- Step 8. Move the cursor onto the **vsadmin** account and then click the more icon .
- Step 9. Select **Unlock**.
The user is unlocked.
- Step 10. Choose **Network → Overview** on the left navigation pane.
- Step 11. In the Network Interfaces card, click the plus icon .
- Step 12. Specify the storage virtual machine and assign a new IP address, subnet mask, and gateway to the virtual machine.
- Step 13. Click **Save**.


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 Storage Manager.

Isolating management network traffic

It is a best practice to configure SP/BMC and the e0M management interface on a subnet dedicated to management traffic. Running data traffic over the management network can cause performance degradation and routing problems.

The management Ethernet port on most storage controllers (indicated by a wrench icon on the rear of the chassis) is connected to an internal Ethernet switch. The internal switch provides connectivity to SP/BMC and to the e0M management interface, which you can use to access the storage system via TCP/IP protocols like Telnet, SSH, and SNMP.



If you plan to use both the remote management device and e0M, you must configure them on the same IP subnet. Since these are low-bandwidth interfaces, the best practice is to configure SP/BMC and e0M on a subnet dedicated to management traffic.

If you cannot isolate management traffic, or if your dedicated management network is unusually large, you should try to keep the volume of network traffic as low as possible. Excessive ingress broadcast or multicast traffic may degrade SP/BMC performance.

Note: Some storage controllers have two external ports, one for BMC and the other for e0M. For these controllers, there is no requirement to configure BMC and e0M on the same IP subnet.

Assigning IP addresses to Service Processors

You can use Storage 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.

Step 1. Click **Cluster → Overview**.

Step 2. Click the more icon  of the node where you wish to set the service processor address.

Step 3. Select **Edit Service Processor**.

Step 4. Select the IP address (DHCP or static IP).

If you select the static IP, you need to specify the IP address, subnet mask and gateway.

Step 5. 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 Storage 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 an IP address manually.
- You cannot edit the settings of a Service Processor that was assigned an IP address through a DHCP server or through a subnet.

Step 1. Click **Cluster → Overview**.

Step 2. Click the more icon  of the node where you wish to set the service processor address.

Step 3. Select **Edit Service Processor**.

Step 4. Select the IP address (DHCP or static IP).

If you select the static IP, you need to specify the IP address, subnet mask and gateway.

Step 5. Click **Save**.

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.

Cluster peers

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

Generating a peering passphrase

Starting with Storage Manager 9.6, you can generate a passphrase for the local cluster IPspace and use the same passphrase in the remote cluster when creating peering relationships.

- Step 1. Click **Protection → Overview**.
- Step 2. In the **Cluster Peers** area, click the more icon .
- Step 3. Select **Generate Passphrase**.
- Step 4. Select the IPspace, validity and remote cluster version.
- Step 5. Copy the generated passphrase.

Modifying the cluster peer passphrase

You can modify the passphrase that is provided during cluster peer creation.

- Step 1. Click **Protection → Overview**.
- Step 2. In the **Cluster Peers** area, click the more icon .
- Step 3. Select **Manage Cluster Peers**.
- Step 4. Select the required cluster peer and click the more icon .
- Step 5. Select **Update Passphrase**.
- Step 6. Specify the value in the passphrase and then click **Save**.

Modifying LIFs that are configured for the remote cluster

You can use Storage 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.

- Step 1. Click **Protection → Overview**.
- Step 2. In the **Cluster Peers** area, click the more icon .
- Step 3. Select **Manage Cluster Peers**.
- Step 4. Select the required cluster peer and click the more icon .
- Step 5. Select **Edit Peer Cluster Settings**.
- Step 6. Choose an IPspace.
- Step 7. Either edit the existing IP addresses or choose **+Add** to add a new IP address value.
- Step 8. Click **Save**.

Deleting cluster peer relationships

You can use Storage 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.

Step 1. Click **Protection → Overview**.

Step 2. In the **Cluster Peers** area, click the more icon .

Step 3. Select **Manage Cluster Peers**.

Step 4. Select the required cluster peer and click the more icon .

Step 5. Select **Delete the selected cluster peer relationships**.

Step 6. Select **Delete**.

Cluster Peers window

You can use the Cluster Peers window to manage peer cluster relationships, which enables 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.

Edit

Displays a drop-down menu with the following choices:

Local Cluster Passphrase

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

Peer Cluster Network Parameters

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

You can add multiple IP addresses, separated by commas.

Change Encryption

Opens the Change Encryption dialog box for the selected peer cluster. While you are changing the encryption of the peered relationship, you can either generate a new passphrase or provide a passphrase that was already generated at the remote peered cluster.

This action is not available if the encryption status is “N/A”.

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 the local cluster IPspace by specifying the IPspace, setting the passphrase validity duration, and specifying which SVMs are given permission.

You use the same passphrase in the remote cluster for 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 with the local cluster peer relationship.

Peer Cluster Intercluster IP Addresses

Displays IP addresses associated with the intercluster peer relationship.

Last Updated Time

Displays the time at which peer cluster was last modified.

Encryption

Displays the status of the encryption of the peering relationship.

Note: Starting with Storage Manager 9.6, peering is encrypted by default when you establish a peering relationship between two clusters

- **N/A:** Encryption is not applicable to the relationship.
- **none:** The peering relationship is not encrypted.
- **tls_psk:** The peering relationship is encrypted.

Licenses

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

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.

- 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 the node 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 (FabricPool only)

An ONTAP Select instance is licensed according to the amount of data that the user wants to manage. For example, the user might 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 9.5, FabricPool-enabled aggregates require 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 cloud tier 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 site license is installed on a cluster, the entitlement risk is never high.

- Medium risk

If a 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 of the nodes, or a site license is installed on the cluster, irrespective of usage.

- Unknown


The risk is unknown if the API is sometimes unable to retrieve the data related to entitlement risk that is associated with a cluster or the nodes in the cluster.

Updating clusters

You can use Storage Manager to update a cluster or the individual nodes in a high-availability (HA) pair. You can also update a cluster in a MetroCluster configuration.


Updating clusters in a non MetroCluster configuration

You can use Storage 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.



Updating clusters in a MetroCluster configuration

You can use Storage Manager to update a cluster in MetroCluster configurations. You must perform each operation on both the clusters except for updating the cluster.



Updating site A automatically updates site B.

Obtaining ONTAP software images

For ONTAP 9.5 and later, you can copy the ONTAP software image from the Lenovo Data Center Support site to a local folder.

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 Lenovo Data Center 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.
- If you are upgrading a system with Volume Encryption to ONTAP 9.5 or later, you must download the ONTAP software image for non-restricted countries, which includes Volume Encryption.

If you use the ONTAP software image for restricted countries to upgrade a system with Volume Encryption, the system panics and you lose access to your volumes.

Step 1. Locate the target ONTAP software in the Software Downloads area of the Lenovo Data Center Support site.

Step 2. Copy the software image.

For ONTAP 9.5 and later versions, copy the software image (for example, 95_q_image.tgz) from the Lenovo Data Center Support site to the directory on the HTTP server or FTP server from which the image will be served or to a local folder.

Updating a cluster nondisruptively

You can use Storage Manager to update a cluster or individual nodes in high-availability (HA) pairs 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 Lenovo Data Center Support site to your local workstation, or an HTTP server or FTP server on your network, so that the nodes can access the image.

[“Obtaining ONTAP software images” on page 44](#)

About this task

- If you try to perform other tasks from Storage 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.

- Starting with Storage Manager 9.6, if the NVMe protocol is configured in Storage Manager 9.5 and you perform an upgrade from Storage Manager 9.5 to Storage Manager 9.6, you no longer have a grace period of 90 days to have the NVMe protocol available without a license. If the grace period is in effect when you upgrade from ONTAP 9.5 to 9.6, the grace period must be replaced with a valid NVMeoF license so you can continue to use the NVMe features.

This feature is not available in MetroCluster configurations.

- If the NVMe protocol is not configured in Storage Manager 9.5 and you perform an update from Storage Manager 9.5 to Storage Manager 9.6, then the grace period is not provided, and you must install the NVMe license to use the NVMe protocol.

This feature is not available in MetroCluster configurations.

- Starting with ONTAP 9.5, at least one NVMe LIF must be configured for each node in an HA pair using the NVMe protocol. You can also create a maximum of two NVMe LIFs per node. When you upgrade to

ONTAP 9.5, you must ensure that a minimum of one NVMe LIF is defined for each node in an HA pair using the NVMe protocol.

Step 1. Click **Cluster → Overview**.

Step 2. In the **Overview** card, click the more icon .

Step 3. Select **ONTAP Update**.

Step 4. Click **+Add Image**.

Step 5. Select an image from a server or a local client.

Step 6. After the image is uploaded, click **Update** and follow the instructions on the screen.

How to update a cluster nondisruptively

You can use Storage 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 Lenovo Data Center 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.

SNMP

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

Enabling or disabling SNMP

You can enable or disable SNMP on your clusters by using Storage 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.

Step 1. Click **Cluster → Settings**.

Step 2. In the **SNMP** card, click the more icon .

Step 3. Select **Disable** or **Enable** as required.

Editing SNMP information

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

About this task

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

Step 1. Click **Cluster → Settings**.

Step 2. In the **SNMP** card, click the more icon .

Step 3. Select **Edit**.

Step 4. Specify a contact name and location.

Step 5. Add a trap host, community string and any security as required.

Step 6. Click **Update**.

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 Storage 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.

Step 1. Click **Cluster → Settings**.

Step 2. In the **SNMP** card, click the more icon .

Step 3. Select **Edit**.

Step 4. Specify a contact name and location.

Step 5. Add a trap host, community string and any security as required.

Step 6. Click **Update**.

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.

LDAP

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

Viewing the LDAP client configuration

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

Step 1. Click **Cluster** → **Settings**.

Step 2. In the **LDAP** card, click the gear icon .

Step 3. Specify the IP address of the LDAP server, schema to use, associated domain name and port.

Step 4. Specify the authentication level, user name and password to use.

Step 5. Click **Save**.

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 and Roles

You can use Storage 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 Storage 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.

Step 1. Click **Cluster → Settings**.

Step 2. In the **Users and Roles** card, click the arrow icon .

Step 3. Click **+Add**.

Step 4. Specify a user name.

Step 5. Specify a login method by selecting values from the drop-down lists in the **User Login Methods** area and then click **+Add**.

Multiple methods can be added.

Step 6. Specify the password and confirm it.

Step 7. Click **Save**.

Editing a cluster user account

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

Step 1. Click **Cluster → Settings**.

Step 2. In the **Users and Roles** card, click the arrow icon .

Step 3. Select an existing user and click the more icon .

Step 4. Select **Edit**.

Step 5. If you need to change the role, select the new role from the list.

Step 6. Specify a login method by selecting values from the drop-down lists in the **User Login Methods** area and then click **+Add**.

