



Manage logical storage

System Manager Classic

NetApp
January 21, 2022

Table of Contents

- Managing logical storage 1
 - Storage Virtual Machines 1
 - Volumes 15
 - Junction Path 65
 - Shares 67
 - LUNs 73
 - Qtrees 100
 - Quotas 105
 - CIFS protocol 111
 - NFS protocol 120
 - NVMe protocol 122
 - iSCSI protocol 132
 - FC/FCoE protocol 137
 - Export policies 139
 - Efficiency policies 144
 - Protection policies 147
 - QoS policy groups 150
 - NIS services 156
 - LDAP client services 157
 - LDAP configuration services 160
 - Kerberos realm services 161
 - Kerberos interface services 164
 - DNS/DDNS Services 166
 - Users 167
 - Roles 170
 - UNIX 171
 - Windows 173
 - Name mapping 184

Managing logical storage

You can use System Manager to manage the logical storage such as storage virtual machines (SVMs), volumes, Qtrees, protocols, policies and so on.

Storage Virtual Machines

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

Related information

[SAN administration](#)

[ONTAP concepts](#)

SVM Dashboard window

The dashboard provides a cumulative at-a-glance information about your storage virtual machine (SVM) and its performance. You can use the Dashboard window to view important information related to your SVM such as the protocols configured, the volumes that are nearing capacity, and the performance.

SVM Details

This window displays details about the SVM through various panels such as the Protocol Status panel, Volumes Nearing Capacity panel, Applications panel, and performance panel.

- **Protocol Status**

Provides an overview of the protocols that are configured for the SVM. You can click the protocol name to view the configuration.

If a protocol is not configured or if a protocol license is not available for the SVM, you can click the protocol name to configure the protocol or to add the protocol license.

- **Volumes Nearing Capacity**

Displays information about the volumes that are nearing capacity utilization of 80 percent or more and that require immediate attention or corrective action.

- **Applications**

Displays information about the top five applications of the SVM. You can view the top five applications based on either IOPS (from low to high or from high to low) or capacity (from low to high or from high to low). You must click the specific bar chart to view more information about the application. For capacity, the total space, used space, and available space are displayed, and for IOPS, the IOPS details are displayed. For L2/L3 applications, latency metrics are also displayed.



The used size displayed in the Applications window does not equal the used size in the CLI.

You can click **View details** to open the Applications window of the specific application. You can click **View**

all applications to view all of the applications for the SVM.

The refresh interval for the Applications panel is one minute.

- **SVM Performance**

Displays the performance metrics of the protocols in the SVM, including latency and IOPS.

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

The refresh interval for the SVM Performance panel is 15 seconds.

Monitoring SVMs

The dashboard in System Manager enables you to monitor the health and performance of a storage virtual machine (SVM).

Steps

1. Click **Storage > SVMs**.
2. Select the name the SVM that you want to monitor.
3. View the details in the dashboard panels.

Editing SVM settings

You can use System Manager to edit the properties of storage virtual machines (SVMs), such as the name service switch, name mapping switch, and aggregate list.

About this task

- You can edit the values of the following SVM properties:
 - Name service switch
 - Protocols that are enabled to serve data



The CIFS protocol that is configured on the SVM continues to serve data even when you disable the protocol on that SVM.

- The list of aggregates that are available to create volumes



For FlexVol volumes, you can assign aggregates only if you have delegated administration to an SVM administrator.

- System Manager does not display the values of the name service switch and the name mapping switch for an SVM that is created through the command-line interface or for the SVM services that are not configured and are not set to the default values by ONTAP.

You can use the command-line interface to view the services because the Services tab is disabled.

System Manager displays the name service switch and the name mapping switch of an SVM only when it is created by using System Manager or when the services of the SVM are set to the default values by ONTAP.

Steps

1. Click **Storage > SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **Details** tab, modify the required data protocols.
4. In the **Resource Allocation** tab, choose one of the following methods to delegate volume creation:

If you want to provision volume creation...	Then...
For all aggregates	Select the Do not delegate volume creation option.
For specific aggregates	<ol style="list-style-type: none">a. Select the Delegate volume creation option.b. Select the required aggregates for delegating volume creation.

5. In the **Service** tab, specify the name service switch sources for the required database types and the order in which they should be consulted to retrieve name service information.

The default values for each of the database types are as follows:

- hosts: files, dns
- namemap: files
- group: files
- netgroup: files
- passwd: files

6. Click **Save and Close**.

Related information

[How ONTAP name service switch configuration works](#)

Deleting SVMs

You can use System Manager to delete storage virtual machines (SVMs) that you no longer require from the storage system configuration.

Before you begin

You must have completed the following tasks:

1. Disabled the Snapshot copies, data protection (DP) mirrors, and load-sharing (LS) mirrors for all the volumes



You must use the command-line interface (CLI) to disable LS mirrors.

2. Deleted all the igroups that belong to the SVM manually if you are deleting SVMs
3. Deleted all the portsets

4. Deleted all the volumes in the SVM, including the root volume
5. Unmapped the LUNs, taken them offline, and deleted them
6. Deleted the CIFS server if you are deleting SVMs
7. Deleted any customized user accounts and roles that are associated with the SVM
8. Deleted any NVMe subsystems associated with the SVM using the CLI.
9. Stopped the SVM

About this task

When you delete SVMs, the following objects associated with the SVM are also deleted:

- LIFs, LIF failover groups, and LIF routing groups
- Export policies
- Efficiency policies

If you delete SVMs that are configured to use Kerberos, or modify SVMs to use a different Service Principal Name (SPN), the original service principal of the SVM is not automatically deleted or disabled from the Kerberos realm. You must manually delete or disable the principal. You must have the Kerberos realm administrator's user name and password to delete or disable the principal.

If you want to move data from an SVM to another SVM before you delete the first SVM, you can use the SnapMirror technology to do so.

Steps

1. Click **Storage > SVMs**.
2. Select the SVM that you want to delete, and then click **Delete**.
3. Select the confirmation check box, and then click **Delete**.

Starting SVMs

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

Steps

1. Click **Storage > SVMs**.
2. Select the SVM that you want to start, and then click **Start**.

Results

The SVM starts serving data to clients.

Stopping SVMs

You can use System Manager to stop a storage virtual machine (SVM) if you want to troubleshoot any issue with the SVM, delete the SVM, or stop data access from the SVM.

Before you begin

All the clients connected to the SVM must be disconnected.



If any clients are connected to the SVM when you stop it, data loss might occur.

About this task

- You cannot stop SVMs during storage failover (SFO).
- When you stop the SVM, an SVM administrator cannot log in to the SVM.

Steps

1. Click **Storage > SVMs**.
2. Select the SVM that you want to stop, and then click **Stop**.

Results

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.



SVM administrators cannot log in to System Manager.

SVM administrators might have all or some of the following administration capabilities:

- Data access protocol configuration

SVM administrators can configure data access protocols, such as NFS, CIFS, iSCSI, and Fibre Channel (FC) protocol (Fibre Channel over Ethernet or FCoE included).

- Services configuration

SVM administrators can configure services such as LDAP, NIS, and DNS.

- Storage management

SVM administrators can manage volumes, quotas, qtrees, and files.

- LUN management in a SAN environment
- Management of Snapshot copies of the volume
- Monitoring SVM

SVM administrators can monitor jobs, network connection, network interface, and the SVM health.

Related information

[ONTAP 9 Documentation Center](#)

Tracing file access to diagnose access errors on SVMs

Starting with System Manager 9.6, you can diagnose CIFS or NFS file access errors on a

storage virtual machine (SVM).


About this task

File access issues, such as an “access denied” error, are likely to occur when there are problems with a share configuration, permissions, or user mapping. You can use System Manager to help you resolve file access problems by viewing the access trace results for the file or share that a user wants to access. System Manager shows whether the file or share has effective read, write, or execute permissions and the reasons why access is or is not effective.

Steps

1. Click **Storage > SVMs**.
2. Select the SVM that contains the files or shares for which file access errors were received.
3. Click **Trace File Access**.

The Trace File Access window for the selected SVM shows the prerequisites and steps required to trace file access permissions.

4. Click **Continue** to begin the file tracing process.
5. Select the protocol that is used to access files or shares on the selected SVM.
6. In the **User Name** field, enter the name of the user who was trying to access the file or share.
7. Click  to specify more details to narrow the scope of the trace.

The Advanced Options dialog window allows you to specify the following details:

- **Client IP Address:** Specify the IP address of the client.
- **File:** Specify the file name or file path to trace.
- **Show in Trace Results:** Specify whether you want to view only access denied entries or all entries. Click **Apply** to apply the details you specified and to return to the Trace File Access window.

8. Click **Start Tracing**.

The trace is initiated and a results table is displayed. The table is empty until users receive errors when requesting file access. The results table is refreshed every 15 seconds and displays messages in reverse chronological order.

9. Notify the affected user or users that they should try accessing the files within the next 60 minutes.

Details of the denied file access requests are shown in the results table when errors occur for the specified username for the duration of the trace. The Reasons column identifies the problems that are preventing the user from accessing files and reasons why they occurred.

10. In the **Reasons** column of the result table, click **View Permissions** to view permissions for the file that the user is trying to access.
 - 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 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.
 - If you specified the CIFS protocol, the Effective File and Share Permissions dialog box displays, listing both file and share permissions associated with the share and file that the user is trying to access.
 - If you specified the NFS protocol, the Effective File Permissions dialog box displays, listing the file

permissions associated with the file that the user is trying to access. A check mark indicates that permissions are granted, and an “X” indicates that permissions are not granted.

Click **OK** to return to the Trace File Access window.

11. The results table displays read-only data. You can perform the following actions with the results of the trace:
 - Click **Copy to Clipboard** to copy the results to the clipboard.
 - Click **Export Trace Results** to export the results to a comma-separated values (CSV) file.
12. When you want to end the tracing operation, click **Stop Tracing**.

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.



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

Database type	Defines name service sources for...	Valid sources are...
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.

Related information

[Editing SVM settings](#)

Storage Virtual Machines window

You can use the Storage Virtual Machines window to manage your storage virtual machines (SVMs) and display information about them.

You cannot manage (create, delete, start, or stop) an SVM configured for disaster recovery (DR) by using System Manager. Also, you cannot view the storage objects associated with the SVM configured for disaster recovery in the application interface.

Command buttons

- **Create**

Opens the Storage Virtual Machine (SVM) Setup wizard, which enables you to create a new SVM.

- **Edit**

Opens the Edit Storage Virtual Machine dialog box, which enables you to modify the properties, such as the name service switch, name mapping switch, and aggregate list, of a selected SVM.

- **Delete**

Deletes the selected SVMs.

- **Start**

Starts the selected SVM.

- **Stop**

Stops the selected SVM.

- **SVM Settings**

Manages the storage, policies, and configuration for the selected SVM.

- **Protection Operations**

Provides the following options:

- **Initialize**

Enables you to initialize the SVM relationship to perform a baseline transfer from the source SVM to the destination SVM.

- **Update**

Enables you to update data from the source SVM to the destination SVM.

- **Activate Destination SVM**

Enables you to activate the destination SVM.

- **Resync from Source SVM**

Enables you to initiate resynchronization of the broken relationship.

- **Resync from Destination SVM (Reverse Resync)**

Enables you to resynchronize the relationship from the destination SVM to the source SVM.

- **Reactivate Source SVM**

Enables you to reactivate the source SVM.

- **Refresh**

Updates the information in the window.

- **Trace File Access**

Enables you to trace the accessibility of a file or share on the selected SVM for a specified username.

SVM list

The SVM list displays the name of each SVM and the allowed protocols on it.

You can view only data SVMs by using System Manager.

- **Name**

Displays the name of the SVM.

- **State**

Displays the SVM state, such as Running, Starting, Stopped, or Stopping.

- **Subtype**

Displays the subtype of the SVM, which can be one of the following:

- default

Specifies that the SVM is a data-serving SVM.

- dp-destination

Specifies that the SVM is configured for disaster recovery.

- sync-source

Specifies that the SVM is in the primary site of a MetroCluster configuration.

- sync-destination

Specifies that the SVM is in the surviving site of a MetroCluster configuration.

- **Allowed Protocols**

Displays the allowed protocols, such as CIFS and NFS, on each SVM.

- **IPspace**

Displays the IPspace of the associated SVM.

- **Volume Type**

Displays the allowed volume type, such as FlexVol volume, on each SVM.

- **Protected**

Displays whether the SVM is protected or not.

- **Configuration State**

Displays whether the configuration state of the SVM is locked or unlocked.

Details area

The area below the SVM list displays detailed information, such as the type of volumes allowed, language, and Snapshot policy, about the selected SVM.


You can also configure the protocols that are allowed on this SVM. If you have not configured the protocols while creating the SVM, you can click the protocol link to configure the protocol.

You cannot configure protocols for an SVM configured for disaster recovery by using System Manager.



If the FCP service is already started for the SVM, clicking the FC/FCoE link opens the Network Interfaces window.

The color indicates the status of the protocol configuration:

Status	Description
Green	LIFs exist and the protocol is configured. You can click the link to view the configuration details.  Configuration might be partially completed. However, service is running. You can create the LIFs and complete the configuration from the Network Interfaces window.
Yellow	Indicates one of the following: <ul style="list-style-type: none">• LIFs exist. Service is created but is not running.• LIFs exist. Service is not created.• Service is created. LIFs do not exist.
Grey	The protocol is not configured. You can click the protocol link to configure the protocol.

Status	Description
Grey border	The protocol license has expired or is missing. You can click the protocol link to add the licenses in the Licenses page.

You can also add the management interface and view details such as the protection relationships, protection policy, NIS domain, and so on.

The **Details** area also includes a link to view the Public SSL Certificate for an SVM. When you click this link, you can perform the following tasks:

- View certificate details, the serial number, the start date, and the expiration date.
- Copy the certificate to the clipboard.
- Email the certificate details.

Peer Storage Virtual Machines area

Displays a list of the SVMs that are peered with the selected SVM along with details of the applications that are using the peer relationship.

Trace File Access window

Starting with System 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.



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.



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.

Related information

[SMB/CIFS management](#)

[SMB/CIFS and NFS multiprotocol configuration](#)

Volumes

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

You can access all the volumes in the cluster by using the Volumes tab or you can access the volumes specific to an SVM by using **SVMs > Volumes**.



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

Related information

[ONTAP concepts](#)

[Logical storage management](#)

Editing volume properties

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

Before you begin

For enabling volume encryption, you must have installed the volume encryption license by using System 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.
- For Data ONTAP 8.3.1 clusters, you can enable both inline compression and background compression for Cloud Volumes ONTAP for AWS (AWS).

Compression is not supported for Data ONTAP Edge.

- You cannot rename a SnapLock Compliance volume.

Steps

1. Click **Storage > Volumes**.
2. From the drop-down menu in the **SVM** field, select the storage virtual machine (SVM) in which the volume that you want to edit resides.
3. Select the volume that you want to modify, and then click **Edit**.

The Edit Volume dialog box is displayed.

4. In the **General** tab, modify the following properties as required:
 - Change the volume name
 - Enable volume encryption

This option is available only if you have enabled the Volume Encryption license and if the corresponding platform supports encryption. You can perform key-manager set up from the CLI.

- Change the security style of the volume
 - Enable or disable thin provisioning
5. Click the **Storage Efficiency** tab, and enable storage efficiency by configuring the following properties:
 - Deduplication
 - Data compression You cannot enable background compression for a volume that is contained by an aggregate with All Flash Optimized personality. You can enable only inline compression for these volumes.

You can enable inline deduplication only on a volume that is contained by an aggregate with All Flash Optimized personality or on a volume in a Flash Pool aggregate.

6. For SnapLock volumes, click the **SnapLock** tab, and perform the following steps:
 - a. Specify the autocommit period.

The autocommit period determines how long a file in the volume must remain unchanged before the file is committed to WORM state.

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

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

- c. Select the default retention period.

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

7. Click the **Advanced** tab, and enable the following properties:

- If you want the volume to automatically grow when the used space in the volume is above the grow threshold, select **Grow**.
- If you want the volume to grow or shrink in size in response to the amount of used space, select **Grow or Shrink**.
 - a. Specify the maximum size to which the volume can grow.
- Enable automatic deletion of older Snapshot copies by choosing one of the following options:
 - Try
Deletes the Snapshot copies that are not locked by any other subsystems.
 - Destroy
Deletes the Snapshot copies that are locked by the data-backing functionality.
 - Disrupt
Deletes the Snapshot copies that can disrupt the data transfer.
- Select the caching policy that you want to assign to the volume.
This option is available only for FlexVol volumes in a Flash Pool aggregate.
- Select the retention priority for cached data in the volume.
This option is available only for FlexVol volumes in a Flash Pool aggregate.
- Specify the fractional reserve that you want to set for the volume.
- Update the access time for reading the file.
This option is disabled for SnapLock volumes.

8. Click **Save and Close**.

Related information

[Volumes window](#)

[Setting up CIFS](#)

Editing data protection volumes

You can use System Manager to modify the volume name for a data protection (DP) volume. If the source volume does not have storage efficiency enabled, you might want to enable storage efficiency only on the destination volume.

About this task

You cannot modify storage efficiency on a mirror DP volume.

Steps

1. Click **Storage > Volumes**.

2. From the drop-down menu in the **SVM** field, select the storage virtual machine (SVM) in which the DP volume that you want to edit resides.
3. Select the volume that you want to modify, and then click **Edit**.
4. In the **Edit Data Protection Volume** dialog box, modify the volume name.
5. Ensure the **Enable Storage Efficiency** option is selected.

If storage efficiency is already enabled on the volume, then the check box is selected by default.

6. Click the **Advanced** tab, and perform the following steps:
 - a. Select the caching policy that you want to assign to the volume.
 - b. Select the retention priority for the cached data in the volume.

These options are available only for data protection FlexVol volumes in a Flash Pool aggregate.

7. Click **Save**.

Deleting volumes

You can use System Manager to delete a 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.

Steps

1. Click **Storage > Volumes**.
2. From the drop-down menu in the **SVM** field, select the storage virtual machine (SVM) in which the volume that you want to delete resides.
3. Select the volumes that you want to delete.

4. [NOTE]

Verify that you have selected the correct volumes that you want to delete. When you delete a volume, all the data in the volume is destroyed, and you cannot recover that data.

+ Click **Delete**.

1. Select the confirmation check box, and then click **Delete**.

Related information

[Volumes window](#)

Create FlexClone volumes

You can use System Manager to create a FlexClone volume when you require a writable, point-in-time copy of an existing FlexVol volume. You might want to create a copy of a volume for testing or to provide access to the volume for additional users without giving them access to the production data.

Before you begin

- The FlexClone license must be installed on the storage system.
- The volume that you want to clone must be online and must be a non-root volume.

About this task

The base Snapshot copy that is used to create a FlexClone volume of a SnapMirror destination is marked as busy and cannot be deleted. If a FlexClone volume is created from a Snapshot copy that is not the most recent Snapshot copy, and that Snapshot copy no longer exists on the source volume, all SnapMirror updates to the destination volume fail.

Steps

1. Click **Storage > Volumes**.
2. From the drop-down menu in the **SVM** field, select **All SVMs**.
3. Select the FlexVol volume that you want to clone from the list of volumes.
4. Click **More Actions > Clone > Create > Volume**.
5. Type the name of the FlexClone volume that you want to create.
6. If you want to enable thin provisioning for the new FlexClone volume, select **Thin Provisioning**.

By default, this setting is the same as that of the parent volume.

7. Create a Snapshot copy or select an existing Snapshot copy that you want to use as the base Snapshot copy for creating the FlexClone volume.
8. Click **Clone**.

Related information

[Volumes window](#)

Create FlexClone files

You can use System Manager to create a FlexClone file, which is a writable copy of a parent file. You can use these copies to test applications.

Before you begin

- The file that is cloned must be part of the active file system.
- The FlexClone license must be installed on the storage system.

About this task

- FlexClone files are supported only for FlexVol volumes.

You can create a FlexClone file of a parent file that is within a volume by accessing the parent file from the volume in which it resides, not from the parent volume.

- You cannot create a FlexClone file on a SnapLock volume.

Steps

1. Click **Storage > Volumes**.
2. From the drop-down menu in the **SVM** field, select **All SVMs**.
3. Select the volume in which you want to create a FlexClone file from the list of volumes.
4. Click **More Actions > Clone > Create > File**.
5. In the **Create FlexClone File** dialog box, select the file that you want to clone, and then specify a name for the FlexClone file.
6. Click **Clone**.

Results

The FlexClone file is created in the same volume as the parent file.

Related information

[Volumes window](#)

Splitting a FlexClone volume from its parent volume

If you want a FlexClone volume to have its own disk space instead of using the disk space of its parent volume, you can split the volume from its parent by using System Manager. After the split, the FlexClone volume becomes a normal FlexVol volume.

Before you begin

The FlexClone volume must be online.

About this task

For systems that are *not* AFF 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.

Steps

1. Click **Storage > Volumes**.
2. From the drop-down menu in the **SVM** field, select **All SVMs**.
3. Select the FlexClone volume that you want to split from its parent volume.
4. Click **More Actions > Clone > Split**.
5. Confirm the FlexClone volume details for the clone-splitting operation, and then click **Start Split** in the confirmation dialog box.

Related information

[Volumes window](#)

Viewing the FlexClone volume hierarchy

You can use System Manager to view the hierarchy of FlexClone volumes and their parent volumes.

Steps

1. Click **Storage > Volumes**.
2. From the drop-down menu in the **SVM** field, select **All SVMs**.
3. Select the required volume from the list of volumes.
4. Click **More Actions > Clone > View Hierarchy**.

Results

Volumes that have at least one child FlexClone volume are displayed. The FlexClone volumes are displayed as children of their respective parent volumes.

Related information

[Volumes window](#)

Changing the status of a volume

You can use System Manager to change the status of a FlexVol volume when you want to take a volume offline, bring a volume back online, or restrict access to a volume.

Before you begin

- If you want a volume to be the target of a volume copy operation or a SnapMirror replication operation, the volume must be in the restricted state.
- If you want to take a NAS volume offline, the NAS volume must be unmounted.

About this task

You can take a volume offline to perform maintenance on the volume, to move the volume, or to destroy the volume. When a volume is offline, the volume is unavailable for read or write access by clients. You cannot take a root volume offline.

Steps

1. Click **Storage > Volumes**.

2. From the drop-down menu in the **SVM** field, select **All SVMs**.
3. Select the volume for which you want to modify the status.
4. From the **More Actions > Change status to** menu, select the required volume status.
5. Click **Ok** in the confirmation dialog box to change the volume status.

Related information

[Volumes window](#)

Viewing the list of saved Snapshot copies

You can use System Manager to view the list of all of the saved Snapshot copies for a selected volume from the Snapshot Copies tab in the lower pane of the Volumes window. You can use the list of saved Snapshot copies to rename, restore, or delete a Snapshot copy.

Before you begin

The volume must be online.

About this task

You can view Snapshot copies for only one volume at a time.

Steps

1. Click **Storage > Volumes**.
2. From the drop-down menu in the **SVM** field, select **All SVMs**.
3. Click the plus sign (+) next to the volume for which you want to view saved Snapshot copies.
4. Click the **Show More Details** link to view more information about the volume.
5. Click the **Snapshot Copies** tab.

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

Create Snapshot copies outside a defined schedule

You can use System Manager to create a Snapshot copy of a volume outside a defined schedule to capture the state of the file system at a specific point in time.

Steps

1. Click **Storage > Volumes**.
2. From the drop-down menu in the **SVM** field, select **All SVMs**.
3. Select the volume from the list of volumes.
4. Click **More Actions > Manage Snapshots > Create**.
5. In the **Create Snapshot Copy** dialog box, if you want to change the default name, specify a new name for the Snapshot copy.

Valid characters are ASCII characters, numerals, hyphens (-), underscores (_), periods (.), and the plus (+) symbol.

The default name of a Snapshot copy consists of the volume name and the timestamp.

6. Click **Create**.
7. Verify that the Snapshot copy that you created is included in the list of Snapshot copies in the **Snapshot Copies** tab.

Related information

[Volumes window](#)

Setting the Snapshot copy reserve

You can use System 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.

Steps

1. Click **Storage > Volumes**.
2. From the drop-down menu in the **SVM** field, select **All SVMs**.
3. Select the volume for which you want to set the Snapshot copy reserve.
4. Click **More Actions > Manage Snapshots > Configuration Settings**.
5. Type or select the percentage of volume space that you want to reserve for the Snapshot copies, and then click **OK**.

Related information

[Volumes window](#)

Hiding the Snapshot copy directory

You can use System Manager to hide the Snapshot copy directory (`.snapshot`) so that the Snapshot copy directory is not visible when you view your volume directories. By default, the `.snapshot` directory is visible.

Steps

1. Click **Storage > Volumes**.
2. From the drop-down menu in the **SVM** field, select **All SVMs**.
3. Select the volume for which you want hide the Snapshot copy directory.
4. Click **More Actions > Manage Snapshots > Configuration Settings**.
5. Ensure that the **Make snapshot directory (.snapshot) visible** option is not selected, and then click **OK**.

Related information

[Volumes window](#)

Scheduling automatic creation of Snapshot copies

You can use System Manager to set up a schedule for the automatic creating automatic Snapshot copies of a volume. You can specify the time and frequency of creating the copies. You can also specify the number of Snapshot copies that are saved.

Steps

1. Click **Storage > Volumes**.
2. From the drop-down menu in the **SVM** field, select **All SVMs**.
3. Select the required volume from the list of volumes.
4. Click **More Actions > Manage Snapshots > Configuration Settings**.
5. In the **Configure Volume Snapshot Copies** dialog box, select **Enable scheduled Snapshot Copies**.
6. Select a Snapshot policy.

You can schedule the creation of only policy-based Snapshot copies.

7. Click **OK** to save your changes and start your Snapshot copy schedule.

Related information

[Volumes window](#)

Restoring a volume from a Snapshot copy

You can use System Manager to restore a volume to a state that is recorded in a previously created Snapshot copy to retrieve lost information. When you restore a volume from a Snapshot copy, the restore operation overwrites the existing volume configuration. Any changes that were made to the data in the volume after the Snapshot copy was created are lost.

Before you begin

- The SnapRestore license must be installed on your system.
- If the FlexVol volume that you want to restore contains a LUN, the LUN must be unmounted or unmapped.
- There must be enough space available for the restored volume.
- Users accessing the volume must be notified that you are going to revert a volume, and that the data from the selected Snapshot copy replaces the current data in the volume.

About this task

- If the volume that you restore contains junction points to other volumes, the volumes that are mounted on these junction points will not be restored.
- You cannot restore Snapshot copies for SnapLock Compliance volumes.

Steps

1. Click **Storage > Volumes**.
2. From the drop-down menu in the **SVM** field, select **All SVMs**.
3. Select the volume that you want to restore from a Snapshot copy.

4. Click **More Actions > Manage Snapshots > Restore**.
5. Select the appropriate Snapshot copy, and then click **Restore**.
6. Select the confirmation check box, and then click **Restore**.

Related information

[Volumes window](#)

Extending the expiry date of Snapshot copies

You can use System Manager to extend the expiry date of the Snapshot copies in a volume.

Before you begin

The SnapLock license must be installed on your system.

About this task

You can extend the expiry date only for Snapshot copies in a data protection (DP) volume that is the destination in a SnapLock for SnapVault relationship.

Steps

1. Click **Storage > Volumes**.
2. From the drop-down menu in the **SVM** field, select **All SVMs**.
3. Select a volume.
4. Click **Show More Details** to view more information about the volume.
5. Click the **Snapshot Copies** tab.

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

6. Select the Snapshot copy that you want to modify, and then click **Extend Expiry Date**.
7. In the **Extend Expiry Date** dialog box, specify the expiry date.

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

8. Click **OK**.

Renaming Snapshot copies

You can use System Manager to rename a Snapshot copy to help you organize and manage your Snapshot copies.

About this task

You cannot rename the Snapshot copies (which are committed to the WORM state) of a SnapLock DP volume that is in a SnapVault relationship.

Steps

1. Click **Storage > Volumes**.
2. From the drop-down menu in the **SVM** field, select **All SVMs**.

3. Click the required volume.
4. Click the **Show More Details** link to view more information about the volume.
5. Click the **Snapshot Copies** tab.

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

6. Select the Snapshot copy that you want to rename, and then click **More Actions > Rename**.
7. Specify a new name, and then click **Rename**.

Valid characters are ASCII characters, numerals, hyphens (-), underscores (_), periods (.), and the plus (+) symbol.

8. Verify the Snapshot copy name in the **Snapshot Copies** tab of the **Volumes** window.

Related information

[Volumes window](#)

Deleting Snapshot copies

You can delete a Snapshot copy to conserve disk space or to free disk space by using System Manager. You can also delete a Snapshot copy if the Snapshot copy is no longer required.

Before you begin

If you want to delete a Snapshot copy that is busy or locked, you must have released the Snapshot copy from the application that was using the Snapshot copy.

About this task

- You cannot delete the base Snapshot copy in a parent volume if a FlexClone volume is using that Snapshot copy.

The base Snapshot copy is the Snapshot copy that is used to create a FlexClone volume. The base Snapshot copy always displays the status `busy` and Application Dependency as `busy, vclone` in the parent volume.

- You cannot delete a locked Snapshot copy that is used in a SnapMirror relationship.

The Snapshot copy is locked and is required for the next update.

- You cannot delete a Snapshot copy from a SnapLock DP volume that is used in a SnapVault relationship before the expiry time of the Snapshot copy.
- You cannot delete the unexpired Snapshot copies (which are committed to WORM state) of a SnapLock DP volume that is in a SnapVault relationship.

Steps

1. Click **Storage > Volumes**.
2. From the drop-down menu in the **SVM** field, select **All SVMs**.
3. Expand the required volume.
4. Click the **Show More Details** link to view more information about the volume.

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

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

6. Select the Snapshot copy that you want to delete.
7. Click **Delete**.
8. Select the confirmation check box, and then click **Delete**.

Related information

[Volumes window](#)

[ONTAP 9 Documentation Center](#)

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 System Manager to provide more free space.

About this task

- For a volume that is configured to grow automatically, you can modify the limit to which the volume can grow automatically based on the increased size of the volume.
- You cannot resize a data protection volume if its mirror relationship is broken or if a reverse resynchronization operation has been performed on the volume.

Instead, you must use the command-line interface (CLI).

Steps

1. Click **Storage > Volumes**.
2. From the drop-down menu in the **SVM** field, select **All SVMs**.
3. Select the volume that you want to resize.
4. Click **More Actions > Resize**.
5. Type or select information as prompted by the wizard.
6. Confirm the details, and then click **Finish** to complete the wizard.
7. Verify the changes that you made to the available space and the total space of the volume in the **Volumes** window.

Related information

[Volumes window](#)

Enabling storage efficiency on a volume

You can use System Manager to enable storage efficiency and to configure both deduplication and data compression or only deduplication on a volume to save storage space. If you have not enabled storage efficiency when you created the volume, you can do so later by editing the volume.

Before you begin

- The volume must be online.
- If you want to use a policy-based deduplication schedule, you must have created an efficiency policy.

About this task

- You can enable background compression only if you have enabled background deduplication.
- You can enable inline compression and inline deduplication with or without enabling background compression and background deduplication, respectively.
- You can enable inline deduplication only on volumes that are contained by an aggregate with All Flash Optimized personality and on volumes that are contained by a Flash Pool aggregate.
- Starting with System Manager 9.6, editing storage efficiency is supported for FlexGroup DP volumes.

Steps

1. Click **Storage > Volumes**.
2. From the drop-down menu in the **SVM** field, select **All SVMs**.
3. Select the volume for which you want to enable storage efficiency, and then click **Edit**.
4. In the **Edit Volume** dialog box, click **Storage Efficiency**.
5. Select the **Background Deduplication** check box.
6. Select one of the following methods to run deduplication:

If you want to run deduplication...	Then...
Based on a storage efficiency policy	<ol style="list-style-type: none">a. Ensure that the Policy based option is selected.b. Click Choose, and then select a storage efficiency policy.c. Click OK.
When required	Select the On-demand option.

7. Select the **Background Compression** check box to enable background compression.

You cannot enable background compression for a volume that is contained by an aggregate with All Flash Optimized personality.

8. Select the **Inline Compression** check box to compress data while it is being written to the volume.

By default, inline compression is enabled on volumes that are contained by an aggregate with All Flash Optimized personality.

9. Select the **Inline Deduplication** check box to run deduplication before data is written to the disk.

By default, inline deduplication is enabled on volumes that are contained by an aggregate with All Flash Optimized personality.

10. Click **Save and Close**.

Related information

Changing the deduplication schedule

You can use System Manager to change the deduplication schedule by choosing to run deduplication manually, automatically, or on a schedule that you specify.

Steps

1. Click **Storage > Volumes**.
2. From the drop-down menu in the **SVM** field, select **All SVMs**.
3. Select the read/write volume for which you want to modify the deduplication schedule.
4. Click **Edit**, and then click the **Storage Efficiency** tab.
5. Change the deduplication schedule as required.
6. Click **Save and Close**.

Related information

Running deduplication operations

You can use System Manager to run deduplication immediately after creating a FlexVol volume or to schedule deduplication to run at a specified time.

Before you begin

- Deduplication must be enabled on the volume.
- The volume must be online and mounted.

About this task

Deduplication is a background process that consumes system resources during the operation; therefore, it might affect other operations that are in progress. You must cancel deduplication before you can perform any other operation.

Steps

1. Click **Storage > Volumes**.
2. From the drop-down menu in the **SVM** field, select **All SVMs**.
3. Select the volume for which you want to run deduplication.
4. Click **More Actions > Storage Efficiency**.
5. If you are running deduplication on the volume for the first time, run deduplication on the entire volume data by selecting **Scan Entire Volume** in the **Storage Efficiency** dialog box.
6. Click **Start**.
7. View the last-run details of the deduplication operation in the **Storage Efficiency** tab of the **Volumes** window.

Related information

Moving FlexVol volumes between aggregates or nodes

You can nondisruptively move a FlexVol volume to a different aggregate or a different node for capacity utilization and improved performance by using System Manager.

Before you begin

If you are moving a data protection (DP) volume, the data protection mirror relationships must be initialized before you move the volume.

About this task

You cannot move SnapLock volumes between aggregates and nodes.