Multiple methods can be added.

Step 7. Optional: Remove roles by clicking on the trash can icon  next to that role.

Step 8. Click **Save**.

Changing passwords for cluster user accounts

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

Step 1. Click **Cluster → Settings**.

Step 2. In the **Users and Roles** card, click the arrow icon .

Step 3. Select an existing user and click the more icon .

Step 4. Select **Change Password**.

Step 5. Specify the password and confirm it.

Step 6. Click **Save**.

Locking or unlocking cluster user accounts

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

Step 1. Click **Cluster → Settings**.

Step 2. In the **Users and Roles** card, click the arrow icon .

Step 3. Select an existing user and click the more icon .

Step 4. Select **Lock**.

Step 5. In the displayed dialog box, click **Lock** to confirm the change.

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

Users and Roles 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.

Chapter 9. Managing the network

You can use Storage Manager to manage the network of your storage system by creating and managing IPspaces, broadcast domains, network interfaces, Ethernet ports, and FC/FCoE adapters.

IPspaces

You can use Storage Manager to create and manage IPspaces.

Editing IPspaces

You can use Storage 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.

Step 1. Click **Network → Overview**.

Step 2. Select the required IPspace and click the more icon .

Step 3. Type the new name and click **Save**.

Deleting IPspaces

You can use Storage 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.

Step 1. Click **Network → Overview**.

Step 2. Select the required IPspace and click the more icon .

Step 3. Select **Delete** and click **OK**.

Broadcast domains

You can use Storage Manager to create and manage broadcast domains.

Editing broadcast domain settings

You can use Storage 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 Storage Manager to edit broadcast domains in the cluster IPspace.

You must use the command-line interface (CLI) instead.

Step 1. Click **Network → Overview**.

Step 2. Select the required broadcast domain and click the more icon  .

Step 3. Select **Edit**.

Step 4. Modify the name, ports assigned or MTU as needed.

Step 5. Click **Save**.

Deleting broadcast domains

You can delete a broadcast domain by using Storage 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

- All ports assigned to the Broadcast domain must be removed before deleting it.
- You cannot use Storage Manager to delete broadcast domains that are in the cluster IPspace.

You must use the command-line interface (CLI) instead.

Step 1. Click **Network → Overview**.

Step 2. In the **Broadcast Domains** card, select the required broadcast domain and click the more icon  .

Step 3. Select **Delete**.

Step 4. Click **OK**.

Network interfaces

You can use Storage Manager to create and manage network interfaces.

Creating network interfaces

You can use Storage 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, or FC/FCoE protocols, or for management access only.

- You can specify an IP address by using a subnet or by not using a subnet.
- You cannot use Storage 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 NVMeoF data LIF the SVM must already be set up, the NVMe service must already exist on the SVM and the NVMeoF capable adapters should be available.
- NVMe protocol is enabled only if the selected SVM has the NVMe service configured.

Step 1. Click **Network → Overview**.

Step 2. In the **Network Interface** card, click the plus icon .

Step 3. Select the interface role and protocol to support.

Step 4. Specify the name and home node.

If you are creating a new iSCSI LIF, you also need specify the IP address and subnet mask.

Step 5. Click **Save**.

Editing network interface settings

You can use Storage 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 Storage Manager.
- You cannot enable management access for an intercluster LIF.

Step 1. Click **Network → Overview**.

Step 2. In the **Network Interfaces** card, select the required interface and click the more icon .

Step 3. Select **Edit**.

Step 4. In the Edit Network Interface dialog box, modify the network interface settings as required.

Step 5. Click **Save**.

Deleting network interfaces

You can use Storage 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.

Step 1. Click **Network → Overview**.

Step 2. In the **Network Interfaces** card, select the required interface and click the more icon .

Step 3. Select **Delete**.

Step 4. Select all check boxes and click **Delete**.

Migrating a LIF

You can use Storage 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.

Step 1. Click **Network → Overview**.

Step 2. In the **Network Interfaces** card, select the required interface and click the more icon .

Step 3. Select **Migrate**.

Step 4. Select the destination port.

Step 5. Optional: Select **Permanently migrate** if you want to permanently migrate the port.

Step 6. Click **Migrate**.

Ethernet ports

You can use Storage Manager to create and manage Ethernet ports.

Creating interface groups

You can use Storage 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.

Step 1. Click **Network → Ethernet Ports**.

Step 2. Click **+Link Aggregation Group**.

Step 3. Specify the node and ports to place in the link aggregation group.

Step 4. Select **Single**, **Multiple** or **LACP** as the link aggregation mode.

Step 5. Selecting **IP based**, **MAC based**, **Sequential**, or **Port** as the load distribution.

Step 6. Click **Save**.

Creating VLAN interfaces

You can create a VLAN to maintain separate broadcast domains within the same network domain by using Storage Manager.

Step 1. Click **Network → Ethernet Ports**.

Step 2. Click **+VLAN**.

Step 3. Specify the VLAN ID to use, the port that will host the VLAN, and the node that will host the port.

You may also need to specify the broadcast domain.

Step 4. Click **Save**.

Editing interface group settings

You can use Storage 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.

Step 1. Click **Network → Ethernet Ports**.

Step 2. Click the expand icon  next to the node.

Step 3. Select the required interface group and click the more icon  under the interface group.

Step 4. Click **Edit**.

Step 5. Make the changes to the interface group.

Step 6. Click **Save**.

Deleting VLANs

You can delete VLANs that are configured on network ports by using Storage 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.

Step 1. Click **Network → Ethernet Ports**.

Step 2. Click the expand icon  next to the node.

Step 3. Select the required VLAN and click the more icon  under the VLAN.

Step 4. Select **Delete** and then click **OK**.

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.

Chapter 10. Managing physical storage

You can use Storage Manager to manage physical storage such as local and cloud tiers, storage pools, disks, nodes, events, system alerts, AutoSupport notifications, and jobs.

Storage tiers

You can use Storage Manager to create aggregates to support the different security requirements, backup requirements, performance requirements, and data sharing requirements of your users.

Renaming a Storage Tier

You can use Storage Manager to rename an existing storage tier after it has been created.

Step 1. Click **Storage → Tiers**.

Step 2. Select the required tier and click the more icon .

Step 3. Select **Rename**.

Step 4. Type the new name.

Step 5. Click **Save**.

Deleting aggregates

You can use Storage 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 FlexVol volumes and the associated storage virtual machines (SVMs) contained by the aggregate must be deleted.
- The aggregate must be offline.

Step 1. Click **Storage → Tiers**.

Step 2. Select the required aggregate and click the more icon .

Step 3. Select **Delete**.

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 Storage 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 cached 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.

Step 1. Click **Storage → Volumes**.

Step 2. Select the required volume and click the more icon  .

Step 3. Select **Move**.

Step 4. Select the destination aggregate.

Step 5. Click **Move**.

Viewing aggregate information

You can use the Aggregates window in Storage Manager to view the name, status, and space information about an aggregate.

Step 1. Click **Storage → Tiers**.

Step 2. Select the required aggregate and click the more icon  .

Step 3. Select **More Details**.

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 9.5, 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 9.5, 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 Storage 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.

How Storage 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. Storage 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 serve 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. Storage 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, Storage 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, Storage Manager leaves 2 HDDs as hot spares and uses the other 8 HDDs for HDD-related operations.

Storage 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 is an exception to the hot spare rule in Storage 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.

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, Storage 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, Storage Manager displays the effective RPM as 10K.

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 cloud 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.

Storage recommendations for creating aggregates

Starting with Storage Manager 9.5, you can create aggregates based on storage recommendations. However, if your environment includes the following configurations, you must create aggregates manually using the CLI.

Storage 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. Storage 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 9.5 or earlier, or if your environment includes the following configurations, you must create aggregates manually:

- MetroCluster configurations
- SyncMirror functionality
- 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 (for ONTAP versions earlier than 9.6)
- Disks that are undergoing maintenance testing


External capacity tier

Configuring and managing cloud tiers

Storing data in tiers can enhance the efficiency of your storage system. You manage storage tiers by using FabricPool-enabled aggregates. Cloud tiers store data in a tier based on whether the data is frequently accessed.

Before you begin

- You must be running ONTAP 9.5 or later.
- You must have all flash (all SSD) aggregates



Adding a cloud tier

You can use Storage Manager to add a cloud tier to an SSD aggregate or a virtual machine disk (VMDK) aggregate. Cloud 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 cloud tier.
- If communication between the cloud tier and the cluster is encrypted using SSL or TLS, the required certificates must be installed.

About this task

The following object stores can be used as cloud tiers:

- StorageGRID
- Alibaba Cloud (Starting with Storage Manager 9.6)
- Amazon Web Services (AWS) Simple Storage Service (S3)

- Amazon Web Services (AWS) Commercial Cloud Service (C2S)
- Microsoft Azure Blob storage
- IBM Cloud
- Google Cloud

Notes:

- Azure Stack, which is an on-premises Azure service, is not supported.
- If you want to use any object store as a cloud tier, other than StorageGRID, you must have the FabricPool capacity license. You can add the license by clicking **Add License**.
- If you want to use an IBM Cloud Object Storage environment (such as Cleversafe), with FabricPool, you should specify a certification authority (CA) certificate. You can specify the CA certificate by moving the **Object Store Certificate** toggle button and specifying the certificate credentials.

Step 1. Click **Storage → Tiers**.

Step 2. Click **+Add Cloud Tier**.

Step 3. Select your cloud tier provider from the drop-down list.

Step 4. Fill in all of the required areas including port to use and bucket name.

Step 5. Select the local tier to add the cloud tier to.

Step 6. Select the intercluster LIFs to use.

Step 7. Click **Save**.

Changing the tiering policy of a volume

You can use Storage Manager to change the default tiering policy of a volume to control whether the data of the volume is moved to the cloud tier when the data becomes inactive.

Step 1. Click **Storage → Volumes**.

Step 2. From the list on the right pane, select the volume to modify.

Step 3. Click the more icon  .

Step 4. Select **Edit Cloud Tier Settings**.

Step 5. Select the new policy to use.

Step 6. Click **Save**.

Editing a cloud tier

You can use Storage Manager to modify the configuration information of cloud 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.

Step 1. Click **Storage → Tiers**.

Step 2. Click **+Add Cloud Tier**.

Step 3. Select your cloud service provider from the list.

Step 4. Fill out the required fields in the next screen.

Step 5. Click **Save**.

Deleting a cloud tier

You can use Storage Manager to delete a cloud tier that you no longer require.

Before you begin

You must have deleted the FabricPool-enabled local tier that is associated with the cloud tier.

- Step 1. Click **Storage → Tiers**.
- Step 2. Select the cloud tier that you want to delete.
- Step 3. Click the more icon .
- Step 4. Select **Delete**.
- Step 5. Select **Delete the cloud tier**.
- Step 6. Click **Delete**.

What cloud tiers and tiering policies are

Cloud tiers provide storage for infrequently accessed data. You can attach an all-flash (all-SSD) aggregate to a cloud tier to store infrequently used data. You can use tiering policies to decide whether data should be moved to a cloud 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.

Auto

Moves the inactive (cold) data and the Snapshot copies from the active file system to the cloud tier.

Backup (for Storage Manager 9.5)

Moves the newly transferred data of a data protection (DP) volume to the cloud tier.

All (starting with Storage Manager 9.6)

Moves all data to the cloud tier.

None

Prevents the data on the volume from being moved to a cloud tier.

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 and SSD aggregates. Inactive data is not displayed for FlexGroups.

Configuring S3 Object Store

You can use Storage Manager to configure S3 Object Store.

Before you begin

You will first need to create an SVM enabled for S3.

- Step 1. Go to **Storage → Storage VMs**.
- Step 2. Click **+Add**.
- Step 3. Specify a name for the new VM in the **STORAGE VM NAME** field.

- Step 4. Click the **SMB/CIFS, NFS, S3** tab.
- Step 5. Select **Enable S3**.
- Step 6. Specify a name for the S3 server.
- Step 7. Choose if you want to use TLS or HTTP.
- Step 8. Select the port to use.
- Step 9. Choose if you are going to use a system generated certificate or external CA signed certificate.
- Step 10. Specify the IP addresses for each node in the S3 virtual machine.
- Step 11. Click **Save**.
- Step 12. When the window stating the access key and certificate is displayed, click **Download**.
- Step 13. Choose **Storage → Buckets**.
- Step 14. Click **+Add**.
- Step 15. Type a name in the **NAME** field.
- Step 16. Verify that you select the SVM that you created previously.
- Step 17. Type a value in the **CAPACITY** field.
- Step 18. Click **Save**.

After you finish

All additional steps need to be completed from the CLI.

Aggregates

You can use Storage Manager to create aggregates to support the differing security, backup, performance, and data sharing requirements of your users.

Disks

You can use Storage Manager to manage disks.

Viewing disk information

You can use the Disks window in Storage Manager to view the name, size, and container details of disks along with graphical information about capacity disks and cache disks.

- Step 1. Click **Cluster → Disks**.
- Step 2. Select the required disk to view the details.

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
FSAS	Capacity	NL-SAS	Near Line SAS
SAS	Performance	SAS	Serial-Attached SCSI
SSD	Ultra-performance	SSD	Solid-state drives

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. These LUNs are referred to as <i>array LUNs</i> to distinguish them from the LUNs that ONTAP serves to clients.

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.

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.

About disks

A disk is the basic unit of storage for storage systems that use ONTAP software to access native disk shelves.

ONTAP software enables you to assign ownership to your disks 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. You can create an aggregate using any of the supported ONTAP disk types.

Events

You can use Storage 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 Storage Manager to monitor different parts of a cluster.

Acknowledging system health alerts

You can use Storage 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.

Step 1. Click **Events & Jobs → System Alerts**.

Step 2. Click the event to acknowledge.

Step 3. Click **Acknowledge**.

Suppressing system health alerts

You can use Storage Manager to suppress system health alerts that do not require any intervention from you.

Step 1. Click **Events & Jobs → System Alerts**.

Step 2. Click the required event.

Step 3. Click **Suppress**.

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.
- Resolve the problem by taking the corrective actions provided in the alert, such as fixing cabling to resolve a connectivity problem.
- 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.

AutoSupport notifications

You can use Storage Manager to configure AutoSupport notifications that help you to monitor your storage system health.

Enabling or disabling AutoSupport settings

You can enable or disable AutoSupport settings on your storage system by using Storage 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 disabled by default.

Step 1. Click **Cluster → Settings**.

Step 2. In the **AutoSupport** card, click the more icon .

Step 3. Choose **Enable** or **Disable** as required.

- If you have clicked **Enable**, no further operation is required.
- If you have clicked **Disable**, the **Disable AutoSupport** dialog box is displayed and you need to click **Disable**.

Adding AutoSupport email recipients

You can use the **Email Recipient** tab of the Edit AutoSupport Settings dialog box in Storage Manager to add email addresses of the recipients of AutoSupport notifications.

Step 1. Click **Cluster → Settings**.

Step 2. In the **AutoSupport** card, click the more icon .

Step 3. Select **More Options**.

Step 4. In the **Email** card, click **Edit**.

Step 5. Specify the email address to send from.

Step 6. Click **+Add** and specify the email address to send to and the recipient category.

Step 7. Click **Save**.

Testing AutoSupport settings

You can use the AutoSupport Test dialog box in Storage Manager to test that you have configured the AutoSupport settings correctly.

Step 1. Click **Cluster → Settings**.

- Step 2. In the **AutoSupport** card, click the more icon  .
- Step 3. Select **Test connectivity**.
- Step 4. In the **Test connectivity** dialog box, specify a subject and click **Send Test AutoSupport Message**.

Generating AutoSupport data

You can use Storage Manager to generate AutoSupport data for a single node or multiple nodes to monitor their health and to send notifications to technical support.

- Step 1. Click **Cluster → Settings**.
- Step 2. In the **AutoSupport** card, click the more icon  .
- Step 3. Select **Generate and Send**.
- Step 4. Specify a subject. Then, specify whether to collect the data from all nodes by selecting or clearing the **All nodes** check box.
- By default, **All nodes** is selected and no further operation is required.
 - If you wish to collect data from a single node, unselect **All nodes** and select the node from which the data is collected.
- Step 5. Click **Send**.

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. **Disable** is the default.

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.

Jobs

You can use Storage 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.

Chapter 11. Managing logical storage

You can use Storage Manager to manage the logical storage such as storage virtual machines (SVMs), volumes, tiers, and applications.

Storage Virtual Machines

You can use Storage Manager to manage the SVMs in your cluster.

Editing SVM settings

You can use Storage Manager to edit the properties of storage virtual machines (SVMs), such as the name, default language, and resource allocation.

Step 1. Click **Storage** → **Storage VMs**.

Step 2. Select the required SVM and click the more icon .

Step 3. Select **Edit**.

Step 4. Specify the SVM name, volume deletion retention period, and the language.

Step 5. Optional: Select where the volume creation is limited to using the **Resource Allocation** option and select **Limit volume creation to preferred local tiers**.

Step 6. Click **Save**.

Deleting SVMs

You can use Storage 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.

Step 1. Click **Storage → Storage VMs**.

Step 2. Select the required SVM and click the more icon .

Step 3. Select **Delete**.

Starting SVMs

You can use Storage Manager to provide data access from a storage virtual machine (SVM) by starting the SVM.

Step 1. Click **Storage → Storage VMs**.

Step 2. Locate the required SVM that is stopped and click the more icon .

Step 3. Select **Start**.

Result

The SVM starts serving data to clients.

Stopping SVMs

You can use Storage 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.

Step 1. Click **Storage → Storage VMs**.

Step 2. Locate the required SVM and click the more icon .

Step 3. Select **Stop**.

Result

The SVM stops serving data to clients.

Managing SVMs

A storage virtual machine (SVM) administrator can administer SVMs and their resources, such as volumes, protocols, and services, depending on the capabilities assigned by the cluster administrator. An SVM administrator cannot create, modify, or delete SVMs.

Note: SVM administrators cannot log in to Storage 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.

- 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.

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.

Note: Unless otherwise specified, the term SVM refers to a data (data-serving) SVM. In the CLI, SVMs are displayed as Vservers.

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	<pre>vserver services name-service unix-user vserver services name-service unix-group vserver services name-service netgroup vserver services name-service dns hosts</pre>
nis	External NIS servers as specified in the NIS domain configuration of the SVM	<pre>vserver services name-service nis-domain</pre>
ldap	External LDAP servers as specified in the LDAP client configuration of the SVM	<pre>vserver services name-service ldap</pre>
dns	External DNS servers as specified in the DNS configuration of the SVM	<pre>vserver services name-service dns</pre>

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.

Trace File Access window

Starting with Storage Manager 9.6, you can use the Trace File Access window to diagnose issues when you have problems accessing files and shares on an SVM using the CIFS or NFS protocol.

Command buttons

Continue

Starts the process of setting up and initiating a file access trace on the selected SVM.

Protocols

Allows you to select the protocol that is used to access files and shares on the selected SVM, either CIFS or NFS.

Advanced Options icon

Allows you to specify additional details to narrow the scope of the trace.

Show in Trace Results

Allows you to specify in the Advanced Options dialog box whether you want the trace results to display only file access requests that were denied or to display all file access requests—those that were successful and those that were denied.

Start Tracing

Allows you to start the trace. The results show access problems for file access requests submitted over the next 60 minutes.

Stop Tracing

Allows you to stop the trace.

View Permissions

Allows you to display permissions. When using the CIFS protocol, you can display effective file and share permissions. When using the NFS protocol, you can display effective file permissions.

Copy to Clipboard

Allows you copy the results table to the clipboard.

Export Trace Results

Allows you to export the trace results to a file in comma-separated-values (.csv) format.

Entry fields

User Name

You enter the name of the user who received file access request errors that you want to trace.

Search trace results

You enter specific information that you want to find in the search results, and then you click **Enter**.

Client IP Address

In the Advanced Options dialog box, you can specify the IP address of the client as an additional detail to narrow the scope of the trace.

File

In the Advanced Options dialog box, you can specify the file or file path that you want to access as an additional detail to narrow the scope of the trace.

Results list for CIFS protocol tracing

When you specify the CIFS protocol, the results list displays the following data for your trace based on the parameters you specified. The data is displayed in reverse chronological order. If you stop tracing, the results remain in the list until you start another trace.

- Share: The name of the share that the system attempted to access, whether successful or not.
- Path: The file path of the file that the system attempted to access, whether successful or not.
- Client IP Address: The IP address of the client from which access requests were initiated.
- Reasons: The reasons the attempt to access the file or share was successful or not.

Note: When the trace result shows a message saying that access is not granted for "Synchronize", "Read Control", "Read Attributes", "Execute", "Read EA", "Write", or "Read", the message is indicating that the desired access has not been not granted for the set of permissions listed. In order to view the actual permissions status, you need to view the permissions using the provided link.

Results list for NFS protocol tracing

When you specify the NFS protocol, the results list displays the following data for your trace based on the parameters you specified. The data is displayed in reverse chronological order. If you stop tracing, the results remain in the list until you start another trace.

- Path: The file path of the file that the system attempted to access, whether successful or not.
- Client IP Address: The IP address of the client from which access requests were initiated.
- Reasons: The reasons the attempt to access the file or share was successful or not.

Note: When the trace result shows a message saying that access is not granted for "Synchronize", "Read Control", "Read Attributes", "Execute", "Read EA", "Write", or "Read", the message is indicating that the desired access has not been not granted for the set of permissions listed. In order to view the actual permissions status, you need to view the permissions using the provided link.