Steps

1. Click **Storage > Volumes**.
2. From the drop-down menu in the **SVM** field, select **All SVMs**.
3. Select the volume that you want to move.
4. Click **More Actions > Move**.
5. In the **Move Volume** dialog box, select the destination aggregate or node for the volume, and then change the tiering policy.



- You cannot change the tiering policy of a root volume.
- You cannot move the root volume to FabricPool.
- For read/write volumes, you can set the tiering policy as “back up” during the volume move.

The tiering policy changes to “snapshot-only” after the move.

- Capacity tier values that are displayed in the “Used After Move” in both the source aggregate and destination aggregate are estimated values.

For the exact values, you must navigate to the Aggregate window and view the details of a specific aggregate.

6. Click **Move**.

Manually triggering the cutover for volume move

For a volume move operation, you can use System Manager to manually trigger the cutover when the volume enters the cutover deferred phase. You can set the duration of the cutover and the cutover action to be performed by the system if the operation fails within that duration.

Steps

1. Click **Storage > SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. Click the **Volumes** tab.
4. Expand the volume for which the volume move operation has been initiated.

5. Click the **Show More Details** link to view more information about the volume.
6. In the **Overview** tab, click **Cutover**.
7. In the **Cutover** dialog box, click **Advanced Options**.
8. Specify the cutover action and the cutover window period.
9. Click **OK**.

Assigning volumes to Storage QoS

You can limit the throughput of FlexVol volumes and FlexGroup volumes by assigning them to storage Quality of Service (QoS) policy groups. You can assign storage QoS for new volumes, or you can modify the storage QoS details of the volumes that are already assigned to a policy group by using System Manager.

About this task

- You can assign storage QoS only to read/write (rw) volumes that are online.
- You cannot assign storage QoS to a volume if the following storage objects are assigned to a policy group:
 - Parent storage virtual machine (SVM) of the volume
 - Child LUNs of the volume
 - Child files of the volume
- You can assign storage QoS or modify the QoS details for a maximum of 10 volumes simultaneously.

Steps

1. Click **Storage > Volumes**.
2. From the drop-down menu in the **SVM** field, select **All SVMs**.
3. Select one or more volumes for which you want to assign storage QoS.
4. Click **More Actions > Storage QoS**.
5. In the **Quality of Service Details** dialog box, select the **Manage Storage Quality of Service** check box if you want to manage the workload performance of the FlexVol volume.

If some of the volumes that you selected are already assigned to a policy group, the changes that you make might affect the performance of these volumes.

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

If you want to...	Do this...
Create a new policy group	<p>a. Select New Policy Group.</p> <p>b. Specify the policy group name.</p> <p>c. Specify the minimum throughput limit.</p> <ul style="list-style-type: none"> ◦ In System Manager 9.5, you can set the minimum throughput limit only on a performance-based All Flash Optimized personality. In System Manager 9.6, you can also set the minimum throughput limit for ONTAP Select Premium systems. ◦ You cannot set the minimum throughput limit for volumes on a FabricPool-enabled aggregate. ◦ If you do not specify the minimum throughput value or if the minimum throughput value is set to 0, the system automatically displays “None” as the value. <p style="padding-left: 40px;">This value is case-sensitive.</p> <p>d. Specify the maximum throughput limit to prevent the workload of the objects in the policy group from exceeding the specified throughput limit.</p> <ul style="list-style-type: none"> ◦ The minimum throughput limit and the maximum throughput limit must be of the same unit type. ◦ If you do not specify the minimum throughput limit, then you can set the maximum throughput limit in IOPS, B/s, KB/s, MB/s, and so on. ◦ If you do not specify the maximum throughput value, the system automatically displays “Unlimited” as the value. <p style="padding-left: 40px;">This value is case-sensitive. The unit that you specify does not affect the maximum throughput.</p>

If you want to...	Do this...
Select an existing policy group	<p>a. Select Existing Policy Group, and then click Choose to select an existing policy group from the Select Policy Group dialog box.</p> <p>b. Specify the minimum throughput limit.</p> <ul style="list-style-type: none"> ◦ In System Manager 9.5, you can set the minimum throughput limit only on a performance-based All Flash Optimized personality. In System Manager 9.6, you can also set the minimum throughput limit for ONTAP Select Premium systems. ◦ You cannot set the minimum throughput limit for volumes on a FabricPool-enabled aggregate. ◦ If you do not specify the minimum throughput value or if the minimum throughput value is set to 0, the system automatically displays “None” as the value. <p style="padding-left: 40px;">This value is case-sensitive.</p> <p>c. Specify the maximum throughput limit to prevent the workload of the objects in the policy group from exceeding the specified throughput limit.</p> <ul style="list-style-type: none"> ◦ The minimum throughput limit and the maximum throughput limit must be of the same unit type. ◦ If you do not specify the minimum throughput limit, then you can set the maximum throughput limit in IOPS, B/s, KB/s, MB/s, and so on. ◦ If you do not specify the maximum throughput value, the system automatically displays “Unlimited” as the value. <p style="padding-left: 40px;">This value is case-sensitive. The unit that you specify does not affect the maximum throughput.</p> <p style="padding-left: 40px;">If the policy group is assigned to more than one object, the maximum throughput that you specify is shared among the objects.</p>

7. Click the link that specifies the number of volumes if you want to review the list of selected volumes, and then click **Discard** if you want to remove any volumes from the list.


The link is displayed only when multiple volumes are selected.

8. Click **OK**.

Create a mirror relationship from a source SVM

You can use System Manager to create a mirror relationship from the source storage virtual machine (SVM), and to assign a mirror policy and schedule to the mirror relationship. The mirror copy enables quick availability of data if the data on the source volume is corrupted or lost.

Before you begin

- The SnapMirror license must be enabled on the source cluster and destination cluster.
- 
- 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.
 - If you are connecting from a cluster running ONTAP 9.2 or earlier to a remote cluster on which Security Assertion Markup Language (SAML) authentication is enabled, password-based authentication must be enabled on the remote cluster.

About this task

- System Manager does not support a cascade relationship.

For example, a destination volume in a relationship cannot be the source volume in another relationship.

- You can create a mirror relationship between SnapLock volumes of the same type only.

For example, if the source volume is a SnapLock Enterprise volume, then the destination volume must also be a SnapLock Enterprise volume.

- If encryption is enabled on the source volume and the destination cluster is running a version of ONTAP software earlier than ONTAP 9.3, then encryption is disabled on the destination volume by default.
- 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 System 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.

Steps


1. Click **Storage > Volumes**.
2. From the drop-down menu in the **SVM** field, select **All SVMs**.

3. Select the volumes for which you want to create mirror relationships, and then click **More Actions > Protect**.

The **Protect** option is available only for a read/write volume.

4. Select the **Replication** type:

If you selected the replication type as...	Do this...
Asynchronous	<ol style="list-style-type: none"> Optional: If you do not know the replication type and relationship type, click Help me Choose, specify the values, and then click Apply. Select the relationship type. The relationship type can be mirror, vault, or mirror and vault. Select a cluster and an SVM. If the selected cluster is running a version of ONTAP software earlier than ONTAP 9.3, then only peered SVMs are listed. If the selected cluster is running ONTAP 9.3 or later, peered SVMs and permitted SVMs are listed. Modify the volume name suffix, if required.
Synchronous	<ol style="list-style-type: none"> Optional: If you do not know the replication type and relationship type, click Help me Choose, specify the values, and then click Apply. Select the synchronization policy. The synchronization policy can be StrictSync or Sync. Select a cluster and an SVM. If the selected cluster is running a version of ONTAP software earlier than ONTAP 9.3, then only peered SVMs are listed. If the selected cluster is running ONTAP 9.3 or later, peered SVMs and permitted SVMs are listed. Modify the volume name suffix, if required.

5. Click , update the protection policy and protection schedule, select **FabricPool-enabled aggregate**, and then initialize the protection relationship.
6. Click **Save**.

Results

A new destination volume of type *dp* is created with the following default settings:

- Autogrow is enabled.
- Compression is disabled.
- The language attribute is set to match the language attribute of the source volume.

If the destination FlexVol volume is on a different SVM than the source FlexVol volume, then a peer relationship is created between the two SVMs if the relationship does not already exist.

A mirror relationship is created between the source volume and the destination volume. The base Snapshot copy is transferred to the destination volume if you have opted to initialize the relationship.

Related information

[Protection window](#)

Create a vault relationship from a source SVM

You can use System Manager to create a vault relationship from the source storage virtual machine (SVM), and to assign a vault policy to the vault relationship to create a backup vault. In the event of data loss or corruption on a system, backed-up data can be restored from the backup vault destination.

Before you begin

- The SnapVault license or SnapMirror license must be enabled on both the source cluster and the destination cluster.



+

- 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.
- If you are connecting from a cluster running ONTAP 9.2 or earlier to a remote cluster on which Security Assertion Markup Language (SAML) authentication is enabled, password-based authentication must be enabled on the remote cluster.

About this task

- System Manager does not support a cascade relationship.

For example, a destination volume in a relationship cannot be the source volume in another relationship.

- You can create a lock-vault relationship only between a non-SnapLock (primary) volume and a SnapLock destination (secondary) volume.
- If encryption is enabled on the source volume and the destination cluster is running a version of ONTAP software earlier than ONTAP 9.3, then encryption is disabled on the destination volume by default.

Steps


1. Click **Storage > Volumes**.
2. From the drop-down menu in the **SVM** field, select **All SVMs**.
3. Select the volumes for which you want to create vault relationships, and then click **More Actions > Protect**.

The **Protect** option is available only for a read/write volume.

4. Select the **Replication** type:

If you selected the replication type as...	Do this...
Asynchronous	<ol style="list-style-type: none"> Optional: If you do not know the replication type and relationship type, click Help me Choose, specify the values, and then click Apply. Select the relationship type. The relationship type can be mirror, vault, or mirror and vault. Select a cluster and an SVM. If the selected cluster is running a version of ONTAP software earlier than ONTAP 9.3, then only peered SVMs are listed. If the selected cluster is running ONTAP 9.3 or later, peered SVMs and permitted SVMs are listed. Modify the volume name suffix, if required.

If you selected the replication type as...	Do this...
Synchronous	<p>a. Optional: If you do not know the replication type and relationship type, click Help me Choose, specify the values, and then click Apply.</p> <p>b. Select the synchronization policy.</p> <p>The synchronization policy can be StrictSync or Sync.</p> <p>c. Select a cluster and an SVM.</p> <p>If the selected cluster is running a version of ONTAP software earlier than ONTAP 9.3, then only peered SVMs are listed. If the selected cluster is running ONTAP 9.3 or later, peered SVMs and permitted SVMs are listed.</p> <p>d. Modify the volume name suffix, if required.</p>

5. Click , update the protection policy and protection schedule, enable SnapLock properties on the destination volume, select a FabricPool-enabled aggregate, and then initialize the protection relationship.
6. Click **Save**.


Related information

[Protection window](#)

Create a mirror and vault relationship from a source SVM

You can use System Manager to create a mirror and vault relationship from the source storage virtual machine (SVM). Creating this relationship enables you to better protect your data by periodically transferring data from the source volume to the destination volume. 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 8.3.2 or later.
 - The SnapMirror license must be enabled on the source cluster and destination cluster.
-  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 9.2 or earlier to a remote cluster on which Security Assertion Markup Language (SAML) authentication is enabled, password-based authentication must be enabled on the remote cluster.

About this task

- System Manager does not support a cascade relationship.

For example, a destination volume in a relationship cannot be the source volume in another relationship.

- If encryption is enabled on the source volume and the destination cluster is running a version of ONTAP software earlier than ONTAP 9.3, then encryption is disabled on the destination volume by default.
- 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 System 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.


Steps

1. Click **Storage > Volumes**.
2. From the drop-down menu in the **SVM** field, select **All SVMs**.
3. Select the volumes for which you want to create mirror and vault relationships, and then click **More Actions > Protect**.

The **Protect** option is available only for a read/write volume.

4. Select the **Replication** type:

If you selected the replication type as...	Do this...
Asynchronous	<p>a. Optional: If you do not know the replication type and relationship type, click Help me Choose, specify the values, and then click Apply.</p> <p>b. Select the relationship type.</p> <p>The relationship type can be mirror, vault, or mirror and vault.</p> <p>c. Select a cluster and an SVM.</p> <p>If the selected cluster is running a version of ONTAP software earlier than ONTAP 9.3, then only peered SVMs are listed. If the selected cluster is running ONTAP 9.3 or later, peered SVMs and permitted SVMs are listed.</p> <p>d. Modify the volume name suffix, if required.</p>
Synchronous	<p>a. Optional: If you do not know the replication type and relationship type, click Help me Choose, specify the values, and then click Apply.</p> <p>b. Select the synchronization policy.</p> <p>The synchronization policy can be StrictSync or Sync.</p> <p>c. Select a cluster and an SVM.</p> <p>If the selected cluster is running a version of ONTAP software earlier than ONTAP 9.3, then only peered SVMs are listed. If the selected cluster is running ONTAP 9.3 or later, peered SVMs and permitted SVMs are listed.</p> <p>d. Modify the volume name suffix, if required.</p>

5. Click , update the protection policy and protection schedule, select **FabricPool-enabled aggregate**, and then initialize the protection relationship.

6. Click **Save**.

Create an NFS datastore for VMware

You can use the Create NFS Datastore for VMware wizard in System Manager to create an NFS datastore for VMware. You can create a volume for the NFS datastore and specify the ESX servers that can access the NFS datastore.

Before you begin

The NFS service must be licensed.

Steps

1. Click **Storage > Volumes**.
2. From the drop-down menu in the **SVM** field, select **All SVMs**.
3. Select the volume, and then click **More Actions > Provision Storage for VMware**.
4. In the **Create NFS Datastore for VMware** wizard, type or select information as required.
5. Confirm the details, and then click **Finish** to complete the wizard.

Changing the tiering policy of a volume

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

Steps

1. Click **Storage > Volumes**.
2. From the drop-down menu in the **SVM** field, select **All SVMs**.
3. Select the volume for which you want to change the tiering policy, and then click **More Actions > Change Tiering Policy**.
4. Select the required tiering policy from the **Tiering Policy** list, and then click **Save**.

Create FlexGroup volumes

A FlexGroup volume can contain many volumes that can be administered as a group instead of individually. You can use System Manager to create a FlexGroup volume by selecting specific aggregates or by selecting system-recommended aggregates.

About this task

- You can create only read/write (rw) FlexGroup volumes.
- Starting with System Manager 9.6, you can create FlexGroup volumes in a MetroCluster configuration.

Steps

1. Click **Storage > Volumes**.
2. Click **Create > Create FlexGroup**.
3. In the **Create FlexGroup** window, specify a name for the FlexGroup volume.

By default, the aggregates are selected according to best practices.

4. Click the **Volume Encryption** button to enable encryption for the volume.

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

Turning on volume encryption might affect the cross-volume storage efficiency when the selected aggregates are encrypted.

5. Specify a size for the FlexGroup volume.



You must also specify the measurement units.

6. Enable the **FabricPool** toggle button to use FabricPool aggregates in the FlexGroup volume.

- When you enable **FabricPool**, you can select the Tiering policy from the following choices in the drop-down menu:
 - **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 System Manager 9.5)**

Moves the newly transferred data of a data protection (DP) volume to the cloud tier.
 - **All (starting with System Manager 9.6)**

Moves all data to the cloud tier.
 - **None**

Prevents the data on the volume from being moved to a cloud tier.
- If you leave **FabricPool** in the “not enabled” position, only non-FabricPool aggregates are included in the created FlexGroup volume, and the tiering policy is set to “None”.
- If no FabricPool aggregates exist in the SVM, then **FabricPool** displays in the “not enabled” position and cannot be changed.
- If only FabricPool aggregates exist in the SVM, then the **FabricPool** button is displays in the “enabled” position and cannot be changed.

7. If you want to specify particular aggregates, click  (advanced options).

The aggregates associated with the FlexGroup volume you are creating are selected by default, according to best practices. They are displayed next to the **Aggregates** label.

8. In the **Protection** section, perform the following actions:

- a. Enable the **Volume Protection** option.
- b. Select the **Replication** type.



The **Synchronous** replication type is not supported for FlexGroup volumes.

- c. Click **Help me Choose**, if you do not know the replication type and relationship type.
 - Specify the values and click **Apply**.

The replication type and the relationship type is automatically selected based on the values specified.

- d. Select the relationship type.

The relationship types can be mirror, vault, or mirror and vault.

- e. Select a cluster and an SVM for the destination volume.

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

- f. Modify the volume name suffix as required.

9. Click **Create** to create the FlexGroup volume.

Related information

[Volumes window](#)

Viewing FlexGroup volume information

You can use System Manager to view information about a FlexGroup volume. You can view a graphical representation of the space allocated, the protection status, and the performance of a FlexGroup volume.

About this task

You can also view the Snapshot copies that are available for the FlexGroup volume, the data protection relationships for the FlexGroup volume, and the average performance metrics, read performance metrics, and write performance metrics of the FlexGroup volume based on latency, IOPS, and throughput.

Steps

1. Click **Storage > Volumes**.
2. From the drop-down menu in the **SVM** field, select **All SVMs**.
3. From the displayed list of FlexGroup volumes, select the FlexGroup volume about which you want to view information.