Volumes

You can use Storage Manager to create, edit, and delete volumes.

You can access all the volumes in the cluster by using the Volumes tab.

Note: The Volumes tab is displayed only if you have enabled the CIFS and NFS licenses.

Editing volume properties

You can modify volume properties such as the volume name, security style, fractional reserve, and space guarantee by using Storage Manager. You can modify storage efficiency settings (deduplication schedule, deduplication policy, and compression) and space reclamation settings.

Before you begin

For enabling volume encryption, you must have installed the volume encryption license by using Storage Manager, and you must have 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 set the fractional reserve to either zero percent or 100 percent.
- Data compression is not supported on 32-bit volumes.
- You cannot rename a SnapLock Compliance volume.

Step 1. Click **Storage → Volumes**.

Step 2. Select the required volume and click the more icon .

Step 3. Select **Edit**.

Step 4. If required, select a new name and capacity for the volume.

Step 5. In the Storage and Optimization area, select the following options as required:

- **Enable think provisioning**
- **Resize automatically**
- **Enable quota**
- **Enforce performance limits**

Step 6. Select a security style and permissions.

Step 7. In the Snapshot Copies(Local) Settings, specify the snapshot reserve (%), the schedule for snapshots, and if the snapshots should be deleted automatically.

Step 8. Click **Save**.

Editing data protection volumes

You can use Storage 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.

Step 1. Click **Storage → Volumes**.

Step 2. Locate the required data protection volume that was created previously and click the more icon .

Step 3. Select **Edit**.

Step 4. Type a new name in the **NAME** field.

Step 5. Click **Save**.

Deleting volumes

You can use Storage Manager to delete a FlexVol 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 that data.

Before you begin

The following conditions must exist before you delete a FlexVol volume:

- The volume must be unmounted and must be in the offline state.
- FlexClone volumes must be either split from the parent volume or destroyed if the FlexVol volume is cloned.
- The SnapMirror relationships must be deleted if the volume is in one or more SnapMirror relationships.

About this task

You should be aware of the following limitations when deleting a FlexVol volume:

- 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.
- If the FlexVol contains both qtrees and volumes, the qtrees appear as directories. You should be careful to not delete the qtrees accidentally when deleting volumes.
- If you have associated FlexCache volumes with an origin volume, then you must delete the FlexCache volumes before you can delete the origin volume.

Step 1. Click **Storage → Volumes**.

Step 2. Select the required volume and click the more icon .

Step 3. Select **Delete**.

Step 4. Choose to take the volume off and delete the data.

Step 5. Click **Delete**.

Creating FlexClone volumes

You can use Storage 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.

Step 1. Click **Storage → Volumes**.

- Step 2. Select the required volume and click the more icon .
- Step 3. Select **Clone**.
- Step 4. Specify a name and then choose if you want to enable snapshots.
- Step 5. Click **Clone**.

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 Storage Manager. After the split, the FlexClone volume becomes a normal FlexVol volume.

Before you begin

The FlexClone volume must be online.

About this task

For systems that are *not* AFA systems, 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.

- Step 1. Click **Storage → Volumes**.
- Step 2. Select the clone volume that was created previously and click the more icon .
- Step 3. Select **Split Clone**.
- Step 4. Select to delete the snapshot copied and split the clone.
- Step 5. Click **Split**.

Setting the Snapshot copy reserve

You can use Storage Manager to reserve space (measured as a 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.

- Step 1. Click **Storage → Volumes**.
- Step 2. Select the required volume and click the more icon .
- Step 3. Select **Edit**.
- Step 4. Scroll down to locate the Snapshot Copies(Local) Settings area.
- Step 5. Specify the percentage of the Snapshot copy reserve.
- Step 6. Click **Save**.

Scheduling automatic creation of Snapshot copies

You can use Storage 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.

Step 1. Click **Storage → Volumes**.

Step 2. Select the required volume and click the more icon  .

Step 3. Select **Edit**.

Step 4. Scroll down to locate the Schedule Snapshot Copies(Local) Settings area.

Step 5. Verify that **Schedule Snapshot copies** is selected and specify a snapshot policy.

Step 6. If required, select **Automatically delete older Snapshot copies**.

Step 7. Click **Save**.

Extending the expiry date of Snapshot copies

You can use Storage 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.

Step 1. Click **Storage → Volumes**.

Step 2. From the list on the right pane, select the volume to modify.

Step 3. Select a volume.

Step 4. Click **Show More Details** to view more information about the volume.

Step 5. Click the **Snapshot Copies** tab.

The list of available Snapshot copies for the selected volume is displayed.

Step 6. Select the Snapshot copy that you want to modify, and then click **Extend Expiry Date**.

Step 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.

Step 8. Click **OK**.

Resizing volumes

When a 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 Storage 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).

Step 1. Click **Storage → Volumes**.

Step 2. Select the required volume and click the more icon  .

Step 3. Select **Edit**.

Step 4. Specify a new value in the **Capacity** field.

Step 5. Click **Save**.

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 Storage 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.

Step 1. Click **Storage → Volumes**.

Step 2. Select the required volume and click the more icon  .

Step 3. Select **Move**.

Step 4. Select the new aggregate.

Step 5. Click **Move** .

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 Storage 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.

Step 1. Click **Storage → Volumes**.

Step 2. Select the required volume and click the more icon  .

Step 3. Select **Edit**.

Step 4. Select **Enforce performance limits**.

Step 5. Specify whether you want to use an existing limit or a new limit.

- If you want to use an existing limit, select **extreme**, **performance**, or **value**.
- If you want to create a new limit, specify the policy group name, IOPs guarantee, bandwidth limit, and IOPs limit.

Step 6. Click **Save**.

Creating a mirror relationship from a source SVM

You can use Storage 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.

Notes:

- 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 of type read/write.
- The SnapLock aggregate type must be the same on both clusters.
- A maximum of 25 volumes can be protected in one selection.

About this task

- Storage 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.

- SVMs that are peered only for FlexCache applications and do not have peering permissions for SnapMirror applications are not shown in the list of SVMs in this task. You can use the ONTAP Storage Manager 9.6 enhanced peering workflow to give permission to, or peer to, these SVMs. You then can select them in this task to create a protection relationship.

Step 1. Click **Protection → Relationships**.

Step 2. Click **Protect** and select **Volumes**.

Step 3. Specify whether the relationship will be **Synchronous** or **Asynchronous** in the **PROTECTION POLICY** field.

Step 4. Specify the source cluster and destination cluster.

Step 5. Choose the storage VM associated with the volume and the volume to mirror.

Step 6. Select the target SVM and any destination settings.

Step 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.

Creating a vault relationship from a source SVM

You can use Storage 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.

Notes:

- 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 the 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 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 of type read/write.
- The SnapLock aggregate type must be the same.
- A maximum of 25 volumes can be protected in one selection.

About this task

- Storage 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.

- Step 1. Click **Protection → Relationships**.
- Step 2. Click **Protect** and select **Volumes**.
- Step 3. Specify whether the relationship will be **Synchronous** or **Asynchronous** in the **PROTECTION POLICY** field.
- Step 4. Specify the source cluster and destination cluster.
- Step 5. Choose the storage VM associated with the volume and the volume to mirror.
- Step 6. Select the target SVM and any destination settings.
- Step 7. Click **Save**.

Creating a mirror and vault relationship from a source SVM

You can use Storage 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. This relationship 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 9.5 or later.
- The SnapMirror license must be enabled on the source cluster and destination cluster.

Notes:

- 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 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 of type 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 9.5 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

- Storage Manager does not support a cascade relationship. For example, a destination volume in a relationship cannot be the source volume in another relationship.
- SVMs that are peered only for FlexCache applications and do not have peering permissions for SnapMirror applications are not shown in the list of SVMs in this task. You can use the ONTAP Storage Manager 9.6 enhanced peering workflow to give permission to, or peer to, these SVMs. You then can select them in this task to create a protection relationship.

- Step 1. Click **Protection → Relationships**.
- Step 2. Click **Protect** and select **Volumes**.

- Step 3. Specify whether the relationship will be **Synchronous** or **Asynchronous** in the **PROTECTION POLICY** field.
- Step 4. Specify the source cluster and destination cluster.
- Step 5. Choose the storage VM associated with the volume and the volume to mirror.
- Step 6. Select the target SVM and any destination settings.
- Step 7. Click **Save**.

Changing the tiering policy of a volume

You can use Storage Manager to change the default tiering policy of a volume to control whether the data of the volume is moved to the cloud tier when the data becomes inactive.

- Step 1. Click **Storage → Volumes**.
- Step 2. Select the volume to change and click the more icon .
- Step 3. Select **Edit Cloud Tier Settings**.
- Step 4. Select the new tiering policy.
- Step 5. Click **Save**.

Creating FlexGroup volumes

A FlexGroup volume can contain many volumes that can be administered as a group instead of individually. You can use Storage 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.
- Starting with Storage Manager 9.6, you can create FlexGroup volumes in a MetroCluster configuration.

- Step 1. Click **Storage → Volumes**.
- Step 2. Click **+Add**.
- Step 3. Specify the name, the storage VM to assign the FlexGroup to, and the capacity.
- Step 4. Click **More Options**.
- Step 5. Select to distribute volume data across the cluster.
- Step 6. Click **Save**.

Viewing FlexGroup volume information

You can use Storage 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 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.

- Step 1. Click **Storage → Volumes**.
- Step 2. Select the required FlexGroup volume.

- Step 3. Click the expand icon  next to the FlexGroup volume name.
All settings are displayed for the volume.

Editing FlexGroup volumes

Starting with Storage Manager 9.6, you can edit the properties of an existing FlexGroup volume.

Before you begin

The FlexGroup volume must be online.

About this task

FabricPool FlexGroup volumes can be expanded under the following conditions:

- A FabricPool FlexGroup volume can be expanded only with FabricPool aggregates.
- A non-FabricPool FlexGroup volume can be expanded only with non-FabricPool aggregates.
- If the FlexGroup volume contains a mix of FabricPool and non-FabricPool volumes, then the FlexGroup volume can be expanded with both FabricPool and non-FabricPool aggregates.

Step 1. Click **Storage → Volumes**.

Step 2. Select the required FlexGroup volume and click the more icon .

Step 3. Select **Edit**.

Step 4. Choose the parameters to modify and make the required changes.

Step 5. Click **Save**.

Resizing FlexGroup volumes

You can use Storage 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.

Step 1. Click **Storage → Volumes**.

Step 2. Select the required FlexGroup volume and click the more icon .

Step 3. Select **Edit**.

Step 4. Specify a new value in **Capacity**.

Step 5. Click **Save**.

Deleting FlexGroup volumes

You can use Storage 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

Storage Manager does not support constituent level of management for FlexGroup volumes.

Step 1. Click **Storage → Volumes**.

Step 2. Select the FlexGroup volume and click the more icon .

Step 3. Select **Delete**.

Step 4. Select the check boxes of unmounting volume, taking volume offline, and deleting data.

Step 5. Click **Delete**.

What Volume Encryption is

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.

When NetApp Aggregate Encryption is enabled on an aggregate, new volumes are encrypted by default. Volume encryption can override the default encryption.

Note: When a selected aggregate is encrypted, volume encryption affects cross-volume storage efficiency.

Snapshot configuration

You can configure Snapshot copies by setting a schedule for an existing Snapshot policy. Starting with ONTAP 9.5, 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.

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.

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-Lenovo systems.

Data compression and deduplication


Beginning with Data ONTAP 9.5, 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.

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.

Shares

You can use Storage Manager to create, edit, and manage shares.

Creating a CIFS share

You can use Storage 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.

- Step 1. Click **Storage → Shares**.
- Step 2. Click **+Add** and select **Share**.
- Step 3. Specify a name for the new share.

- Step 4. Browse to the volume that was created previously and add it under folder name.
- Step 5. Add permissions and choose if you want the data encrypted.
- Step 6. Click **Save**.

Result

The CIFS share is created with the access permissions set to “Full Control for Everyone” in the group.

Creating home directory shares

You can use Storage Manager to create a home directory share and to manage home directory search paths.

Before you begin

CIFS must be set up and started.

- Step 1. Click **Storage → Shares**.
- Step 2. Select **+Add** and then **Home Directory**.
- Step 3. Specify the name of the share and add a relative path.
- Step 4. In the home directory search path, add a path to the volume that was created previously.
- Step 5. Add permissions if needed.
- Step 6. Click **Save**.

Editing share settings

You can use Storage 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). Starting with Storage Manager 9.6, continuous availability is supported for FlexGroup volumes.

- Step 1. Click **Storage → Shares**.
- Step 2. Select the share that you want to modify and click the more icon  .
- Step 3. Select **Edit**.
- Step 4. Change the permissions if needed.
- Step 5. Choose symbolic links to use.
- Step 6. Set the properties as required.
- Step 7. Click **Save**.

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 who 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: `/vol0home/home`
- Search path #2: `/vol1home/home`
- Search path #3: `/vol2home/home`
- Home directory: `/vol1home/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:

- `/vol0home/home/jsmith` does not exist; moving on to search path #2.
- `/vol1home/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:

- /vol0home/home/acme/jsmith does not exist; moving on to search path #2.
- /vol1home/home/acme/jsmith does not exist; moving on to search path #3.
- /vol2home/home/acme/jsmith does not exist; the home directory does not exist; therefore, the connection fails.

LUNs

You can use Storage 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 clicking **SVMs → LUNs**.

Note: The LUNs tab is displayed only if you have enabled the FC/FCoE and iSCSI licenses.

Creating LUNs

You can use Storage Manager to create a new volume and LUNs for an existing aggregate when there is available free space.

About this task

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, Storage 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.

Step 1. Click **Storage → LUNs**.

Step 2. Click **+Add**.

Step 3. Type a name in the **NAME PREFIX** field.

Step 4. Type the number of required LUNs in the **NUMBER OF LUNS** field.

Step 5. Type the capacity in the **CAPACITY** field.

Step 6. Select an operating system from the drop-down list of the **HOST OPERATING SYSTEM** field.

Step 7. Type initiator port names in the **HOST INITIATORS** field.

Step 8. Click **Save**.

Deleting LUNs

You can use Storage 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.

Step 1. Click **Storage → LUNs**.

Step 2. Select the required LUN and click the more icon .

Step 3. Select **Delete**.

Step 4. Select to unmap the LUN and take the LUN offline.

Step 5. Click **Delete**.

Creating initiator groups

You can use Storage 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.

Step 1. Click **Hosts → SAN Initiator Groups**.

Step 2. Click **+Add**.

Step 3. Type a name in the provided area.

Step 4. Select an SVM to which the host is assigned.

Step 5. Select an operating system from the provided list.

Step 6. Click **+Add Initiator** and type the initiator identifier in the area provided.

Step 7. Click **Delete**.

Deleting initiator groups

You can use Storage Manager to delete an existing initiator group.

Before you begin

All the LUNs mapped to the initiator group must be manually unmapped.

Step 1. Click **Hosts → SAN Initiator Groups**.

Step 2. Select the required host and click the more icon .

Step 3. Select **Delete**.

Step 4. In the displayed dialog box, click **Delete**.

Adding initiators

You can use Storage 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.

Step 1. Click **Hosts → SAN Initiator Groups**.

Step 2. Select the host and click the more icon .

Step 3. Select **Edit**.

Step 4. Select **+Add Initiator**.

Step 5. Type the initiator address to be added in the provided area.

Step 6. Click **Save**.

Deleting initiators from an initiator group

You can use Storage Manager to delete an initiator from an 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.

Step 1. Click **Hosts → SAN Initiator Groups**.

Step 2. Select the host and click the more icon .

Step 3. Select **Edit**.

Step 4. Locate the initiator port to be removed and click the trash can icon  in the same row.

Step 5. Click **Save**.

Cloning LUNs

LUN clones enable you to create multiple readable and writable copies of a LUN. You can use Storage 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 Storage Manager.

The LUN clone is deleted when ONTAP triggers automatic deletion to conserve space.

- You cannot clone LUNs that are on SnapLock volumes.

Step 1. Click **Storage → LUNs**.

Step 2. Select the LUN and click the more icon .

Step 3. Select **Clone**.

Step 4. Select a name for the cloned LUN using the provided area.

Step 5. Select an operating system for the new LUN.

Step 6. Select to use either a new initiator group or an existing initiator group.

- To use an new initiator group, you need to provide the initiator port address.
- To use an existing group, you need to select it from the list.

Step 7. Click **Clone**.

Editing LUNs

You can use the LUN properties dialog box in Storage Manager to change the name, description, size, 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.

Step 1. Click **Storage → LUNs**.

Step 2. Select the LUN to modify and click the more icon .

Step 3. Select **Edit**.

Step 4. Make the desired changes.

Step 5. Click **Save**.

Moving LUNs

You can use Storage 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 Storage Manager.
- The LUN move operation is nondisruptive; it can be performed when the LUN is online and serving data.
- You cannot use Storage 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.

Step 1. Click **Storage → LUNs**.

Step 2. Select the LUN to move and click the more icon .

Step 3. Select **Move**.

Step 4. In the MOVE TO VOLUME dialog box, select the volume to which the LUN is removed.

Step 5. Click **Move**.

Editing initiator groups

You can use the Edit Initiator Group dialog box in Storage 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.

Step 1. Click **Hosts → SAN Initiator Groups**.

Step 2. Select the initiator group to modify and click the more icon .

Step 3. Select **Edit**.

Step 4. Make the desired changes.

Step 5. Click **Save**.

Viewing LUN information

You can use the LUN Management tab in Storage Manager to view details about a LUN, such as its name, status, size, and type.

Step 1. Click **Storage → LUNs**.

Step 2. Select the LUN that you want to view.

Step 3. Click the expand icon  next to the LUN name to view the details.

Viewing initiator groups

You can use the Initiator Groups tab in Storage 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.

Step 1. Click **Storage → LUNs**.

Step 2. Click **Initiator Groups** and review the initiator groups that are listed in the upper pane.

Step 3. Select an initiator group to view the initiators that belong to it, which are listed in the Initiators tab in the lower pane.

Step 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.

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 Lenovo Storage Interoperation Center (LSIC). <https://datacentersupport.lenovo.com/lsc>

The following table describes the LUN multiprotocol type values and the guidelines for using each type:

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. 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.
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. 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. 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.

LUN multiprotocol type	When to use
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. Note: For raw LUNs, you can use the type of guest operating system that the LUN multiprotocol type uses.

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 ostyles of initiators are solaris, windows, xen, hyper_v, vmware, and linux.

You must select an ostyle for the igroup.

Qtrees

You can use Storage Manager create, edit, and delete Qtrees.

Creating qtrees

Qtrees enable you to manage and partition your data within a volume. You can use the Create Qtree dialog box in Storage Manager to add a new qtree to a volume on your storage system.

- Step 1. Click **Storage → Qtrees**.
- Step 2. Click **+Add**.
- Step 3. Provide a name in the required area.
- Step 4. Select a volume from the list.
- Step 5. Choose a security style.
- Step 6. Select if you want to enable quotas.
- Step 7. Click **Save**.

Deleting qtrees

You can delete a qtree and reclaim the disk space that the qtree uses within a volume by using Storage 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.

- Step 1. Click **Storage → Qtrees**.
- Step 2. Select the required Qtree and click the more icon .
- Step 3. Select **Delete**.
- Step 4. Select the check box stating that you wish to continue.
- Step 5. Click **Delete**.

Editing qtrees

You can use Storage Manager to modify the properties of a qtree such as the security style and assign a new or existing export policy.

- Step 1. Click **Storage → Qtrees**.

- Step 2. Select the required Qtree and click the more icon  .
- Step 3. Select **Edit**.
- Step 4. Make the desired changes.
- Step 5. Click **Save**.

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 Storage 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.

- Step 1. Click **Storage → Qtrees**.
- Step 2. Select the Qtree to modify and click the more icon  .
- Step 3. Select **Edit Export Policy**.
- Step 4. Choose if you want to inherit the policy from the volume, if you want to add a new export policy, or if you will use an existing export policy.
- To add a new export policy, you need to assign clients, assign the read-only and read-write rule, and assign a rule index.
 - To use an existing export policy, you need to select it from the list.
- Step 5. Click **Save**.

Viewing qtree information

You can use the Qtrees window in Storage Manager to view the volume that contains the qtree, the name, security style, and status of the qtree, and the oplocks status.

- Step 1. Click **Storage → Qtrees**.
- Step 2. Select the Qtree to view.
- Step 3. Click the expand icon  next to the Qtree to view the details.

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. Qtrees are used to manage and partition data within the volume.

If you create qtrees on a FlexVol that contains volumes, the qtrees appear as directories. Therefore, you need to be careful to not delete the qtrees accidentally when deleting volumes.

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 for each qtree.

- 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

You can 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

Quotas

You can use Storage Manager to create, edit, and delete quotas.

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 Storage Manager to create a quota and to apply the quota to a specific volume or qtree.

About this task

Using Storage 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 100. If you want to specify a value lower than 100, you should use the command-line interface (CLI).

Step 1. Click **Storage → Quotas**.

Step 2. Click **+Add**.

Step 3. Select the volume for the quota target.

Step 4. If **Enable Quota** is selected, select the Qtree from the list.

If **User** or **Group** is selected, you need to select the user or group to apply the quota to.