The information about the FlexGroup volume, the space allocated to the FlexGroup volume, the protection status of the FlexGroup volume, and the performance information about the FlexGroup volume are displayed.

4. Click the **Show More Details** link to view more information about the FlexGroup volume.
5. Click the **Snapshot Copies** tab to view the Snapshot copies of the FlexGroup volume.
6. Click the **Data Protection** tab to view the data protection relationships for the FlexGroup volume.
7. Click the **Storage Efficiency** tab to view the storage efficiency settings.
8. Click the **Performance** tab to view the average performance metrics, read performance metrics, and write performance metrics of the FlexGroup volume based on latency, IOPS, and throughput.

Related information

[Volumes window](#)

Editing FlexGroup volumes

Starting with System 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.


Steps

1. Click **Storage > Volumes**.
2. From the drop-down menu in the **SVM** field, select **All SVMs**.
3. Select the FlexGroup volume that you want to modify, and click **Edit**.
4. If you want to rename the FlexGroup volume, enter the new name in the **Name** field.

Starting with System Manager 9.6, you can also rename FlexGroup DP volumes.


5. Enable the **Encrypted** option to enable encryption for the volume.

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

6. Specify the percentage of the Snapshot copy reserve.
7. Click  to modify the FlexGroup volume settings. Refer to [Specifying advanced options for a FlexGroup volume](#).
8. Specify the size to which you want to resize the FlexGroup volume.

By default, existing aggregates are used to resize the FlexGroup volume. The minimum size that is allowed for the volume is displayed next to the size fields.



If you want to expand the FlexGroup volume by adding new resources, click  (advanced options). Refer to [Specifying advanced options for a FlexGroup volume](#).

9. Click **Save** to save the changes.

Related information

[Volumes window](#)

Specifying advanced options for a FlexGroup volume

When you create a FlexGroup volume, you can specify options you want to associate with the FlexGroup volume.

Steps

1. In the **Create FlexGroup** window, click  to specify the advanced options.

The Advanced Options window displays. It contains sections (the headings in the left column), in which you can specify various options.

2. In the **General Details** section, select the space reserve and security style, and then set the UNIX permission for the volume.

You should note the following limitations:

- The Space Reserve option is not available for FabricPool aggregates.
 - When thin provisioning is enabled, space is allocated to the volume from the aggregate only when data is written to the volume.
 - For All-Flash Optimized storage systems, thin provisioning is enabled by default, and for other storage systems, thick provisioning is enabled by default.
3. In the **Aggregates** section, you can enable the **Select Aggregates** button to override the best practices defaults and select your choices from a list of FabricPool aggregates.
 4. In the **Optimize Space** section, you can enable deduplication on the volume.

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

For systems with All Flash Optimized personality, inline compression and the `auto` deduplication schedule is enabled by default.

5. In the **QoS** (Quality of Service) section, specify the policy group to control the input/output (I/O) performance of the FlexGroup volume.
6. Click **Apply** to update the changes.

Resizing FlexGroup volumes

You can use System Manager to resize a FlexGroup volume by resizing existing resources or by adding new resources.

Before you begin


- To resize a FlexGroup volume, there must be enough free space on the existing aggregates.
- To expand a FlexGroup volume, there must be enough free space on the aggregate that you are using for expansion.

Steps

1. Click **Storage > Volumes**.
2. From the drop-down menu in the **SVM** field, select **All SVMs**.
3. Select the FlexGroup volume that you want to resize, and then click **More Actions > Resize**.
4. In the **Resize FlexGroup Volume** window, specify the size to which you want to resize the FlexGroup volume.

By default, existing aggregates are used to resize the FlexGroup volume. Starting with System Manager 9.6, the minimum size that is allowed for the volume is displayed next to the size fields.



If you want to expand the FlexGroup volume by adding new resources, click  (advanced options).

5. Specify the percentage of the Snapshot copy reserve.
6. Click **Resize** to resize the FlexGroup volume.

Related information

[Volumes window](#)

Changing the status of a FlexGroup volume

You can use System Manager to change the status of a FlexGroup volume when you want to take a FlexGroup volume offline, bring a FlexGroup volume back online, or restrict access to a FlexGroup volume.

About this task

System Manager does not support constituent-level management for FlexGroup volumes.

Steps

1. Click **Storage > Volumes**.
2. From the drop-down menu in the **SVM** field, select **All SVMs**.
3. Select the FlexGroup volume for which you want to modify the status.
4. Click **More Actions > Change status to**, and then update the FlexGroup volume status by selecting the required status.

Related information

[Volumes window](#)

Deleting FlexGroup volumes

You can use System Manager to delete a FlexGroup volume when you no longer require the FlexGroup volume.

Before you begin

- The junction path of the FlexGroup volume must be unmounted.
- The FlexGroup volume must be offline.

About this task

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

Steps

1. Click **Storage > Volumes**.
2. From the drop-down menu in the **SVM** field, select **All SVMs**.
3. Select the FlexGroup volume that you want to delete, and then click **Delete**.
4. Select the confirmation check box, and then click **OK**.

Related information

[Volumes window](#)

Create FlexCache volumes

Starting with System Manager 9.6, you can create a FlexCache volume.

About this task

You must have a FlexCache capacity license before you can create a FlexCache volume.

Steps

1. Click **Storage > Volumes**.
2. In the **Volumes** window, click **Create > FlexCache**.

The Create FlexCache volume window displays.

3. The following fields in the **Origin Volume** area display values for the origin volume for which you want to create a FlexCache volume. You can modify them.

- **Cluster:** Use the drop-down menu to select the cluster associated with the origin volume.
- **SVM:** Use the drop-down menu to select the SVM that contains the origin volume.

If you choose an SVM that is not peered, but is permitted to peer, System Manager allows you to peer it explicitly.

- **Volume:** Use the drop-down menu to select the volume name, or enter the name into the field.

4. The following fields in the **FlexCache Volume** area display default values for the FlexCache volume you are creating. You can modify them.

- **SVM:** Use the drop-down menu to select the SVM in which you want to create the FlexCache volume. If the FlexCache license capacity is full or almost full, you can select **Manage FlexCache license** to modify your license.
- **New Volume Name:** Enter a name for the FlexCache volume.
- **Size:** Specify the size for the FlexCache volume, including the measurement units.

The size field is initially set by default. The size you specify cannot exceed the licensed capacity size.

5. Click **Save** to create the FlexCache volume.

You can return to the **Volumes** window to view the FlexCache volume in the list of volumes.

Related information

[Volumes window](#)

Viewing FlexCache volume information

Starting with System Manager 9.6, you can view information about a FlexCache volume. You can view a graphical representation of the space allocated and the performance of a FlexCache volume.

Steps

1. Click **Storage > Volumes**.
2. From the drop-down menu in the **SVM** field, select **All SVMs**.
3. From the displayed list of volumes, select the FlexCache volume about which you want to view information.

The **Style** column displays "FlexCache" for a FlexCache volume.

When you make a selection, the Volume window for the selected FlexCache volume displays.

4. Initially, the **Volume** window displays the **Overview** tab. Click the tabs to view additional details about the FlexCache volume:

Click this tab...	To view these details...
Overview	General information about the FlexCache volume, the space allocated to the FlexCache volume, and performance information about the FlexCache volume.
Storage Efficiency	The storage efficiency settings of the FlexCache volume.
Performance	The average performance metrics, read performance metrics, and write performance metrics of the FlexCache volume based on latency, IOPS, and throughput. Also, the percentage of cache hits or cache misses is displayed.


5. Click **More actions** to view additional information and take actions from the selections in the drop-down menu:

Action	Description
Change status	Enables you to change the status of the FlexCache volume. Refer to Changing the status of a FlexCache volume .
Resize	Enables you to resize the FlexCache volume. Refer to Resizing FlexCache volumes .
Storage Efficiency	Enables you to adjust parameters to improve the storage efficiency of the FlexCache volume.
Storage QoS	Enables you to adjust the minimum and maximum storage limits for the FlexCache volume.
Encryption rekey	Enables you to reset the encryption key (only if you have enabled encryption on the peer cluster that includes the FlexCache volume)

Editing FlexCache volumes

Starting with System Manager 9.6, you can edit the properties of an existing FlexCache volume.

Steps

1. Click **Storage > Volumes**.
2. From the drop-down menu in the **SVM** field, select **All SVMs**.
3. Select the FlexCache volume that you want to modify, and click **Edit**.
4. Enter a new name for the FlexCache volume in the **Volume** field under **FlexCache Volume**.
5. Enter a new size for the FlexCache volume in the **Size** field under **FlexCache Volume**, and select the measurement unit from the drop-down menu.
6. Enable or disable encryption.
7. Click  to modify the FlexCache volume advanced settings. Refer to [Setting advanced settings for FlexCache volumes](#).
8. Click **Save** to save the changes.

Related information

[Volumes window](#)

Specifying advanced options for a FlexCache volume

Starting with System Manager 9.6, when you edit a FlexCache volume, you can specify the advanced options that you want to associate with the FlexCache volume.

Steps

1. In the **Edit FlexCache volume** window, click  to specify the advanced options.

The Advanced Options window displays. It contains sections (the headings in the left column), in which you can specify various options.

2. In the **General Details** section, you can edit the permissions for the volume.
3. In the **Aggregates** section, you can enable the **Select Aggregates** toggle button to override the best practices defaults and select your choices from a list of aggregates.
4. In the **Storage Efficiency** section, you can enable compression and deduplication on the volume.

Deduplication is not enabled by default for FlexCache volumes. System Manager uses the default deduplication schedule if the specified volume size exceeds the limit that is required for running deduplication.

5. Click **Apply** to update the changes.

Resizing FlexCache volumes

Starting with System Manager 9.6, you can resize a FlexCache volume by resizing existing resources or by adding new resources.

Before you begin


- To resize a FlexCache volume, there must be enough free space on the existing aggregates.
- To expand a FlexCache volume, there must be enough free space on the aggregate that you are using for expansion.

Steps

1. Click **Storage > Volumes**.
2. From the drop-down menu in the **SVM** field, select **All SVMs**.
3. Select the FlexCache volume that you want to resize, and then click **More Actions > Resize**.
4. In the **Resize FlexCache Volume** window, specify the size to which you want to resize the FlexCache volume.

By default, existing aggregates are used to resize the FlexCache volume. Starting with System Manager 9.6, the maximum size that is allowed for the volume is displayed next to the size field.



If you want to expand the FlexCache volume by adding new resources, click  (advanced options). Refer to [Specifying advanced options for FlexCache volumes](#).

5. Click **Save** to resize the FlexCache volume.

Related information

[Volumes window](#)

Changing the status of a FlexCache volume

Starting with System Manager 9.6, you can change the status of a FlexCache volume when you want to take it offline, bring a FlexCache volume back online, or restrict access to a FlexCache volume.

Steps

1. Click **Storage > Volumes**.
2. From the drop-down menu in the **SVM** field, select **All SVMs**.
3. Select the FlexCache volume for which you want to modify the status.
4. Click **More Actions > Change status to**, and then update the FlexCache volume status by selecting the required status.



To take a FlexCache volume offline and to change the status to "restricted", you must first unmount the volume.

Deleting FlexCache volumes

Starting with System Manager 9.6, you can delete a FlexCache volume when you no longer require it.

Before you begin

- The junction path of the FlexCache volume must be unmounted.
- The FlexCache volume must be offline.

Steps

1. Click **Storage > Volumes**.
2. From the drop-down menu in the **SVM** field, select **All SVMs**.
3. Select the FlexCache volume that you want to delete, and then click **Delete**.
4. Select the confirmation check box, and then click **OK**.

Related information

[Volumes window](#)

What NetApp Volume Encryption is

NetApp Volume Encryption is the process of protecting the user data, including the metadata, by encrypting the data before storing it on the disk. The data is decrypted and provided to the user only after proper authentication is provided.

To encrypt data, an encryption key is required. Each volume is assigned an encryption key to encrypt/decrypt operations of its data.

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



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 9.4, you can have less than 1024 Snapshot copies of a FlexVol volume.

How volume guarantees work for FlexVol volumes

Volume guarantees (sometimes called *space guarantees*) determine how space for a volume is allocated from its containing aggregate—whether or not the space is preallocated for the volume.

The guarantee is an attribute of the volume.

You set the guarantee when you create a new volume; you can also change the guarantee for an existing volume, provided that sufficient free space exists to honor the new guarantee.

Volume guarantee types can be `volume` (the default type) or `none`.

- A guarantee type of `volume` allocates space in the aggregate for the entire volume when you create the volume, regardless of whether that space is used for data yet.

The allocated space cannot be provided to or allocated for any other volume in that aggregate.

- A guarantee of `none` allocates space from the aggregate only as it is needed by the volume.

The amount of space consumed by volumes with this guarantee type grows as data is added instead of being determined by the initial volume size, which might leave space unused if the volume data does not grow to that size. The maximum size of a volume with a guarantee of `none` is not limited by the amount of free space in its aggregate. It is possible for the total size of all volumes associated with an aggregate to exceed the amount of free space for the aggregate, although the amount of space that can actually be used is limited by the size of aggregate.

Writes to LUNs or files (including space-reserved LUNs and files) contained by that volume could fail if the containing aggregate does not have enough available space to accommodate the write.

When space in the aggregate is allocated for a `volume` guarantee for an existing volume, that space is no longer considered free in the aggregate, even if the volume is not yet using the space. Operations that consume free space in the aggregate, such as creation of aggregate Snapshot copies or creation of new volumes in the containing aggregate, can occur only if there is enough available free space in that aggregate; these operations are prevented from using space already allocated to another volume.

When the free space in an aggregate is exhausted, only writes to volumes or files in that aggregate with preallocated space are guaranteed to succeed.

Guarantees are honored only for online volumes. If you take a volume offline, any allocated but unused space for that volume becomes available for other volumes in that aggregate. When you try to bring that volume back online, if there is insufficient available space in the aggregate to fulfill its guarantee, it will remain offline. You must force the volume online, at which point the volume's guarantee will be disabled.

Related information

[NetApp Technical Report 3965: NetApp Thin Provisioning Deployment and Implementation Data ONTAP 8.1 \(7-Mode\)](#)

FlexClone volumes and space guarantees

A FlexClone volume inherits its initial space guarantee from its parent volume. For example, if you create a FlexClone volume from a parent volume with a space guarantee of `volume`, then the FlexClone volume's initial space guarantee will be `volume` also. You can change the FlexClone volume's space guarantee.

For example, suppose that you have a 100-MB FlexVol volume with a space guarantee of `volume`, with 70 MB used and 30 MB free, and you use that FlexVol volume as a parent volume for a new FlexClone volume. The new FlexClone volume has an initial space guarantee of `volume`, but it does not require a full 100 MB of space from the aggregate, as it would if you had copied the volume. Instead, the aggregate needs to allocate only 30 MB (100 MB minus 70 MB) of free space to the clone.

If you have multiple clones with the same parent volume and a space guarantee of `volume`, they all share the same shared parent space with each other, so the space savings are even greater.



The shared space depends on the existence of the shared Snapshot copy (the base Snapshot copy that was used to create the FlexClone volume). If you delete this shared Snapshot copy, you lose the space savings provided by the FlexClone volume.

Thin provisioning for greater efficiencies using FlexVol volumes

With thin provisioning, when you create volumes and LUNs in a given aggregate, you do

not actually allocate any space for those in advance. The space is allocated as data is written to the volumes or LUNs.

The unused aggregate space is available to other volumes and LUNs. By allowing as-needed provisioning and space reclamation, thin provisioning can improve storage utilization and decrease storage costs.

A FlexVol volume can share its containing aggregate with other FlexVol volumes. Therefore, a single aggregate is the shared source of all the storage used by the FlexVol volumes it contains. Flexible volumes are no longer bound by the limitations of the disks on which they reside. A FlexVol volume can be sized based on how much data you want to store in it, rather than on the size of your disk. This flexibility enables you to maximize the performance and capacity utilization of the storage systems. Because FlexVol volumes can access all available physical storage in the system, improvements in storage utilization are possible.

Example

A 500-GB volume is allocated with only 100 GB of actual data; the remaining 400 GB allocated has no data stored in it. This unused capacity is assigned to a business application, even though the application might not need all 400 GB until later. The allocated but unused 400 GB of excess capacity is temporarily wasted.

With thin provisioning, the storage administrator provisions 500 GB to the business application but uses only 100 GB for the data. The difference is that with thin provisioning, the unused 400 GB is still available to other applications. This approach allows the application to grow transparently, and the physical storage is fully allocated only when the application needs it. The rest of the storage remains in the free pool to be used as needed.

Using space reservations with FlexVol volumes

Using space reservation, you can provision FlexVol volumes. Thin provisioning appears to provide more storage than is actually available from a given aggregate, as long as not all of that storage is currently being used.

Thick provisioning sets aside enough storage from the aggregate to ensure that any block in the volume can be written to at any time.

Aggregates can provide storage to volumes contained by more than one storage virtual machine (SVM). If you are using thin provisioning, and you need to maintain strict separation between your SVMs (for example, if you are providing storage in a multi-tenancy environment), you should either use fully allocated volumes (thick provisioning) or ensure that your aggregates are not shared between tenants.

When the space reserve is set to “Default”, the ONTAP space reservation settings apply to the volumes.