Step 5. Specify the file limit and the space limit.

Step 6. Click **Save**.

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.

Deleting quotas

You can use Storage Manager to delete one or more quotas when your users and their storage requirements and limitations change.

Step 1. Click **Storage → Quotas**.

Step 2. Select the required Quota and click the more icon .

Step 3. Select **Delete**.

Step 4. In the displayed dialog box, click **Delete**.

Editing quota limits

You can use Storage 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.

Step 1. Click **Storage** → **Quotas**.

Step 2. Select the Quota to modify and click the more icon  .

Step 3. Select **Edit**.

Step 4. Choose the setting to modify.

Step 5. Click **Save**.

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.

Quotas can be soft or hard. Soft quotas cause Data ONTAP to send a notification when specified limits are exceeded, and hard quotas prevent a write operation from succeeding when specified limits are exceeded.

Hard quotas impose a hard limit on system resources; any operation that would result in exceeding the limit fails. The following settings create hard quotas:

- Disk Limit parameter
- Files Limit parameter

Soft quotas send a warning message when resource usage reaches a certain level, but do not affect data access operations, so you can take appropriate action before the quota is exceeded. The following settings create soft quotas:

- Threshold for Disk Limit parameter
- Soft Disk Limit parameter
- Soft Files Limit parameter

Threshold and Soft Disk quotas enable administrators to receive more than one notification about a quota. Typically, administrators set the Threshold for Disk Limit to a value that is only slightly smaller than the Disk Limit, so that the threshold provides a "final warning" before writes start to fail.

Disk space hard limit

Disk space limit applied to hard quotas.

Disk space soft limit

Disk space limit applied to soft quotas.

Threshold limit

Disk space limit applied to threshold quotas.

Files hard limit

The maximum number of files on a hard quota.

Files soft limit

The maximum number of files on a soft quota.

Quota management

Storage 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

CIFS protocol

You can use Storage Manager to enable and configure CIFS servers to allow CIFS clients to access files on the cluster.

Setting up CIFS

You can use Storage 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, Storage Manager sets the security style as UNIX.

- Superuser access must be set to Any for the CIFS protocol.

- Step 1. Click **Storage → Storage VMs**.
- Step 2. Click **+Add**.
- Step 3. Specify the name of the new SVM to create.
- Step 4. Click **Enable SMB/CIFS**.
- Step 5. Specify the following information:
 - a. Specify the administrator name and password to add the system to either a domain or workgroup.
 - b. Specify the server name to create in the domain or workgroup.
 - c. Specify the domain or workgroup to add the system to.
- Step 6. Ensure that the DNS domain is correct and that the name server is valid.
- Step 7. Specify the LIF interface for each controller to use for the new CIFS SVM.
- Step 8. Click **Save**.

Editing the name for a CIFS SVM

You can modify the name of the CIFS SVM using the **Edit** option.

- Step 1. Click **Storage → Storage VMs**.
- Step 2. Click the CIFS SVM to modify and click the more icon .
- Step 3. Select **Edit**.
- Step 4. Type the new name in the provided field.
- Step 5. Click **Save**.

Adding home directory paths

You can use Storage 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.

- Step 1. Click **Storage → Shares**.
- Step 2. Click **+Add**.
- Step 3. Select **Home Directory** from the drop-down list.
- Step 4. Specify the name for the home directory in the name field.
- Step 5. Specify a relative path in the provided field.
- Step 6. Add a search path for the home directory.
- Step 7. Add user access permissions.
- Step 8. Click **Save**.

Deleting home directory paths

You can use Storage 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.

Step 1. Click **Storage → Shares**.

Step 2. Select the home directory to delete and click the more icon  .

Step 3. Select **Delete**.

Step 4. Select the check box to delete the share.

Step 5. Click **Delete**.

NFS protocol

You can use Storage Manager to authenticate NFS clients to access data on the SVM.

Editing NFS settings

You can use Storage 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.

Step 1. Click **Storage → Volumes**.

Step 2. Select the volume that is exported and click the more icon  .

Step 3. Select **Edit**.

Step 4. Select an export policy and then change the rule assigned with that export policy to either enable or disable NFS3 or NFS4.

You also need to select the user authentication and read-and-write or read-only permissions.

Step 5. Click **Save**.

NVMe protocol


You can use Storage 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 Storage 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.

Starting with ONTAP 9.5, you must configure at least one NVMe LIF for each node in an HA pair that uses the NVMe protocol. You can also define a maximum of two NVMe LIFs per node. You configure the NVMe LIFs when you create or edit the SVM settings using Storage Manager.

The following illustration shows the workflow for setting up NVMe:



Creating an NVMe namespace

You can use Storage 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.

- Step 1. Click **Storage → NVMe namespaces**.
- Step 2. Click **+Add**.
- Step 3. Specify the name in the provided list.
- Step 4. Specify the number of namespaces to create and the capacity for the namespaces.
- Step 5. Select the host operating system for the namespace and NVMe subsystem to assign the namespace to.
- Step 6. Click **Save**.

Editing an NVMe namespace

You can use Storage 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.

Step 1. Click **Storage → NVMe namespaces**.

Step 2. Select the required namespace and click the more icon  .

Step 3. Select **Edit**.

Step 4. Specify whether you want to add a new subsystem or use an existing namespace.

To use an existing namespace, you need to select it from the drop-down list.

Step 5. Click **Save**.

Cloning an NVMe namespace

You can use Storage 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.

Step 1. Click **Storage → NVMe namespaces**.

Step 2. Select the required namespace and click the more icon  .

Step 3. Select **Clone**.

Step 4. Specify a name for the new namespace.

Step 5. Choose a subsystem to map the namespace to.

Step 6. Click **Clone**.

What NVMe is

The nonvolatile memory express (NVMe) protocol is a transport protocol used for accessing nonvolatile storage media.

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.

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
- NVMe 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

What an NVMe subsystem is

An NVMe subsystem includes one or more controllers, one or more namespaces, one or more non-volatile memory (NVM) subsystem ports (FC-NVMe or RDMA transport ports), an NVM storage medium, and an interface between the controllers and the NVM storage medium. For controller mapping and management, an NVM subsystem maps to a vservers in ONTAP.

An NVMe subsystem can be created using Storage Manager. You can associate the NVMe subsystem with different hosts and namespaces within the vservers. Also, each vservers can support more than one NVMe subsystem. However, you cannot configure a NVMe subsystem to be used on multiple vservers.

An NVMe over Fabric (NVMeoF) subsystem is a separate kernel object that resides in the FreeBSD kernel. The NVMeoF subsystem interfaces with the following components:

- SAN components, such as BCOMKA, FCT, and VDOM
- WAFL
- RAS components, such as CM, ASUP, and EMS

All interfaces with NVMeoF subsystems adhere to the current definitions and patterns found in ONTAP.

Creating NVMe subsystems

You can use Storage Manager to create an NVMe subsystem.

- Step 1. Click **Storage → NVMe namespaces**.
- Step 2. Select **+Add**.
- Step 3. Specify the name for the new NVMe namespace.
- Step 4. Specify the number of namespaces to add.
- Step 5. Specify the size for the new namespace.
- Step 6. Click **More Options**.
- Step 7. Choose to add a new subsystem.
- Step 8. Specify the name and host NQN for the new subsystem.
- Step 9. Click **Save**.

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 and are the equivalent of LUNs for FC and iSCSI protocols.

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 subsystem provisioning for NVMe namespaces

An NVMe 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. When you create an NVMe namespace, you can choose to map an NVMe subsystem to the namespace, as follows:

None (default)

No NVMe subsystems are mapped to the namespace.

Existing subsystem

You can select an existing NVMe subsystem to map to the namespace. NVMe subsystems are listed based on the host OS and SVM fields. When you hover the pointer over the NVMe subsystem name, more details are shown about the subsystem.

New subsystem

You can create a new NVMe subsystem and map it to the namespace. The subsystem is created on the host OS and SVM.

You provision a subsystem by providing the following details:

- **The NVMe subsystem name** The NVMe subsystem name is case sensitive. It must contain 1 to 96 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 contain comma-separated values because there can be from one to many hosts attached to a subsystem.

iSCSI protocol

You can use Storage Manager to configure the iSCSI protocol that enables you to transfer block data to hosts using SCSI protocol over TCP/IP.

Starting or stopping the iSCSI service

You can use Storage Manager to start or stop the iSCSI service on your storage system.

Step 1. Click **Storage → Storage VMs**.

Step 2. Select the required SVM that you want to stop or start and click the more icon .

Step 3. Select **Start** or **Stop** as required.

- If you have clicked **Start**, no further operation is required.
- If you have clicked **Stop**, the **Stop Storage VM** dialog box is displayed and you need to click **Stop**.

FC/FCoE protocol

You can use Storage Manager to configure FC/FCoE protocols.

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 Storage 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.

Step 1. Click **Storage → Storage VMs**.

Step 2. Select the SVM that you want to stop or start and click the more icon .

Step 3. Select **Start** or **Stop** as required.

- If you have clicked **Start**, no further operation is required.
- If you have clicked **Stop**, the **Stop Storage VM** dialog box is displayed and you need to click **Stop**.

Export policies

You can use Storage Manager to create, edit, and manage export policies.

Creating an export policy

You can use Storage Manager to create an export policy so that clients can access specific volumes.

Step 1. Click **Storage → Storage VMs**.

Step 2. Select either an NFS or CIFS enabled SVM by clicking the SVM name.

Step 3. Click **Settings** on the displayed right pane.

Step 4. Scroll down to locate the **Export Policies** card and click the arrow icon .

Step 5. Click **+Add**.

Step 6. Specify a name for the new export policy.

Step 7. Select one of the followings and then perform subsequent operations:

- **Copy rules from existing policy:** Specify the storage SVM that it is applied to and the export policy.
- **Add New Rules:**
 1. Click **+Add**.
 2. Specify the client specification, access protocols, and the read/write rules. Then, specify whether to enable the superuser access.

Step 8. Click **Save**.

Renaming export policies

You can use Storage Manager to rename an existing export policy.

Step 1. Click **Storage → Storage VMs**.

Step 2. Click the SVM with the assigned export policy.

Step 3. Click **Settings** on the displayed right pane.

- Step 4. Scroll down to locate the **Export Policies** card and click the arrow icon .
- Step 5. Select the required export policy and type the new name.

Deleting export policies

You can use Storage Manager to delete export policies that are no longer required.

- Step 1. Click **Storage → Storage VMs**.
- Step 2. Click the required SVM with the assigned export policy.
- Step 3. Click **Settings** on the displayed right pane.
- Step 4. Scroll down to locate the **Export Policies** card and click the arrow icon .
- Step 5. Select the required export policy and click **Delete** above.

Adding rules to an export policy

You can use Storage 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.

- Step 1. Click **Storage → Storage VMs**.
- Step 2. Click the SVM with the assigned export policy.
- Step 3. Click **Settings** on the displayed right pane.
- Step 4. Scroll down to locate the **Export Policies** card and click the arrow icon .
- Step 5. Select the export policy that you want to modify. Then, click **+Add** above.
- Step 6. Specify the client specification, access protocols, and the read/write rules. Then, specify whether to enable the Superuser access.
- Step 7. Click **Save**.

Modifying export policy rules

You can use Storage Manager to modify the specified client, access protocols, and access permissions of an export policy rule.

- Step 1. Click **Storage → Storage VMs**.
- Step 2. Click the required SVM with the assigned export policy.
- Step 3. Click **Settings** on the displayed right pane.
- Step 4. Scroll down to locate the **Export Policies** card and click the arrow icon .
- Step 5. Select the export policy that you want to modify. Make sure that the policy is selected on the current rule that you wish to modify. Then, click the more icon .
- Step 6. Select **Edit**.
- Step 7. Specify the client specification, access protocols, and the read/write rules. Then, specify whether to enable the superuser access.
- Step 8. Click **Save**.

Deleting export policy rules

You can use Storage Manager to delete export policy rules that are no longer required.

- Step 1. Click **Storage → Storage VMs**.
- Step 2. Click the required SVM with the assigned export policy.
- Step 3. Click **Settings** on the displayed right pane.
- Step 4. Scroll down to locate the **Export Policies** card and click the arrow icon .
- Step 5. Select the export policy that you want to modify. Make sure that the policy is selected on the current rule that you wish to modify. Then, click the more icon .
- Step 6. Select **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.

QoS policy groups

You can use Storage Manager to create, edit, and delete QoS policy groups.

Creating QoS policy groups

You can use Storage Manager to create storage Quality of Service (QoS) policy groups to limit the throughput of workloads and to monitor workload performance.

- Step 1. Click **Storage → Volumes**.
- Step 2. Select the LUN to which the new QoS group is applied and click the more icon .
- Step 3. Select **Edit**.
- Step 4. Scroll down to locate the **Storage and Optimization** area.
- Step 5. Select **Enforce performance limits**.
- Step 6. Click **New** to create a new policy.
- Step 7. Specify a name for the policy and limits of IOPs and bandwidth.

Step 8. 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 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.



LDAP client services

You can use Storage Manager to add, edit, and delete LDAP client configurations.

Adding an LDAP client configuration

You can use Storage 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.

- Step 1. Click **Cluster** → **Settings**.
- Step 2. Locate the **LDAP** area and click the gear icon  .
The Edit LDAP page is displayed.
- Step 3. In the **LDAP Servers** area, specify a value in **BASE DN**.
- Step 4. In the **Binding** area, specify values in **MINIMUM AUTHENTICATION LEVEL**, **USERNAME**, and **PASSWORD**.
- Step 5. Click **Save**.

Deleting an LDAP client configuration

You can use Storage 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.

- Step 1. To delete an LDAP client configuration:

Cluster level: Click  → **LDAP**.

SVM level: Click **SVM** → **SVM Settings** → **LDAP Client**.

- Step 2. Select the LDAP client that you want to delete, and then click **Delete**.
- Step 3. Select the confirmation check box, and then click **Delete**.

Step 4. Verify that the LDAP client that you deleted is no longer displayed.

Editing an LDAP client configuration

You can use Storage 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.

Step 1. To edit an LDAP client configuration:

Cluster level: Click  → **LDAP**.

SVM level: Click **SVM** → **SVM Settings** → **LDAP Client**.

Step 2. Select the LDAP client that you want to modify, and then click **Edit**.

Step 3. In the Edit LDAP Client dialog box, edit the LDAP client configuration as required.

Step 4. Click **Save and Close**.

Step 5. Verify that the changes that you made to the LDAP client configuration are displayed.

LDAP configuration services

You can use Storage Manager to manage LDAP configurations.

Editing active LDAP clients

You can use Storage 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.

Step 1. Click **Storage** → **SVMs**.

Step 2. Select the SVM, and then click **SVM Settings**.

Step 3. In the Services pane, click **LDAP Configuration**.

Step 4. In the LDAP Configuration window, click **Edit**.

Step 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.

Step 6. Click **OK**.

Step 7. Verify that the changes that you made are updated in the LDAP Configuration window.

Deleting active LDAP clients

You can use Storage Manager to delete an active LDAP client when you do not want a storage virtual machine (SVM) to be associated with it.

Step 1. Click **Storage** → **SVMs**.

Step 2. Select the SVM, and then click **SVM Settings**.

Step 3. Click the **SVM Settings** tab.

Step 4. In the Services pane, click **LDAP Configuration**.

Step 5. Click **Delete**.

Step 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.

DNS Services

You can use Storage Manager to manage DNS services.

Enabling or disabling DNS

You can use Storage Manager to configure the DNS settings for your system.

About this task

- DNS is enabled by default.
- Storage Manager does not perform any validation checks for the DNS settings.

Step 1. Click **Cluster → Overview**.

Step 2. In the **Overview** card, click the more icon 

Step 3. Select **Edit**.

Step 4. Add or remove DNS entries in the **NAME SERVERS** field.

Step 5. Click **Save**.

Chapter 12. Managing data protection

You can use Storage 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 Storage Manager to create and manage mirror relationships by using the mirror policy.

Creating a mirror relationship from a destination SVM

You can use ONTAP Storage 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 9.4 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.

About this task

- Storage 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.
- You cannot create a mirror relationship between SnapLock volumes if the destination cluster is running a version of ONTAP that is older than the ONTAP version that the source cluster is running.

Step 1. Click **Protection → Relationships**.

Step 2. Click **Protect → Volumes**.

Step 3. In the **Source** area, select the cluster to use, the SVM to mirror, and the source volume.

Step 4. In the **Destination Settings**, specify the destination SVM.

You can also change the volume source and destination name if required.

Step 5. Click **Save**.

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.

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.

Step 1. Click **Protection → Relationships**.

Step 2. Select the relationship to delete and click the more icon .

Step 3. Select **Delete**.

Step 4. In the displayed dialog box, click **Delete**.

Result

The relationship is deleted, and the base Snapshot copy on the source volume is deleted.

Editing mirror relationships

You can use Storage 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 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.

Step 1. Click **Protection → Volume Relationships**.

- Step 2. Select the mirror relationship for which you want to modify the policy or schedule, and then click **Edit**.
- Step 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	1. Click Create Policy. 2. Specify a name for the policy. 3. Set the priority for scheduled transfers. 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. 4. 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. 5. Select the Enable Network Compression check box to compress the data that is being transferred. 6. Click Create.

- Step 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	1. Click Create Schedule. 2. Specify a name for the schedule. 3. Select either Basic or Advanced. • Basic specifies only the day of the week, time, and the transfer interval. • Advanced creates a cron-style schedule. 4. Click Create.
You do not want to assign a schedule	Select None .

- Step 5. Click **OK** to save the changes.

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.

Step 1. Click **Protection → Relationships**.

Step 2. Select the mirror relationship to update and click the more icon .

Step 3. Select **Update**.

Step 4. In the displayed dialog box, click **Update**.

Quiescing mirror relationships

You can use Storage Manager to quiesce a mirror destination to stabilize it 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.

Step 1. Click **Protection → Relationships**.

Step 2. Select the mirror relationship to pause and click the more icon .

Step 3. Select **Pause**.

Step 4. In the displayed dialog, click **Pause**.

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 Storage Manager. You must use the CLI to resume the relationship.

Step 1. Click **Protection → Relationships**.

Step 2. Select the mirror relationship to resume and click the more icon .

Step 3. Select **Resume**.

Step 4. In the displayed dialog box, click **Resume**.

Result

Data transfer to the mirror destination resumes for the selected mirror relationship.

Breaking SnapMirror relationships

You must break a SnapMirror relationship if a SnapMirror source becomes unavailable and you want client applications to be able to access the data from the mirror destination. After the SnapMirror relationship is broken, the destination volume type changes from "data protection" (DP) to "read/write" (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 breaking a FlexGroup volume relationship, you must refresh the page to view the updated status of the relationship.

Step 1. Click **Protection → Relationships**.

Step 2. Select the mirror relationship to break and click the more icon  .

Step 3. Select **Break**.

Step 4. In the displayed dialog box, click **Break**.

Result

The data protection SnapMirror relationship is broken. The destination volume type changes from data protection (DP), read-only, to read/write (RW). The system stores the base Snapshot copy for the data protection mirror relationship for later use.

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 cluster and destination cluster and the source SVM and destination SVM 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 volume.

Attention:

- For SnapLock compliance volumes, all data changes in the active file system with reference to the common Snapshot copy are preserved in a locked Snapshot copy until the expiry time that is set for the current volume.

If the volume expiry time is in the past or has not been set, then the Snapshot copy and the common Snapshot copy are locked for a duration of 30 days. All of the intermediate Snapshot copies between the common Snapshot copy and the latest locked Snapshot copy are deleted.

- For all volumes other than SnapLock compliance volumes, the resynchronization operation might cause loss of newer data that is 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.

Step 1. Click **Protection → Volume Relationships**.

- Step 2. Select the mirror relationship that you want to resynchronize.
- Step 3. Click **Operations → Resync**.
- Step 4. Select the confirmation checkbox, and then click **Resync**.

Reverse resynchronizing mirror relationships

You can use Storage Manager to reestablish a mirror relationship that was previously broken. In a reverse resynchronization operation, you reverse the functions of the source volume and destination volume.

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 volume, update the source volume, 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 volume.

Attention:

- For SnapLock compliance volumes, all data changes in the active file system with reference to the common Snapshot copy are preserved in a locked Snapshot copy until the expiry time that is set for the current volume.

If the volume expiry time is in the past or has not been set, then the Snapshot copy and the common Snapshot copy are locked for a duration of 30 days. All of the intermediate Snapshot copies between the common Snapshot copy and the latest locked Snapshot copy are deleted.

- For all volumes other than SnapLock compliance volumes, the resynchronization operation might cause loss of newer data that is written to the source volume after the base Snapshot copy was created.
- When you perform reverse resynchronization, the mirror policy of the relationship is set to DPDefault, and the mirror schedule is set to None.

- Step 1. Click **Protection → Relationships**.

- Step 2. Select the mirror relationship to update and click the more icon  .

- Step 3. Select **Resync**.

- Step 4. In the displayed dialog box, click **Resync**.

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.

- Step 1. Click **Protection → Relationships**.

- Step 2. Select the mirror relationship to update and click the more icon  .

- Step 3. Select **Abort**.

- Step 4. In the displayed dialog box, click **Abort**.

Restoring a volume in a mirror relationship

For a version-independent mirror relationship, you can use Storage 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.

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 cannot perform a restore operation on SnapLock volumes.
- 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.

Step 1. Click **Protection → Relationships**.