Related information

[NetApp Technical Report 3563: NetApp Thin Provisioning Increases Storage Utilization With On Demand Allocation](#)

[NetApp Technical Report 3483: Thin Provisioning in a NetApp SAN or IP SAN Enterprise Environment](#)

Benefits of storage efficiency

Storage efficiency enables you to store the maximum amount of data for the lowest cost and accommodate rapid data growth while consuming less space. You can use technologies such as RAID-DP, FlexVol, Snapshot copies, deduplication, data

compression, SnapMirror, and FlexClone to increase storage utilization and decrease storage costs. When used together, these technologies help to achieve increased performance.

- High-density disk drives, such as serial advanced technology attachment (SATA) drives mitigated with RAID-DP technology, provide increased efficiency and read performance.
- RAID-DP is a double-parity RAID6 implementation that protects against dual disk drive failures.
- Thin provisioning enables you to maintain a common unallocated storage space that is readily available to other applications as required.

It is based on FlexVol technology.

- Snapshot copies are a point-in-time, read-only view of a data volume, which consume minimal storage space.

Two Snapshot copies created in sequence differ only by the blocks added or changed in the time interval between the two. This block incremental behavior limits the associated consumption of storage capacity.

- Deduplication saves storage space by eliminating redundant data blocks within a FlexVol volume.
- Data compression stores more data in less space and reduces the time and bandwidth required to replicate data during volume SnapMirror transfers.

You have to choose the type of compression (inline or background) based on your requirement and the configurations of your storage system. Inline compression checks if data can be compressed, compresses data, and then writes data to the volume. Background compression runs on all the files, irrespective of whether the file is compressible or not, after all the data is written to the volume.

- SnapMirror technology is a flexible solution for replicating data over local area, wide area, and Fibre Channel networks.

It can serve as a critical component in implementing enterprise data protection strategies. You can replicate your data to one or more storage systems to minimize downtime costs in case of a production site failure. You can also use SnapMirror technology to centralize the backup of data to disks from multiple data centers.

- FlexClone technology copies data volumes, files, and LUNs as instant virtual copies.

A FlexClone volume, file, or LUN is a writable point-in-time image of the FlexVol volume or another FlexClone volume, file, or LUN. This technology enables you to use space efficiently, storing only data that changes between the parent and the clone.

- The unified architecture integrates multiprotocol support to enable both file-based and block-based storage on a single platform.

With FlexArray Virtualization, you can virtualize your entire storage infrastructure under one interface, and you can apply all the preceding efficiencies to your non-NetApp systems.

Data compression and deduplication

Beginning with Data ONTAP 8.0.1, data compression is supported with deduplication.


When both data compression and deduplication are enabled on a FlexVol volume, the data is first compressed

and then deduplicated. Therefore, deduplication can further increase the space savings by removing duplicate blocks in the FlexVol volume.

Though data compression and deduplication can be enabled on a FlexVol volume, the savings might not be the sum of the savings when each is run individually on a data set. The combined savings can yield higher savings than running deduplication or data compression individually.

You can achieve better savings when you run the data compression scanner before deduplication. This is because data compression scanner cannot run on data that is locked by deduplication, but deduplication can run on compressed data.

The following illustration shows how data is first compressed and then deduplicated:



When you run deduplication on a FlexVol volume that contains uncompressed data, it scans all the uncompressed blocks in the FlexVol volume and creates a digital fingerprint for each of the blocks.



If a FlexVol volume has compressed data, but the compression option is disabled on that volume, then you might lose the space savings when you run the `sis undo` command.

Guidelines for using deduplication

You must remember certain guidelines about system resources and free space when using deduplication.

The guidelines are as follows:

- If you have a performance-sensitive solution, you must carefully consider the performance impact of deduplication and measure the impact in a test setup before using deduplication.
- Deduplication is a background process that consumes system resources while it is running.

If the data does not change very often in a FlexVol volume, it is best to run deduplication less frequently. Multiple concurrent deduplication operations running on a storage system lead to a higher consumption of system resources.

- You must ensure that sufficient free space exists for deduplication metadata in the volumes and aggregates.
- If deduplication is used on the source volume, you must use deduplication on the destination volume.

- You must use automatic mode when possible so that deduplication runs only when significant additional data has been written to each FlexVol volume.
- You must run deduplication before creating a Snapshot copy to obtain maximum savings.
- You must set the Snapshot reserve to greater than 0 if Snapshot copies are used.

Options for resizing volumes

You can use the Volume Resize wizard to change your volume size, adjust the Snapshot reserve, delete Snapshot copies, and dynamically view the results of your changes.

The Volume Resize wizard displays a bar graph that displays the current space allocations within the volume, including the amount of used and free space. When you make changes to the size or Snapshot reserve of the volume, this graph is updated dynamically to reflect the changes.

You can also use the **Calculate space** button to determine the amount of space that is freed by deleting selected Snapshot copies.

You can use the Volume Resize wizard to make the following changes to your volume:

- **Change the volume size**

You can change the total volume size to increase or decrease storage space.

- **Adjust Snapshot reserve**

You can adjust the amount of space reserved for Snapshot copies to increase or decrease storage space.

- **Delete Snapshot copies**

You can delete Snapshot copies to reclaim volume space.



Snapshot copies that are in use cannot be deleted.

- **Autogrow**

You can specify the limit to which the volume can be grown automatically, if required.

Considerations when moving volumes

Moving a volume has many considerations and recommendations that are influenced by the volume you are moving or by the system configuration. You should understand the considerations associated with moving volumes.

- If you move a volume that has inline deduplication enabled from an aggregate with All Flash Optimized personality or a Flash Pool aggregate to an HDD aggregate, inline deduplication is disabled on the volume.
- If you move a volume that has background deduplication and inline compression enabled from an aggregate with All Flash Optimized personality to an HDD aggregate, then background compression, background deduplication, and inline compression are automatically enabled on the volume.
- If you move a volume that has background compression enabled from an HDD aggregate to an aggregate with All Flash Optimized personality, background compression is disabled on the volume.

- If you move a volume from a Flash Pool aggregate to a non-Flash Pool aggregate, the caching policies and retention priority are disabled.
- If you move a volume from a non-Flash Pool aggregate to a Flash Pool aggregate, the `default` caching policy and the `default` retention priority are automatically assigned to the volume.

Volumes window

You can use the Volumes window to manage your FlexVol volumes and FlexGroup volumes. Starting with System Manager 9.6, you can also manage FlexCache volumes.

You cannot view or manage volumes that are in storage virtual machines (SVMs) that are configured for disaster recovery by using System Manager. You must use the CLI instead.



The command buttons and list of columns will differ based on the type of volume that is selected. You can view only those command buttons and columns that are applicable for the selected volume.

Selection field

- **SVM selection pull-down menu**

Enables you to select all SVMs or a specific SVM to display in the list.

Command buttons

- **Create**

Provides the following options:

- **FlexVol**

Opens the Create Volume dialog box, which enables you to add FlexVol volumes.

- **FlexGroup**

Opens the Create FlexGroup window, which enables you to create FlexGroup volumes.

- **FlexCache**

Opens the Create FlexCache Volume window, which enables you to create FlexCache volumes.

- **Edit**

Enables you to edit the properties of the selected volume.

- **Delete**

Deletes the selected volume or volumes.

- **More Actions**

Provides the following options:

- **Change status to**

Changes the status of the selected volume to one of the following statuses:

- Online
- Offline
- Restrict

- **Resize**

Enables you to change the size of the volume.

For FlexGroup volumes, you can use existing resources to resize the volumes or you can add new resources to expand the volumes.

For FlexCache volumes, you can also add or remove an aggregate.

- **Protect**

Opens the Create Protection Relationship window for the volumes that are selected as source.

- **Manage Snapshots**

Provides a list of Snapshot options, including the following:

- **Create**

Displays the Create Snapshot dialog box, which you can use to create a Snapshot copy of the selected volume.

- **Configuration Settings**

Configures the Snapshot settings.

- **Restore**

Restores a Snapshot copy of the selected volume.

- **Clone**

Provides a list of clone options, including the following:

- **Create**

Creates a clone of the selected volume or a clone of a file from the selected volume.

- **Split**

Splits the clone from the parent volume.

- **View Hierarchy**

Displays information about the clone hierarchy.

- **Storage Efficiency**

Opens the Storage Efficiency dialog box, which you can use to manually start deduplication or to abort a running deduplication operation. This button is displayed only if deduplication is enabled on the storage system.

- **Move**

Opens the Move Volume dialog box, which you can use to move volumes from one aggregate or node to another aggregate or node within the same SVM.

- **Storage QoS**

Opens the Quality of Service details dialog box, which you can use to assign one or more volumes to a new or existing policy group.

- **Change Tiering Policy**

Enables you to change the tiering policy of the selected volume.

- **Volume Encryption Rekey**

Changes the data encryption key of the volume.

The data in the volume is re-encrypted using the new key that is automatically generated. The old key is automatically deleted after the rekey operation finishes.

Starting with System Manager 9.6, volume encryption rekey is supported for FlexGroup DP volumes and FlexCache volumes. Rekey is disabled for volumes that have inherited encryption from an NAE aggregate.



If you initiate a volume move operation when the rekey operation of the same volume is in progress, the rekey operation is aborted. In System Manager 9.5 and earlier version, if you try to move a volume when a conversion or rekey operation of a volume is in progress, then the operation is aborted without warning. Starting with System Manager 9.6, if you attempt a volume move during a conversion or rekey operation, a message is displayed warning that the conversion or rekey operation will be aborted if you continue.

- **Provision Storage for VMware**

Enables you to create a volume for the NFS datastore and to specify the ESX servers that can access the NFS datastore.

- **View Missing Protection Relationship**

Displays the read/write volumes that are online and are not protected, and displays the volumes that have protection relationships but are not initialized.

- **Reset Filters**

Enables you to reset the filters that were set to view missing protection relationships.

- **Refresh**

Updates the information in the window.



Enables you to select which details you want to display in the list on the Volumes window.

Volume list

- **Status**

Displays the status of the volume.

- **Name**

Displays the name of the volume.

- **Style**

In System Manager 9.5, this column displays the type of volume, such as FlexVol or FlexGroup. FlexCache volumes created by using the CLI are displayed as FlexGroup volumes.

In System Manager 9.6, this column displays the type of volume: FlexVol, FlexGroup, or FlexCache.

- **SVM**

Displays the SVM that contains the volume.

- **Aggregates**

Displays the name of the aggregates belonging to the volume.

- **Thin Provisioned**

Displays whether a space guarantee is set for the selected volume. Valid values for online volumes are Yes and No.

- **Root volume**

Displays whether the volume is a root volume.

- **Available Space**

Displays the available space in the volume.

- **Total Space**

Displays the total space in the volume, which includes the space that is reserved for Snapshot copies.

- **% Used**

Displays the amount of space (in percentage) that is used in the volume.

- **Logical Used %**

Displays the amount of logical space (in percentage), including space reserves, that is used in the volume.



This field is displayed only if you have enabled logical space reporting by using the CLI.

- **Logical Space Reporting**

Displays whether logical space reporting is enabled on the volume.



This field is displayed only if you have enabled logical space reporting by using the CLI.

- **Logical Space Enforcement**

Displays whether to perform logical space accounting on the volume.

- **Type**

Displays the type of volume: `rw` for read/write, `ls` for load sharing, or `dp` for data protection.

- **Protection Relationship**

Display whether the volume has a protection relationship initiated.

If the relationship is between an ONTAP system and a non-ONTAP system, the value is displayed as `No` by default.

- **Storage Efficiency**

Displays whether deduplication is enabled or disabled for the selected volume.

- **Encrypted**

Displays whether the volume is encrypted or not.

- **QoS Policy Group**

Displays the name of the Storage QoS policy group to which the volume is assigned. By default, this column is hidden.

- **SnapLock Type**

Displays the SnapLock type of the volume.

- **Clone**

Displays whether the volume is a FlexClone volume.

- **Is Volume Moving**

Displays whether a volume is being moved from one aggregate to another aggregate or from one node to another node.

- **Tiering Policy**

Displays the tiering policy of a FabricPool-enabled aggregate. The default tiering policy is “snapshot-only”.

- **Application**

Displays the name of the application that is assigned to the volume.

Overview area

You can click the plus sign (+) to the left in the row in which a volume is listed to view an overview of the details about that volume.

- **Protection**

Displays the **Data Protection** tab of the Volume window for the selected volume.

- **Performance**

Displays the **Performance** tab of the Volume window for the selected volume.

- **Show More Details**

Displays the Volume window for the selected volume.

Volume window for the selected volume

You can display this window by either of these methods:

- Clicking the volume name in the list of volumes on the Volumes window.
- Clicking **Show More Details** on the **Overview** area displayed for the selected volume.

The Volume window displays the following tabs:

- **Overview tab**

Displays general information about the selected volume, and displays a pictorial representation of the space allocation of the volume, the protection status of the volume, and the performance of the volume. The Overview tab displays details about the encryption of the volume, such as the encryption status and the encryption type, the conversion status or rekey status, information about a volume that is being moved, such as the state and phase of the volume move, the destination node and aggregate to which the volume is being moved, the percentage of volume move that is complete, the estimated time to complete the volume move operation, and details of the volume move operation. This tab also displays information about whether the volume is blocked for input/output (I/O) operations and the application blocking the operation.

For FlexCache volumes, details about the origin of the FlexCache volume are displayed.

The refresh interval for performance data is 15 seconds.

This tab contains the following command button:

- **Cutover**

Opens the Cutover dialog box, which enables you to manually trigger the cutover.

The **Cutover** command button is displayed only if the volume move operation is in the “replication” or “hard deferred” state.

- **Snapshot Copies tab**

Displays the Snapshot copies of the selected volume. This tab contains the following command buttons:

- **Create**

Opens the Create Snapshot Copy dialog box, which enables you to create a Snapshot copy of the selected volume.

- **Configuration Settings**

Configures the Snapshot settings.

- **More Actions › Rename**

Opens the Rename Snapshot Copy dialog box, which enables you to rename a selected Snapshot copy.

- **More Actions › Restore**

Restores a Snapshot copy.

- **More Actions › Extend Expiry Period**

Extends the expiry period of a Snapshot copy.

- **Delete**

Deletes the selected Snapshot copy.

- **Refresh**

Updates the information in the window.

- **Data Protection tab**

Displays data protection information about the selected volume.

If the source volume (read/write volume) is selected, the tab displays all of the mirror relationships, vault relationships, and mirror and vault relationships that are related to the destination volume (DP volume). If the destination volume is selected, the tab displays the relationship with the source volume.

If some or all of the cluster peer relationships of the local cluster are in an unhealthy state, the Data Protection tab might take some time to display the protection relationships relating to a healthy cluster peer relationship. Relationships relating to unhealthy cluster peer relationships are not displayed.

- **Storage Efficiency tab**

Displays information in the following panes:

- **Bar graph**

Displays (in graphical format) the volume space that is used by data and Snapshot copies. You can view details about the space used before and after applying settings for storage efficiency savings.

- **Details**

Displays information about deduplication properties, including whether deduplication is enabled on the volume, the deduplication mode, the deduplication status, type, and whether inline or background compression is enabled on the volume.

- Last run details

Provides details about the last-run deduplication operation on the volume. Space savings resulting from compression and deduplication operations that are applied on the data on the volume are also displayed.

- **Performance tab**

Displays information about the average performance metrics, read performance metrics, and write performance metrics of the selected volume, including throughput, IOPS, and latency.

Changing the client time zone or the cluster time zone impacts the performance metrics graphs. You must refresh your browser to view the updated graphs.

- **FlexCache tab**

Displays details about FlexCache volumes only if the volume you selected is an origin volume that has FlexCache volumes associated with it. Otherwise, this tab does not appear.

Related information

[Creating FlexVol volumes](#)

[Creating FlexClone volumes](#)

[Creating FlexClone files](#)

[Deleting volumes](#)

[Setting the Snapshot copy reserve](#)

[Deleting Snapshot copies](#)

[Creating Snapshot copies outside a defined schedule](#)

[Editing volume properties](#)

[Changing the status of a volume](#)

[Enabling storage efficiency on a volume](#)

[Changing the deduplication schedule](#)

[Running deduplication operations](#)

[Splitting a FlexClone volume from its parent volume](#)

[Resizing volumes](#)

[Restoring a volume from a Snapshot copy](#)

[Scheduling automatic creation of Snapshot copies](#)

[Renaming Snapshot copies](#)

[Hiding the Snapshot copy directory](#)

[Viewing the FlexClone volume hierarchy](#)

[Creating FlexGroup volumes](#)

[Editing FlexGroup volumes](#)

[Resizing FlexGroup volumes](#)

[Changing the status of a FlexGroup volume](#)

[Deleting FlexGroup volumes](#)

[Viewing FlexGroup volume information](#)

[Creating FlexCache volumes](#)

[Editing FlexCache volumes](#)

[Resizing FlexCache volumes](#)

[Deleting FlexCache volumes](#)

Junction Path

You can use the Junction Path window in System Manager to mount or unmount FlexVol volumes to a junction in the SVM namespace.

Mounting volumes

You can use System Manager to mount volumes to a junction in the storage virtual machine (SVM) namespace.

About this task

- If you mount a volume to a junction path with a language setting that is different from that of the immediate parent volume in the path, NFSv3 clients cannot access some of the files because some characters might not be decoded correctly.

This issue does not occur if the immediate parent directory is the root volume.

- You can mount a SnapLock volume only under the root of the SVM.
- You cannot mount a regular volume under a SnapLock volume.

Steps

1. Click **Storage > Junction Path**.
2. From the drop-down menu in the **SVM** field, select the SVM on which you want to mount a volume.
3. Click **Mount**, and then select the volume that is to be mounted.
4. If you want to change the default junction name, specify a new name.
5. Click **Browse**, and then select the junction path to which you want to mount the volume.
6. Click **OK**, and then click **Mount**.

7. Verify the new junction path in the **Details** tab.

Unmounting FlexVol volumes

You can use the Junction Path option of Storage pane in System Manager to unmount FlexVol volumes from a junction in the storage virtual machine (SVM) namespace.

Steps

1. Click **Storage > Junction Path**.
2. From the drop-down menu in the **SVM** field, select the SVM from which you want to unmount a volume.
3. Select the volumes that have to be unmounted, and then click **Unmount**.
4. Select the confirmation check box, and then click **Unmount**.

Changing export policies

When a volume is created, the volume automatically inherits the default export policy of the root volume of the storage virtual machine (SVM). You can use System Manager to change the default export policy that is associated with the volume to redefine the client access to data.

Steps

1. Click **Storage > Junction Path**.
2. From the drop-down menu in the **SVM** field, select the SVM in which the volume that you want to modify resides.
3. Select the volume, and then click **Change Export Policy**.
4. Select the export policy, and then click **Change**.
5. Verify that the **Export Policy** column in the **Junction Path** window displays the export policy that you applied to the volume.

Results

The default export policy is replaced with the export policy that you selected.

Junction Path window

You can use the Junction Path menu to manage the NAS namespace of storage virtual machines (SVMs).

Command buttons

- **Mount**

Opens the Mount Volume dialog box, which enables you to mount a volume to the junction in an SVM namespace.

- **Unmount**

Opens the Unmount Volume dialog box, which enables you to unmount a volume from its parent volume.

- **Change Export Policy**

Opens the Change Export Policy dialog box, which enables you to change the existing export policy associated with the volume.

- **Refresh**

Updates the information in the window.

Junction Path list

- **Path**

Specifies the junction path of the mounted volume. You can click the junction path to view the related volumes and qtrees.

- **Storage Object**

Specifies the name of the volume mounted on the junction path. You can also view the qtrees that the volume contains.

- **Export Policy**

Specifies the export policy of the mounted volume.

- **Security Style**

Specifies the security style for the volume. Possible values include UNIX (for UNIX mode bits), NTFS (for CIFS ACLs), and Mixed (for mixed NFS and CIFS permissions).

Details tab

Displays general information about the selected volume or qtree, such as the name, type of storage object, junction path of the mounted object, and export policy. If the selected object is a qtree, details about the space hard limit, space soft limit, and space usage are displayed.

Shares

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

Create a CIFS share

You can use System Manager to create a CIFS share that enables you to specify the folder, qtree, or volume that CIFS users can access.

Before you begin

You must have installed the CIFS license before you set up and start CIFS.

Steps

1. Click **Storage > Shares**.
2. From the drop-down menu in the **SVM** field, select the SVM on which you want to create a CIFS share.

3. Click **Create Share**.
4. In the **Create Share** window, click **Browse**, and then select the folder, qtree, or volume that should be shared.
5. Specify a name for the new CIFS share.
6. Select the **Enable continuous availability for Hyper-V and SQL** check box to permit clients that support SMB 3.0 and later to open files persistently during nondisruptive operations.

Files that are opened by using this option are protected from disruptive events such as failover, giveback, and LIF migration.

Starting with System Manager 9.6, continuous availability is supported for FlexGroup volumes.

7. Select the **Encrypt data while accessing this share** check box to enable SMB 3.0 encryption.
8. Provide a description or comment for the share, and then click **Create**.

Results

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

Related information

[Setting up CIFS](#)

[Shares window](#)

Stopping share access

You can use System Manager to stop a share when you want to remove the shared network access to a folder, qtree, or volume.

Before you begin

You must have installed the CIFS license.

Steps

1. Click **Storage > Shares**.
2. From the drop-down menu in the **SVM** field, select the SVM on which the CIFS share that you want to stop resides.
3. From the list of shares, select the share that you want to stop sharing, and then click **Stop Sharing**.
4. Select the confirmation check box, and then click **Stop**.
5. Verify that the share is no longer listed in the **Shares** window.

Related information

[Shares window](#)

Create home directory shares

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

Before you begin

CIFS must be set up and started.

Steps

1. Click **Storage > Shares**.
2. Click **Create Home Directory**, and then provide the pattern information that determines how a user is mapped to a directory.
3. Click **Create**.
4. Verify that the home directory that you created is listed in the **Shares** window.

Editing share settings

You can use System Manager to modify the settings of a share such as the symbolic link settings, share access permissions of users or groups, and the type of access to the share. You can also enable or disable continuous availability of a share over Hyper-V, and enable or disable access-based enumeration (ABE). Starting with System Manager 9.6, continuous availability is supported for FlexGroup volumes.

Steps

1. Click **Storage > Shares**.
2. Select the share that you want to modify from the list of shares, and then click **Edit**.
3. In the **Edit Share Settings** dialog box, modify the share settings as required:
 - a. In the **General** tab, enable continuous availability of a share over Hyper-V.

Enabling continuous availability permits SMB 3.0 and clients that support SMB 3.0 to open files persistently during nondisruptive operations. Files that are opened persistently are protected from disruptive events such as failover, giveback, and LIF migration.
 - b. In the **Permissions** tab, add users or groups, and then assign permissions to specify the type of access.
 - c. In the **Options** tab, select the required options.
4. Click **Save and Close**.
5. Verify the changes that you made to the selected share in the **Shares** window.

Related information

[Shares window](#)

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.

Shares window

You can use the Shares window to manage your shares and to view information about the shares.

Command buttons

- **Create Share**

Opens the Create Share dialog box, which enables you to create a share.

- **Create Home Directory**

Opens the Create Home Directory Share dialog box, which enables you to create a new home directory share.

- **Edit**

Opens the Edit Settings dialog box, which enables you to modify the properties of a selected share.

- **Stop Sharing**

Stops the selected object from being shared.

- **Refresh**

Updates the information in the window.

Shares list

The shares list displays the name and path of each share.

- **Share Name**

Displays the name of the share.

- **Path**

Displays the complete path name of an existing folder, qtree, or volume that is shared. Path separators can be backward slashes or forward slashes, although ONTAP displays all path separators as forward slashes.

- **Home Directory**

Displays the name of the home directory share.

- **Comment**

Displays additional descriptions of the share, if any.

- **Continuously Available Share**

Displays whether the share is enabled for continuous availability. Starting with System Manager 9.6, continuous availability is supported for FlexGroup volumes.

Details area

The area below the shares list displays the share properties and the access rights for each share.

- **Properties**

- Name

Displays the name of the share.

- Oplocks status

Specifies whether the share uses opportunistic locks (oplocks).

- Browsable

Specifies whether the share can be browsed by Windows clients.

- Show Snapshot

Specifies whether Snapshot copies can be viewed by clients.

- Continuously Available Share

Specifies whether the share is enabled or disabled for continuous availability. Starting with System Manager 9.6, continuous availability is supported for FlexGroup volumes.

- Access-Based Enumeration

Specifies whether access-based enumeration (ABE) is enabled or disabled on the share.

- BranchCache

Specifies whether BranchCache is enabled or disabled on the share.

- SMB Encryption

Specifies whether data encryption using SMB 3.0 is enabled at the storage virtual machine (SVM) level

or at the share level. If SMB encryption is enabled at the SVM level, SMB encryption applies for all of the shares and the value is shown as Enabled (at the SVM level).

- Previous Versions

Specifies whether the previous versions can be viewed and restored from the client.

- **Share access control**

Displays the access rights of the domain users, domain groups, local users, and local groups for the share.

Related information

[Creating a CIFS share](#)

[Stopping share access](#)

[Editing share settings](#)

LUNs

You can use System Manager to manage LUNs.

You can access all the LUNs in the cluster by using the LUNs tab or you can access the LUNs specific to the SVM by using **SVMs > LUNs**.



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

Related information

[SAN administration](#)

Create FC SAN optimized LUNs

You can use System Manager to create one or more FC SAN optimized LUNs during the initial setup of a cluster on an AFF platform.

Before you begin

- You must ensure that only one storage virtual machine (SVM) has been created with the name `AFF_SAN_DEFAULT_SVM`, and that this SVM does not contain any LUNs.
- You must have verified that the hardware setup has been completed successfully.

[ONTAP 9 Documentation Center](#)

About this task

- This method is available only during the initial setup of a cluster with two or more nodes.

System Manager uses only the first two nodes to create LUNs.

- Each LUN is created on a separate volume.
- Volumes are thin provisioned.

- Space reservation is disabled on the created LUNs.
- Most of the cluster configurations are already completed at the factory and are optimized for optimum storage efficiency and performance.

You must not modify these configurations.

Steps

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

After you create LUNs using this method, you cannot use this method again.

If you close the dialog box without creating LUNs, you must navigate to the LUNs tab and click **Create** to access the dialog box again.

2. In the **LUN details** area of the **Create LUNs** dialog box, specify the application type:

If the application type is...	Then...
Oracle	a. Specify the database name and size. b. If you have deployed Oracle Real Application Clusters (RAC), then select the Oracle RAC check box. Only two RAC nodes are supported. You must ensure that Oracle RAC has a minimum of two initiators added to the initiator group.
SQL	Specify the number of databases and the size of each database.
Other	a. Specify the name and size of each LUN. b. If you want to create more LUNs, click Add more LUNs , and then specify the name and size for each LUN.

Data, log, binary, and temporary LUNs are created based on the selected application type.

3. In the **Map to these Initiators** area, perform these steps:
 - a. Specify the initiator group name and the type of operating system.
 - b. Add the host initiator WWPN by selecting it from the drop-down list or by typing the initiator in the text box.
 - c. Add the alias for the initiator.

Only one initiator group is created.

4. Click **Create**.

A summary table is displayed with the LUNs that are created.

5. Click **Close**.

Related information

[ONTAP 9 Documentation Center](#)

Application-specific LUN settings

System Manager supports Oracle, SQL, and other application types while creating FC SAN optimized LUNs on an AFF cluster. LUN settings such as the LUN size are determined by rules specific to the application type. For SQL and Oracle, LUN settings are automatically created.

If your cluster contains two or more nodes, System Manager uses only the first two nodes selected by the API to create LUNs. Data aggregates are already created in each of the two nodes. The size of each volume created is equal to the available capacity of the aggregate. The volumes are thin-provisioned and space reservation is disabled on the LUNs.

Storage efficiency policy is enabled by default with the schedule set to “daily” and quality of service (QoS) set to “best_effort”. By default, access time (atime) update is enabled on the cluster. However, access time updates are disabled by System Manager while creating volumes and therefore every time a file is read or written, the access time field in the directory is not updated.



Enabling the access time update causes performance degradation to the data-serving capability of the cluster.

LUN settings for SQL

By default, LUNs and volumes are provisioned for a single instance of the SQL server with 2 databases of 1 TB each and 24 physical cores. Space is provisioned for LUNs and volumes according to specific rules for the SQL server. Load balancing is performed for LUNs across the HA pair. You can modify the number of databases. For each database, eight data LUNs and one log LUN is created. One temporary LUN is created for each SQL instance.

The following table provides information about how space is provisioned for the default values of SQL:

Node	Aggregate	LUN type	Volume name	LUN name	Formula for LUN size	LUN size (GB)
node1	node1_aggr1	data	db01_data01	db01_data01	Database size ÷ 8	125
		data	db01_data02	db01_data02	Database size ÷ 8	125
		data	db01_data03	db01_data03	Database size ÷ 8	125
		data	db01_data04	db01_data04	Database size ÷ 8	125

Node	Aggregate	LUN type	Volume name	LUN name	Formula for LUN size	LUN size (GB)
		data	db02_data01	db02_data01	Database size ÷ 8	125
		data	db02_data02	db02_data02	Database size ÷ 8	125
		data	db02_data03	db02_data03	Database size ÷ 8	125
		data	db02_data04	db02_data04	Database size ÷ 8	125
		log	db01_log	db01_log	Database size ÷ 20	50
		temp	sql_temp	sql_temp	Database size ÷ 3	330
node2	node2_aggr1	data	db01_data05	db01_data05	Database size ÷ 8	125
		data	db01_data06	db01_data06	Database size ÷ 8	125
		data	db01_data07	db01_data07	Database size ÷ 8	125
		data	db01_data08	db01_data08	Database size ÷ 8	125
		data	db02_data05	db02_data05	Database size ÷ 8	125
		data	db02_data06	db02_data06	Database size ÷ 8	125
		data	db02_data07	db02_data07	Database size ÷ 8	125
		data	db02_data08	db02_data08	Database size ÷ 8	125
		log	db02_log	db02_log	Database size ÷ 20	50

LUN settings for Oracle

By default, LUNs and volumes are provisioned for one database of 2 TB. Space is provisioned for LUNs and volumes according to specific rules for Oracle. By default, Oracle Real Application Clusters (RAC) is not selected.

The following table provides information about how space is provisioned for the default values of Oracle:

Node	Aggregate	LUN type	Volume name	LUN name	Formula for LUN size	LUN size (GB)
node1	node1_aggr1	data	ora_vol01	ora_lundata01	Database size ÷ 8	250
		data	ora_vol02	ora_lundata02	Database size ÷ 8	250
		data	ora_vol03	ora_lundata03	Database size ÷ 8	250
		data	ora_vol04	ora_lundata04	Database size ÷ 8	250
		log	ora_vol05	ora_lunlog1	Database size ÷ 40	50
		binaries	ora_vol06	ora_orabin1	Database size ÷ 40	50
node2	node2_aggr1	data	ora_vol07	ora_lundata05	Database size ÷ 8	250
		data	ora_vol08	ora_lundata06	Database size ÷ 8	250
		data	ora_vol09	ora_lundata07	Database size ÷ 8	250
		data	ora_vol10	ora_lundata08	Database size ÷ 8	250
		log	ora_vol11	ora_lunlog2	Database size ÷ 40	50

For Oracle RAC, LUNs are provisioned for grid files. Only two RAC nodes are supported for Oracle RAC.

The following table provides information about how space is provisioned for the default values of Oracle RAC:

Node	Aggregate	LUN type	Volume name	LUN name	Formula for LUN size	LUN size (GB)
node1	node1_aggr1	data	ora_vol01	ora_lundata01	Database size ÷ 8	250
		data	ora_vol02	ora_lundata02	Database size ÷ 8	250
		data	ora_vol03	ora_lundata03	Database size ÷ 8	250
		data	ora_vol04	ora_lundata04	Database size ÷ 8	250
		log	ora_vol05	ora_lunlog1	Database size ÷ 40	50
		binaries	ora_vol06	ora_orabin1	Database size ÷ 40	50
		grid	ora_vol07	ora_lungrid1	10 GB	10
node2	node2_aggr1	data	ora_vol08	ora_lundata05	Database size ÷ 8	250
		data	ora_vol09	ora_lundata06	Database size ÷ 8	250
		data	ora_vol10	ora_lundata07	Database size ÷ 8	250
		data	ora_vol11	ora_lundata08	Database size ÷ 8	250
		log	ora_vol12	ora_lunlog2	Database size ÷ 40	50
		binaries	ora_vol13	ora_orabin2	Database size ÷ 40	50

LUN settings for other application type

Each LUN is provisioned in a volume. The space is provisioned in the LUNs based on the specified size. Load balancing is performed across the nodes for all the LUNs.

Create LUNs

You can use System Manager to create LUNs for an existing aggregate, volume, or qtree when there is available free space. You can create a LUN in an existing volume or create a new FlexVol volume for the LUN. You can also enable storage Quality of Service (QoS) to manage the workload performance.

About this task

If you specify the LUN ID, System Manager checks the validity of the LUN ID before adding it. If you do not specify a LUN ID, ONTAP software automatically assigns one.

While selecting the LUN multiprotocol type, you should have considered the guidelines for using each type. The LUN Multiprotocol Type, or operating system type, determines the layout of data on the LUN, and the minimum and maximum sizes of the LUN. After the LUN is created, you cannot modify the LUN host operating system type.

In a MetroCluster configuration, System Manager displays only the following aggregates for creating FlexVol volumes for the LUN:

- In normal mode, when you create volumes on sync-source SVMs or data-serving SVMs in the primary site, only those aggregates that belong to the cluster in the primary site are displayed.
- In switched-over mode, when you create volumes on sync-destination SVMs or data-serving SVMs in the surviving site, only switched-over aggregates are displayed.