Step 2. Select the mirror relationship to update and click the more icon .

Step 3. Select **Restore**.

Step 4. Select the volume to restore to.

Step 5. Click **Save**.

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.

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, ThinkSystem Storage Manager for DM Series provides different types of protection relationships that enable you to protect data against accidental, malicious, or disaster-induced loss of data.

Asynchronous replication type

Mirror relationship

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.

Notes:

- The SnapMirror license is required to enable mirror relationship.
- 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.

Vault relationship

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.

Note: The SnapMirror or SnapVault license is required to enable vault relationship.

Mirror and vault relationship

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.

Notes:

- The SnapMirror license is required to enable mirror and vault relationship.
- 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.

Synchronous replication policy (SnapMirror Synchronous license required)**StrictSync**

A StrictSync replication policy will impose input/output (I/O) restrictions on the source volume in case of a replication failure post initialization. A StrictSync replication policy provides data protection by ensuring that the source volume and the destination volume are up to date.

Notes:

- If the destination is not Data Protection Optimization (DPO), then the SnapMirror license is required on the source cluster and the destination cluster and the SnapMirror Synchronous license is required on the source cluster.
- If the destination is DPO, then the SnapMirror Synchronous license and the SnapMirror license is required on the source cluster and the DPO license is required on the destination cluster.

Sync

A Sync replication policy does not impose I/O restrictions on the source volume in case of a replication failure post initialization. A Sync replication policy does not transfer data to destination volume after the failure. You need to perform a resynchronization operation to ensure that the source volume and destination volume are up to date.

Notes:

- If the destination is not Data Protection Optimization (DPO), then the SnapMirror license is required on the source cluster and the destination cluster and the SnapMirror Synchronous license is required on the source cluster.
- If the destination is DPO, then the SnapMirror Synchronous license and the SnapMirror license is required on the source cluster and the DPO license is required on the destination cluster.

Understanding workloads supported by StrictSync and Sync policies

StrictSync and Sync policies support all LUN-based applications with FC and iSCSI protocols, as well as NFSv3 and NFSv4 protocols for enterprise applications such as databases, VMWare, quota, CIFS, and so on. Starting with ONTAP 9.6, SnapMirror Synchronous can be used for enterprise file services such as electronic design automation (EDA), home directories, and software build workloads.

For a Sync policy, you need to consider a few important aspects while selecting the NFSv3 or NFSv4 workloads. The amount of data read or write operations by workloads is not a consideration, as Sync policy can handle high read or write IO workloads. In ONTAP 9.5, workloads that have excessive file creation, directory creation, file permission changes, or directory permission changes may not be suitable (these are referred to as high-metadata workloads). A typical example of a high-metadata workload is a DevOps workload in which you create multiple test files, run automation, and delete the files. Another example is parallel build workload that generate multiple temporary files during compilation. The impact of a high rate of write metadata activity is that it can cause synchronization between mirrors to temporarily break which stalls the read and write IOs from the client.

Starting with ONTAP 9.6, these limitations are removed and SnapMirror Synchronous can be used for enterprise file services workloads that include multiuser environments, such as home directories and software build workloads.

For information about best practices and sizing of StrictSync policy and Sync policy, see [Lenovo ONTAP & Storage Manager Documentation Resources](#) page.

SnapMirror licensing

A SnapMirror license is required on both the source and destination clusters, with limited exceptions as defined below. A SnapVault license is not required if a SnapMirror license is already installed.

DP_Optimized (DPO) license

Starting with ONTAP 9.4, 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.

SnapMirror Synchronous license

Starting with ONTAP 9.5, SnapMirror Synchronous relationships are supported. You require the following licenses for creating a SnapMirror Synchronous relationship:

- The SnapMirror Synchronous license is required on the source cluster.
- The SnapMirror license is required on both the source cluster and the destination cluster.

SVM Relationships

Storage virtual machine (SVM) disaster recovery (DR) provides disaster recovery capability at the SVM level by enabling the recovery of the data that is present in the constituent volumes of the SVM and the recovery of the SVM configuration.

You can use Storage Manager to create and manage mirror relationships and mirror and vault relationships between SVMs.

Creating SVM relationships

You can use Storage Manager to create SVM relationships to transfer data from the source SVM to the destination SVM. Creating an SVM relationship helps in recovering from a disaster as data is available on the source SVM and on the destination SVM.

Before you begin

- The destination cluster and source cluster must be running ONTAP 9.5 or later.
- The destination cluster must not be in a MetroCluster configurations.
- Starting with Storage Manager 9.6, Fabric Pool is supported.

Step 1. Click **Protection → SVM Relationship → Create**.

Step 2. Select the SVM relationship type from the **SVM Relationship Type** list.

Step 3. From the Source Storage Virtual Machine pane, select the cluster and the SVM.

Step 4. Optional: To view SVMs that do not have the required permissions, click **Navigate to the source cluster**, and then provide the required permissions.

Step 5. From the Destination Storage Virtual Machine pane, specify the name of the SVM that will be created on the destination cluster.

Step 6. Select the option to copy the source SVM configuration.

- Step 7. Optional: Click , update the protection policy and protection schedule, select aggregate, and then initialize the protection relationship.
- Step 8. Click **Save** to create the SVM relationship.
The SVM Relationships: Summary window is displayed.
- Step 9. Click **Done** to complete the process.

Editing SVM relationships

You can use Storage Manager to modify the properties of an SVM relationship.

- Step 1. Click **Protection → SVM Relationship**.
- Step 2. Select the SVM relationship that you want to modify, and then click **Edit**.
- Step 3. Select the SVM relationship type.
If the SVM relationships were created before ONTAP 9.3, then changing the SVM relationship type from mirror to mirror and vault is not allowed.
- Step 4. Modify the protection policy, the protection schedule, and the option to copy the source SVM configuration, as required.
- Step 5. Click **Save** to save the changes.

Managing SVM relationships

You can use Storage Manager to perform various operations on SVM relationships such as initializing SVM relationships, updating SVM relationships, activating the destination SVM, resynchronizing data from the source SVM, resynchronizing data from the destination SVM, and reactivating the source SVM.

Before you begin

- To initialize the SVM relationship, the source and destination clusters must be in a healthy peer relationship.
- To update the SVM relationship, the SVM relationship must be initialized and in a Snapmirrored state.
- To reactivate the source SVM, the resynchronize data from the destination SVM (reverse resync) operation must have been performed.
- If you had selected the option to copy the source SVM configuration while creating the SVM relationship, then to activate the SVM relationship, the source SVM must be stopped.
- SnapMirror license must be enabled on the source cluster and destination cluster.
- The source cluster and destination cluster must be in a healthy peer relationship.
- The destination cluster must have space available.
- The source SVM must have permission for SVM peering.
- You must break the SVM relationship to activate destination SVM, resync from source SVM, resync from destination SVM (Reverse Resync), and reactivate source SVM.
- To reactivate the source SVM, the SVM reverse relationship must exist and be in a Snapmirrored state.

- Step 1. Click **Protection → SVM Relationship**.
- Step 2. Select the SVM relationship, and then perform the appropriate action:

If you want to...	Do the following...
Initialize the SVM relationship	1. Click Operations → Initialize. The Initialize dialog box is displayed. 2. Click Initialize.
Update the SVM relationship	1. Click Operations → Update. The Update dialog box is displayed. 2. Click Update.
Activate the destination SVM Activating the destination SVM involves quiescing scheduled SnapMirror transfers, aborting any ongoing SnapMirror transfers, breaking the SVM relationship, and starting the destination SVM.	1. Click Operations → Activate Destination SVM. The Activate Destination SVM dialog box is displayed. 2. Select the Ok to activate destination SVM and break the relationship checkbox. 3. Click Activate.
Resynchronize data from the source SVM The resync operation performs a rebaseline of the SVM configuration. You can resync from the source SVM to reestablish a broken relationship between the two SVMs. When the resync is complete, the destination SVM contains the same information as the source SVM and is scheduled for further updates.	1. Click Operations → Resync from Source SVM. The Resync from Source SVM dialog box is displayed. 2. Select the Ok to delete any newer data in the destination SVM checkbox. 3. Click Resync.
Resynchronize data from the destination SVM (Reverse Resync) You can resync from the destination SVM to create a new relationship between the two SVMs. During this operation, the destination SVM continues to serve data with the source SVM backing up the configuration and data of the destination SVM.	1. Click Operations → Resync from Destination SVM (Reverse ReSync). The Resync from Destination SVM (Reverse Resync) dialog box is displayed. 2. If the SVM has multiple relationships, select the This SVM has multiple relationships, Ok to release to other relationships checkbox. 3. Select the Ok to delete the new data in the source SVM checkbox. 4. Click Reverse Resync.
Reactivate the source SVM Reactivating the source SVM involves protecting and recreating the SVM relationships between the source and destination SVM. If you had selected the option to copy the source SVM configuration while creating the SVM relationship, then the destination SVM will stop processing data.	1. Click Operations → Reactivate Source SVM. The Reactivate Source SVM dialog box is displayed. 2. Click Initiate Reactivation to initiate reactivation to the destination SVM. 3. Click Done.

Appendix A. Contacting Support

You can contact Support to obtain help for your issue.

You can receive hardware service through a Lenovo Authorized Service Provider. To locate a service provider authorized by Lenovo to provide warranty service, go to <https://datacentersupport.lenovo.com/serviceprovider> and use filter searching for different countries. For Lenovo support telephone numbers, see <https://datacentersupport.lenovo.com/supportphonelist> for your region support details.

Appendix B. Notices

Lenovo may not offer the products, services, or features discussed in this document in all countries. Consult your local Lenovo representative for information on the products and services currently available in your area.

Any reference to a Lenovo product, program, or service is not intended to state or imply that only that Lenovo product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any Lenovo intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any other product, program, or service.

Lenovo may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document is not an offer and does not provide a license under any patents or patent applications. You can send inquiries in writing to the following:

*Lenovo (United States), Inc.
8001 Development Drive
Morrisville, NC 27560
U.S.A.
Attention: Lenovo Director of Licensing*

LENOVO PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. Lenovo may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

The products described in this document are not intended for use in implantation or other life support applications where malfunction may result in injury or death to persons. The information contained in this document does not affect or change Lenovo product specifications or warranties. Nothing in this document shall operate as an express or implied license or indemnity under the intellectual property rights of Lenovo or third parties. All information contained in this document was obtained in specific environments and is presented as an illustration. The result obtained in other operating environments may vary.

Lenovo may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Any references in this publication to non-Lenovo Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this Lenovo product, and use of those Web sites is at your own risk.

Any performance data contained herein was determined in a controlled environment. Therefore, the result obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Trademarks

LENOVO, LENOVO logo, and THINKSYSTEM are trademarks of Lenovo. All other trademarks are the property of their respective owners. © 2023 Lenovo.