Steps

1. Click **Storage > LUNs**.
2. In the **LUN Management** tab, click **Create**.
3. Browse and select an SVM in which you want to create the LUNs.
4. In the **Create LUN Wizard**, specify the name, size, type, description for the LUN, and select the **Space Reserve**, and then click **Next**.
5. Create a new FlexVol volume for the LUN or select an existing volume or qtree, and then click **Next**.
6. Add initiator groups if you want to control host access to the LUN, and then click **Next**.
7. Select the **Manage Storage Quality of Service** check box if you want to manage the workload performance of the LUN.
8. Create a new storage QoS policy group or select an existing policy group to control the input/output (I/O) performance of the LUN:

If you want to...	Do this...
Create a new policy group	<p>a. Select New Policy Group</p> <p>b. Specify the policy group name.</p> <p>c. Specify the minimum throughput limit.</p> <ul style="list-style-type: none"> ◦ In System Manager 9.5, you can set the minimum throughput limit only on a performance-based All Flash Optimized personality. In System Manager 9.6, you can also set the minimum throughput limit for ONTAP Select Premium systems. ◦ You cannot set the minimum throughput limit for volumes on a FabricPool-enabled aggregate. ◦ If you do not specify the minimum throughput value or if the minimum throughput value is set to 0, the system automatically displays “None” as the value. <p style="padding-left: 40px;">This value is case-sensitive.</p> <p>d. Specify the maximum throughput limit to ensure that the workload of the objects in the policy group does not exceed the specified throughput limit.</p> <ul style="list-style-type: none"> ◦ The minimum throughput limit and the maximum throughput limit must be of the same unit type. ◦ If you do not specify the minimum throughput limit, then you can set the maximum throughput limit in IOPS and B/s, KB/s, MB/s, and so on. ◦ If you do not specify the maximum throughput value, the system automatically displays “Unlimited” as the value and this value is case-sensitive. <p style="padding-left: 40px;">The unit that you specify does not affect the maximum throughput.</p>

If you want to...	Do this...
<p>Select an existing policy group</p>	<p>a. Select Existing Policy Group, and then click Choose to select an existing policy group from the Select Policy Group dialog box.</p> <p>b. Specify the minimum throughput limit.</p> <ul style="list-style-type: none"> ◦ In System Manager 9.5, you can set the minimum throughput limit only on a performance-based All Flash Optimized personality. In System Manager 9.6, you can also set the minimum throughput limit for ONTAP Select Premium systems. ◦ You cannot set the minimum throughput limit for volumes on a FabricPool-enabled aggregate. ◦ If you do not specify the minimum throughput value or if the minimum throughput value is set to 0, the system automatically displays “None” as the value. <p style="padding-left: 40px;">This value is case-sensitive.</p> <p>c. Specify the maximum throughput limit to ensure that the workload of the objects in the policy group does not exceed the specified throughput limit.</p> <ul style="list-style-type: none"> ◦ The minimum throughput limit and the maximum throughput limit must be of the same unit type. ◦ If you do not specify the minimum throughput limit, then you can set the maximum throughput limit in IOPS and B/s, KB/s, MB/s, and so on. ◦ If you do not specify the maximum throughput value, the system automatically displays “Unlimited” as the value and this value is case-sensitive. <p style="padding-left: 40px;">The unit that you specify does not affect the maximum throughput.</p> <p>If the policy group is assigned to more than one object, the maximum throughput that you specify is shared among the objects.</p>

9. Review the specified details in the **LUN summary** window, and then click **Next**.

10. Confirm the details, and then click **Finish** to complete the wizard.

Related information

Deleting LUNs

You can use System Manager to delete LUNs and return the space used by the LUNs to their containing aggregates or volumes.

Before you begin

- The LUN must be offline.
- The LUN must be unmapped from all initiator hosts.

Steps

1. Click **Storage > LUNs**.
2. In the **LUN Management** tab, select one or more LUNs that you want to delete, and then click **Delete**.
3. Select the confirmation check box, and then click **Delete**.

Related information

Create initiator groups

You can use System Manager to create an initiator group. Initiator groups enable you to control host access to specific LUNs. You can use port sets to limit which LIFs an initiator can access.

Steps

1. Click **Storage > LUNs**.
2. In the **Initiator Groups** tab, click **Create**.
3. In the **General** tab of the **Create Initiator Group** dialog box, specify the initiator group name, operating system, host alias name, port set, and supported protocol for the group.
4. Click **Create**.

Related information

Deleting initiator groups

You can use the Initiator Groups tab in System Manager to delete initiator groups.

Before you begin

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

Steps

1. Click **Storage > LUNs**.

2. In the **Initiator Groups** tab, select one or more initiator groups that you want to delete, and then click **Delete**.
3. Click **Delete**.
4. Verify that the initiator groups you deleted are no longer displayed in the **Initiator Groups** tab.

Related information

[LUNs window](#)

Add initiators

You can use System Manager to add initiators to an initiator group. An initiator provides access to a LUN when the initiator group that it belongs to is mapped to that LUN.

Steps

1. Click **Storage > LUNs**.
2. In the **LUN Management** tab, select the initiator group to which you want to add initiators and click **Edit**.
3. In the **Edit Initiator Group** dialog box, click **Initiators**.
4. Click **Add**.
5. Specify the initiator name and click **OK**.
6. Click **Save and Close**.

Related information

[LUNs window](#)

Deleting initiators from an initiator group

You can use the Initiator Groups tab in System Manager to delete an initiator. To delete an initiator from an initiator group, you must disassociate the initiator from the initiator group.

Before you begin

All of the LUNs that are mapped to the initiator group that contains the initiator that you want to delete must be manually unmapped.

Steps

1. Click **Storage > LUNs**.
2. In the **Initiator Groups** tab, select the initiator group from which you want to delete the initiator, and then click **Edit**.
3. In the **Edit Initiator Group** dialog box, click the **Initiators** tab.
4. Select and delete the initiator from the text box, and click **Save**.

The initiator is disassociated from the initiator group.

Related information

[LUNs window](#)

Create port sets

You can use System Manager to create port sets to limit access to your LUNs.

Steps

1. Click **Storage > LUNs**.
2. In the **Portsets** tab, click **Create**.
3. In the **Create Portset** dialog box, select the type of protocol.
4. Choose the network interface that you want to associate with the port set.
5. Click **Create**.

Deleting port sets

You can use System Manager to delete a port set when it is no longer required.

Steps

1. Click **Storage > LUNs**.
2. In the **Portsets** tab, select one or more port sets and click **Delete**.
3. Confirm the deletion by clicking **Delete**.

Cloning LUNs

LUN clones enable you to create multiple readable and writable copies of a LUN. You can use System Manager to create a temporary copy of a LUN for testing or to make a copy of your data available to additional users without providing them access to the production data.

Before you begin

- You must have installed the FlexClone license on the storage system.
- When space reservation is disabled on a LUN, the volume that contains the LUN must have enough space to accommodate changes to the clone.

About this task

- When you create a LUN clone, automatic deletion of the LUN clone is enabled by default in System Manager.

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

- You cannot clone LUNs that are on SnapLock volumes.

Steps

1. Click **Storage > LUNs**.
2. In the **LUN Management** tab, select the LUN that you want to clone, and then click **Clone**.
3. If you want to change the default name, specify a new name for the LUN clone.
4. Click **Clone**.
5. Verify that the LUN clone that you created is listed in the **LUNs** window.

Related information

[LUNs window](#)

Editing LUNs

You can use the LUN properties dialog box in System Manager to change the name, description, size, space reservation setting, or the mapped initiator hosts of a LUN.

About this task

When you resize a LUN, you have to perform the steps on the host side that are recommended for the host type and the application that is using the LUN.

Steps

1. Click **Storage > LUNs**.
2. In the **LUN Management** tab, select the LUN that you want to edit from the list of LUNs, and click **Edit**.
3. Make the required changes.
4. Click **Save and Close**.

Related information

[LUNs window](#)

Bringing LUNs online

You can use the **LUN Management** tab in System Manager to bring selected LUNs online and make them available to the host.

Before you begin

Any host application accessing the LUN must be quiesced or synchronized.

Steps

1. Click **Storage > LUNs**.
2. In the **LUN Management** tab, select one or more LUNs that you want to bring online.
3. Click **Status > Online**.

Related information

[LUNs window](#)

Taking LUNs offline

You can use the **LUN Management** tab in System Manager to take selected LUNs offline and make them unavailable for block protocol access.

Before you begin

Any host application accessing the LUN must be quiesced or synchronized.

Steps

1. Click **Storage > LUNs**.
2. In the **LUN Management** tab, select one or more LUNs that you want to take offline.
3. Click **Status > Offline**.

Related information

[LUNs window](#)

Moving LUNs

You can use System Manager to move a LUN from its containing volume to another volume or qtree within a storage virtual machine (SVM). You can move the LUN to a volume that is hosted on an aggregate containing high-performance disks, thereby improving the performance when accessing the LUN.

About this task

- You cannot move a LUN to a qtree within the same volume.
- If you have created a LUN from a file using the command-line interface (CLI), you cannot move the LUN using System Manager.
- The LUN move operation is nondisruptive; it can be performed when the LUN is online and serving data.
- You cannot use System Manager to move the LUN if the allocated space in the destination volume is not sufficient to contain the LUN, and even if autogrow is enabled on the volume.

You should use the CLI instead.

- You cannot move LUNs on SnapLock volumes.

Steps

1. Click **Storage > LUNs**.
2. In the **LUN Management** tab, select the LUN that you want to move from the list of LUNs, and then click **Move**.
3. In the **Move Options** area of the **Move LUN** dialog box, specify a new name for the LUN if you want to change the default name.
4. Select the storage object to which you want to move the LUN and perform one of the following actions:

If you want to move the LUN to...	Then...
A new volume	<ol style="list-style-type: none"> a. Select an aggregate in which you want to create the new volume. b. Specify a name for the volume.
An existing volume or qtree	<ol style="list-style-type: none"> a. Select a volume to which you want to move the LUN. b. If the selected volume contains any qtrees, select the qtree to which you want to move the LUN.

5. Click **Move**.
6. Confirm the LUN move operation, and click **Continue**.

For a brief period of time, the LUN is displayed on both the origin and destination volume. After the move operation is complete, the LUN is displayed on the destination volume.

The destination volume or qtree is displayed as the new container path for the LUN.

Assigning LUNs to storage QoS

You can use System Manager to limit the throughput of LUNs by assigning them to storage Quality of Service (QoS) policy groups. You can assign storage QoS for new LUNs or modify storage QoS details for LUNs that are already assigned to a policy group.

About this task

- You cannot assign storage QoS to a LUN if the following storage objects are assigned to a policy group:
 - Parent volume of the LUN
 - Parent storage virtual machine (SVM) of the LUN
- You can assign storage QoS or modify the QoS details for a maximum of 10 LUNs simultaneously.

Steps

1. Click **Storage > LUNs**.
2. In the **LUN Management** tab, select one or more LUNs for which you want to assign storage QoS.
3. Click **Storage QoS**.
4. In the **Quality of Service Details** dialog box, select the **Manage Storage Quality of Service** check box if you want to manage the workload performance of the LUN.

If some of the LUNs that you selected are already assigned to a policy group, the changes that you make might affect the performance of these LUNs.

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

If you want to...	Do this...
Create a new policy group	<p>a. Select New Policy Group.</p> <p>b. Specify the policy group name.</p> <p>c. Specify the minimum throughput limit.</p> <ul style="list-style-type: none"> ◦ In System Manager 9.5, you can set the minimum throughput limit only on a performance-based All Flash Optimized personality. In System Manager 9.6, you can also set the minimum throughput limit for ONTAP Select Premium systems. ◦ You cannot set the minimum throughput limit for volumes on a FabricPool-enabled aggregate. ◦ If you do not specify the minimum throughput value or if the minimum throughput value is set to 0, the system automatically displays “None” as the value. <p style="padding-left: 40px;">This value is case-sensitive.</p> <p>d. Specify the maximum throughput limit to ensure that the workload of the objects in the policy group does not exceed the specified throughput limit.</p> <ul style="list-style-type: none"> ◦ The minimum throughput limit and the maximum throughput limit must be of the same unit type. ◦ If you do not specify the minimum throughput limit, then you can set the maximum throughput limit in IOPS and B/s, KB/s, MB/s, and so on. ◦ If you do not specify the maximum throughput value, the system automatically displays “Unlimited” as the value and this value is case-sensitive. <p style="padding-left: 40px;">The unit that you specify does not affect the maximum throughput.</p>

If you want to...	Do this...
<p>Select an existing policy group</p>	<ol style="list-style-type: none"> a. Select Existing Policy Group, and then click Choose to select an existing policy group from the Select Policy Group dialog box. b. Specify the minimum throughput limit. <ul style="list-style-type: none"> ◦ In System Manager 9.5, you can set the minimum throughput limit only on a performance-based All Flash Optimized personality. In System Manager 9.6, you can also set the minimum throughput limit for ONTAP Select Premium systems. ◦ You cannot set the minimum throughput limit for volumes on a FabricPool-enabled aggregate. ◦ If you do not specify the minimum throughput value or if the minimum throughput value is set to 0, the system automatically displays “None” as the value. <p style="margin-left: 40px;">This value is case-sensitive.</p> c. Specify the maximum throughput limit to ensure that the workload of the objects in the policy group does not exceed the specified throughput limit. <ul style="list-style-type: none"> ◦ The minimum throughput limit and the maximum throughput limit must be of the same unit type. ◦ If you do not specify the minimum throughput limit, then you can set the maximum throughput limit in IOPS and B/s, KB/s, MB/s, and so on. ◦ If you do not specify the maximum throughput value, the system automatically displays “Unlimited” as the value and this value is case-sensitive. <p style="margin-left: 40px;">The unit that you specify does not affect the maximum throughput.</p> <p style="margin-left: 40px;">If the policy group is assigned to more than one object, the maximum throughput that you specify is shared among the objects.</p>

6. Click the link that specifies the number of LUNs to review the list of selected LUNs, and click **Discard** if you want to remove any LUNs from the list.

The link is displayed only when multiple LUNs are selected.

7. Click **OK**.

Editing initiator groups

You can use the Edit Initiator Group dialog box in System Manager to change the name of an existing initiator group and its operating system. You can add initiators to or remove initiators from the initiator group. You can also change the port set associated with the initiator group.

Steps

1. Click **Storage > LUNs**.
2. In the **Initiator Groups** tab, select the initiator group that you want to modify, and then click **Edit**.
3. Make the necessary changes.
4. Click **Save and Close**.
5. Verify the changes you made to the initiator group in the **Initiator Groups** tab.

Related information

[LUNs window](#)

Editing initiators

You can use the Edit Initiator Group dialog box in System Manager to change the name of an existing initiator in an initiator group.

Steps

1. Click **Storage > LUNs**.
2. In the **Initiator Groups** tab, select the initiator group to which the initiator belongs, and then click **Edit**.
3. In the **Edit Initiator Group** dialog box, click **Initiators**.
4. Select the initiator that you want to edit and click **Edit**.
5. Change the name and click **OK**.
6. Click **Save and Close**.

Related information

[LUNs window](#)

Editing port sets

You can use the Portsets tab in System Manager to edit settings related to port sets.

Steps

1. Click **Storage > LUNs**.
2. In the **Portsets** tab, select the port set you want to edit and click **Edit**.
3. In the **Edit Portset** dialog box, make the necessary changes.
4. Click **Save and Close**.

Related information

[Configuring iSCSI protocol on SVMs](#)

Viewing LUN information

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

Steps

1. Click **Storage > LUNs**.
2. In the **LUN Management** tab, select the LUN that you want to view information about from the displayed list of LUNs.
3. Review the LUN details in the **LUNs** window.

Viewing initiator groups

You can use the Initiator Groups tab in System Manager to view all the initiator groups and the initiators mapped to these initiator groups, and the LUNs and LUN ID mapped to the initiator groups.

Steps

1. Click **Storage > LUNs**.
2. Click **Initiator Groups** and review the initiator groups that are listed in the upper pane.
3. Select an initiator group to view the initiators that belong to it, which are listed in the **Initiators** tab in the lower pane.
4. Select an initiator group to view the LUNs mapped to it, which are listed in the **Mapped LUNs** in the lower pane.

Guidelines for working with FlexVol volumes that contain LUNs

When you work with FlexVol volumes that contain LUNs, you must change the default settings for Snapshot copies. You can also optimize the LUN layout to simplify administration.

Snapshot copies are required for many optional features such as SnapMirror, SyncMirror, dump and restore, and ndmcopy.

When you create a volume, ONTAP automatically performs the following:

- Reserves 5 percent of the space for Snapshot copies
- Schedules Snapshot copies

Because the internal scheduling mechanism for creating Snapshot copies within ONTAP does not ensure that the data within a LUN is in a consistent state, you should change these Snapshot copy settings by performing the following tasks:

- Turn off the automatic Snapshot copy schedule.

- Delete all of the existing Snapshot copies.
- Set the percentage of space reserved for Snapshot copies to zero.

You should use the following guidelines to create volumes that contain LUNs:

- Do not create any LUNs in the system's root volume.

ONTAP uses this volume to administer the storage system. The default root volume is /vol/vol0.

- You should use a SAN volume to contain the LUN.
- You should ensure that no other files or directories exist in the volume that contains the LUN.

If this is not possible and you are storing LUNs and files in the same volume, you should use a separate qtree to contain the LUNs.

- If multiple hosts share the same volume, you should create a qtree on the volume to store all of the LUNs for the same host.

This is a best practice that simplifies LUN administration and tracking.

- To simplify management, you should use naming conventions for LUNs and volumes that reflect their ownership or the way that they are used.

Related information

[ONTAP 9 Documentation Center](#)

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.



Not all ONTAP versions support all LUN multiprotocol types. For the latest information, see the Interoperability Matrix Tool.

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

LUN multiprotocol type	When to use
AIX	If your host operating system is AIX.
HP-UX	If your host operating system is HP-UX.
Hyper-V	<p>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 <code>hyper_v</code> for your LUN type, you should also use <code>hyper_v</code> for your igroup OS type.</p> <p> For raw LUNs, you can use the type of child operating system that the LUN multiprotocol type uses.</p>
Linux	If your host operating system is Linux.
NetWare	If your host operating system is NetWare.
OpenVMS	If your host operating system is OpenVMS.
Solaris	If your host operating system is Solaris and you are not using Solaris EFI labels.
Solaris EFI	<p>If you are using Solaris EFI labels.</p> <p> Using any other LUN multiprotocol type with Solaris EFI labels might result in LUN misalignment problems.</p>

LUN multiprotocol type	When to use
VMware	<p>If you are using an ESX Server and your LUNs will be configured with VMFS.</p> <p> If you configure the LUNs with RDM, you can use the guest operating system as the LUN multiprotocol type.</p>
Windows 2003 MBR	If your host operating system is Windows Server 2003 using the MBR partitioning method.
Windows 2003 GPT	If you want to use the GPT partitioning method and your host is capable of using it. Windows Server 2003, Service Pack 1 and later are capable of using the GPT partitioning method, and all 64-bit versions of Windows support it.
Windows 2008 or later	If your host operating system is Windows Server 2008 or later; both MBR and GPT partitioning methods are supported.
Xen	<p>If you are using Xen and your LUNs will be configured with Linux LVM with Dom0.</p> <p> For raw LUNs, you can use the type of guest operating system that the LUN multiprotocol type uses.</p>

Related information

[Creating LUNs](#)

[NetApp Interoperability](#)

[Solaris Host Utilities 6.1 Installation and Setup Guide](#)

[Solaris Host Utilities 6.1 Quick Command Reference](#)

[Solaris Host Utilities 6.1 Release Notes](#)

Understanding LUN clones

LUN clones are writable, space-efficient clones of parent LUNs. Creating LUN clones is highly space-efficient and time-efficient because the cloning operation does not involve physically copying any data. Clones help in space storage utilization of the physical aggregate space.

You can clone a complete LUN without the need of a backing Snapshot copy in a SAN environment. The cloning operation is instantaneous and clients that are accessing the parent LUN do not experience any

disruption or outage. Clients can perform all normal LUN operations on both parent entities and clone entities. Clients have immediate read/write access to both the parent and cloned LUN.

Clones share the data blocks of their parent LUNs and occupy negligible storage space until clients write new data either to the parent LUN, or to the clone. By default, the LUN clone inherits the space reserved attribute of the parent LUN. For example, if space reservation is disabled on the parent LUN, then space reservation is also disabled on the LUN clone.



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.



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`, `hpux`, `aix`, `netware`, `xen`, `hyper_v`, `vmware`, and `linux`.

You must select an ostyle for the igroup.

LUNs window

You can use the LUNs window to create and manage LUNs and to display information about LUNs. You can also add, edit, or delete initiator groups and initiator IDs.

LUN Management tab

This tab enables you to create, clone, delete, move, or edit the settings of LUNs. You can also assign LUNs to a Storage Quality of Service (QoS) policy group.

Command buttons

- **Create**

Opens the Create LUN wizard, which enables you to create LUNs.

In a cluster on an AFF platform that does not contain any existing LUNs, the Create FC SAN optimized LUNs dialog box is opened, which enables you to set up one or more FC SAN optimized LUNs.

- **Clone**

Opens the Clone LUN dialog box, which enables you to clone the selected LUNs.

- **Edit**

Opens the Edit LUN dialog box, which enables you to edit the settings of the selected LUN.

- **Delete**

Deletes the selected LUN.

- **Status**

Enables you to change the status of the selected LUN to either Online or Offline.

- **Move**

Opens the Move LUN dialog box, which enables you to move the selected LUN to a new volume or an existing volume or qtree within the same storage virtual machine (SVM).

- **Storage QoS**

Opens the Quality of Service details dialog box, which enables you to assign one or more LUNs to a new or existing policy group.

- **Refresh**

Updates the information in the window.

LUNs list

- **Name**

Displays the name of the LUN.

- **SVM**

Displays the name of the storage virtual machine (SVM) in which the LUN is created.

- **Container Path**

Displays the name of the file system (volume or qtree) that contains the LUN.

- **Space Reservation**

Specifies whether space reservation is enabled or disabled.

- **Available Size**

Displays the space available in the LUN.

- **Total Size**

Displays the total space in the LUN.

- **%Used**

Displays the total space (in percentage) that is used.

- **Type**

Specifies the LUN type.

- **Status**

Specifies the status of the LUN.

- **Policy Group**

Displays the name of the Storage QoS policy group to which the LUN is assigned. By default, this column is hidden.

- **Application**

Displays the name of the application that is assigned to the LUN.

- **Description**

Displays the description of the LUN.

Details area

The area below the LUNs list displays details related to the selected LUN.

- **Details tab**

Displays details related to the LUN such as the LUN serial number, whether the LUN is a clone, LUN description, the policy group to which the LUN is assigned, minimum throughput of the policy group, maximum throughput of the policy group, details about the LUN move operation, and the application assigned to the LUN. You can also view details about the initiator groups and initiators that are associated with the selected LUN.

- **Performance tab**

Displays performance metrics graphs of the LUNs, including data rate, IOPS, and response time.

Changing the client time zone or the cluster time zone impacts the performance metrics graphs. Refresh your browser to see the updated graphs.

Initiator Groups tab

This tab enables you to create, delete, or edit the settings of initiator groups and initiator IDs.

Command buttons

- **Create**

Opens the Create Initiator Group dialog box, which enables you to create initiator groups to control host access to specific LUNs.

- **Edit**

Opens the Edit Initiator Group dialog box, which enables you to edit the settings of the selected initiator group.

- **Delete**

Deletes the selected initiator group.

- **Refresh**

Updates the information in the window.

Initiator Groups list

- **Name**

Displays the name of the initiator group.

- **Type**

Specifies the type of protocol supported by the initiator group. The supported protocols are iSCSI, FC/FCoE, or Mixed (iSCSI and FC/FCoE).

- **Operating System**

Specifies the operating system for the initiator group.

- **Portset**

Displays the port set that is associated with the initiator group.

- **Initiator Count**

Displays the number of initiators added to the initiator group.

Details area

The area below the Initiator Groups list displays details about the initiators that are added to the selected initiator group and the LUNs that are mapped to the initiator group.

Portsets tab

This tab enables you to create, delete, or edit the settings of port sets.

Command buttons

- **Create**

Opens the Create Portset dialog box, which enables you to create port sets to limit access to your LUNs.

- **Edit**

Opens the Edit Portset dialog box, which enables you to select the network interfaces that you want to associate with the port set.

- **Delete**

Deletes the selected port set.

- **Refresh**

Updates the information in the window.

Portsets list

- **Portset Name**

Displays the name of the port set.

- **Type**

Specifies the type of protocol supported by the port set. The supported protocols are iSCSI, FC/FCoE, or Mixed (iSCSI and FC/FCoE).

- **Interface Count**

Displays the number of network interfaces that are associated with the port set.

- **Initiator Group Count**

Displays the number of initiator groups that are associated with the port set.

Details area

The area below the Portsets list displays details about the network interfaces and initiator groups associated with the selected port set.

Related information

[Creating LUNs](#)

[Deleting LUNs](#)

[Creating initiator groups](#)

[Deleting initiator groups](#)

[Adding initiators](#)

[Deleting initiators from an initiator group](#)

[Editing LUNs](#)

[Editing initiator groups](#)

[Editing initiators](#)

[Bringing LUNs online](#)

[Taking LUNs offline](#)

[Cloning LUNs](#)

Qtrees

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

Related information

[ONTAP concepts](#)

[Logical storage management](#)

[NFS management](#)

[SMB/CIFS management](#)

Create qtrees

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

Steps

1. Click **Storage > Qtrees**.
2. From the drop-down menu in the **SVM** field, select the storage virtual machine (SVM) on which you want to create a qtree.
3. Click **Create**.
4. In the **Details** tab of the **Create Qtree** dialog box, type a name for the qtree.
5. Select the volume to which you want to add the qtree.

The Volume browse list includes only the volumes that are online.

6. If you want to disable opportunistic locks (oplocks) for the qtree, clear the **Enable Oplocks for files and directories in this Qtree** check box.

By default, oplocks are enabled for each qtree.

7. If you want to change the default inherited security style, select a new security style.

The default security style of the qtree is the security style of the volume that contains the qtree.

8. If you want to change the default inherited export policy, either select an existing export policy or create an export policy.

The default export policy of the qtree is the export policy that is assigned to the volume that contains the qtree.

9. If you want to restrict the disk space usage, click the **Quotas** tab.

- a. If you want to apply quotas on the qtree, click **Qtree quota**, and then specify the disk space limit.
- b. If you want to apply quotas for all the users on the qtree, click **User quota**, and then specify the disk space limit.

10. Click **Create**.

11. Verify that the qtree that you created is included in the list of qtrees in the **Qtrees** window.

Related information

[Qtrees window](#)

Deleting qtrees

You can delete a qtree and reclaim the disk space that the qtree uses within a volume by using System Manager. When you delete a qtree, all of the quotas that are applicable to that qtree are no longer applied by ONTAP.

Before you begin

- The qtree status must be normal.
- The qtree must not contain any LUN.

Steps

1. Click **Storage > Qtrees**.
2. In the **Qtrees** window, select one or more qtrees that you want to delete, and then click **Delete**.
3. Select the confirmation check box, and then click **Delete**.
4. Verify that the qtree that you deleted is no longer included in the list of qtrees in the **Qtrees** window.

Related information

[Qtrees window](#)

Editing qtrees

You can use System Manager to modify the properties of a qtree such as the security

style, enable or disable opportunistic locks (oplocks), and assign a new or existing export policy.

Steps

1. Click **Storage > Qtrees**.
2. Select the qtree that you want to edit, and then click **Edit**.
3. In the **Edit Qtree** dialog box, edit the following properties as required:
 - Oplocks
 - Security style
 - Export policy
4. Click **Save**.
5. Verify the changes that you made to the selected qtree in the **Qtrees** window.

Related information

[Qtrees window](#)

Assigning export policies to qtrees

Instead of exporting an entire volume, you can export a specific qtree on a volume to make it directly accessible to clients. You can use System Manager to export a qtree by assigning an export policy to the qtree. You can assign an export policy to one or more qtrees from the Qtrees window.

Steps

1. Click **Storage > Qtrees**.
2. From the drop-down menu in the SVM field, select the storage virtual machine (SVM) on which the qtrees that you want to export reside.
3. Select one or more qtrees for which you want to assign an export policy, and then click **Change Export Policy**.
4. In the **Export Policy** dialog box, either create an export policy or select an existing export policy.

[Creating an export policy](#)

5. Click **Save**.
6. Verify that the export policy and its related export rules that you assigned to the qtrees are displayed in the **Details** tab of the appropriate qtrees.

Viewing qtree information

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

Steps

1. Click **Storage > Qtrees**.
2. From the drop-down menu in the **SVM** field, select the storage virtual machine (SVM) on which the qtree

about which you want to view information resides.

3. Select the qtree from the displayed list of qtrees.
4. Review the qtree details in the **Qtrees** window.

Qtree options

A qtree is a logically defined file system that can exist as a special subdirectory of the root directory within a FlexVol volume. 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

Qtrees window

You can use the Qtrees window to create, display, and manage information about qtrees.

Command buttons

- **Create**

Opens the Create Qtree dialog box, which enables you to create a new qtree.

- **Edit**

Opens the Edit Qtree dialog box, which enables you to change the security style and to enable or disable oplocks (opportunistic locks) on a qtree.

- **Change Export Policy**

Opens the Export Policy dialog box, which enables you to assign one or more qtrees to new or existing

export policies.

- **Delete**

Deletes the selected qtree.

This button is disabled unless the status of the selected qtree is normal.

- **Refresh**

Updates the information in the window.

Qtree list

The qtree list displays the volume in which the qtree resides and the qtree name.

- **Name**

Displays the name of the qtree.

- **Volume**

Displays the name of the volume in which the qtree resides.

- **Security Style**

Specifies the security style of the qtree.

- **Status**

Specifies the current status of the qtree.

- **Oplocks**

Specifies whether the oplocks setting is enabled or disabled for the qtree.

- **Export Policy**

Displays the name of the export policy to which the qtree is assigned.

Details area

- **Details tab**

Displays detailed information about the selected qtree, such as the mount path of the volume containing the qtree, details about the export policy, and the export policy rules.

Related information

[Creating qtrees](#)

[Deleting qtrees](#)

[Editing qtrees](#)

Quotas

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

Related information

[Logical storage management](#)

Create quotas

Quotas enable you to restrict or track the disk space and number of files that are used by a user, group, or qtree. You can use the Add Quota wizard in System Manager to create a quota and to apply the quota to a specific volume or qtree.

About this task

Using System Manager, the minimum value that you can specify for the hard limit and soft limit on the number of files that the quota can own is 1000. If you want to specify a value lower than 1000, you should use the command-line interface (CLI).

Steps

1. Click **Storage > Quotas**.
2. From the drop-down menu in the **SVM** field, select the storage virtual machine (SVM) on which you want to create a quota.
3. In the **User Defined Quotas** tab, click **Create**.

The Create Quota Wizard is displayed.

4. Type or select information as prompted by the wizard.
5. Confirm the details, and then click **Finish** to complete the wizard.

What to do next

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/passwdfile` and the `/etc/groupfile` must be updated, respectively.

Related information

[Quotas window](#)

Deleting quotas

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

Steps

1. Click **Storage > Quotas**.
2. From the drop-down menu in the **SVM** field, select the storage virtual machine (SVM) on which the quotas that you want to delete reside.

3. Select one or more quotas that you want to delete, and then click **Delete**.
4. Select the confirmation check box, and then click **Delete**.

Related information

[Quotas window](#)

Editing quota limits

You can use System Manager to edit the disk space threshold, the hard limit and soft limit on the amount of disk space that the quota target can use, and the hard limit and soft limit on the number of files that the quota target can own.

Steps

1. Click **Storage > Quotas**.
2. From the drop-down menu in the **SVM** field, select the storage virtual machine (SVM) on which the quota that you want to edit resides.
3. Select the quota that you want to edit, and click **Edit Limits**.
4. In the **Edit Limits** dialog box, edit the quota settings as required.

One hundred (100) is the minimum value that you can specify for the hard limit and soft limit on the number of files that the quota can own. If you want to specify a value lower than 100, you should use the command-line interface (CLI).

5. Click **Save and Close**.
6. Verify the changes that you made to the selected quota in the **User Defined Quotas** tab.

Related information

[Quotas window](#)

Activating or deactivating quotas

You can use System Manager to activate or deactivate quotas on one or more volumes that you select on your storage system. You can activate or deactivate quotas when you users and their storage requirements and limitations change.

Steps

1. Click **Storage > Quotas**.
2. From the drop-down menu in the SVM field, select the storage virtual machine (SVM) on which the quotas that you want to activate or deactivate reside.
3. In the **Quota Status on Volumes** tab, select one or more volumes for which you want to activate or deactivate quotas.
4. Click **Activate** or **Deactivate**, as required.
5. If you are deactivating a quota, select the confirmation check box, and then click **OK**.
6. Verify the quota status on the volumes in the **Status** column.

Related information

Resizing quotas

You can use the Resize Quota dialog box in System Manager to adjust the active quotas in the specified volume so that they reflect the changes that you have made to a quota.

Before you begin

Quotas must be enabled for the volumes for which you want to resize quotas.

Steps

1. Click **Storage > Quotas**.
2. In the **Quota Status on Volumes** tab of the **Quotas** window, select one or more volumes for which you want to resize the quotas.
3. Click **Resize**.

Related information

Viewing quota information

You can use the Quotas window in System Manager to view quota details such as the volume and qtrees to which the quota is applied, the type of quota, the user or group to which the quota is applied, and the space and file usage.

Steps

1. Click **Storage > Quotas**.
2. From the drop-down menu in the **SVM** field, select the storage virtual machine (SVM) on which the quota that you want to view information about resides.
3. Perform the appropriate action:

If...	Then...
You want to view details of all of the quotas that you created	Click the User Defined Quotas tab.
You want to view details of the quotas that are currently active	Click the Quota Report tab.

4. Select the quota that you want to view information about from the displayed list of quotas.
5. Review the quota details.

Types of quotas

Quotas can be classified on the basis of the targets to which they are applied.

The following are the types of quotas based on the targets to which they are applied:

- **User quota**

The target is a user.

The user can be represented by a UNIX user name, UNIX UID, a Windows SID, a file or directory whose UID matches the user, Windows user name in pre-Windows 2000 format, and a file or directory with an ACL owned by the user's SID. You can apply it to a volume or a qtree.

- **Group quota**

The target is a group.

The group is represented by a UNIX group name, a GID, or a file or directory whose GID matches the group. ONTAP does not apply group quotas based on a Windows ID. You can apply a quota to a volume or a qtree.

- **Qtree quota**

The target is a qtree, specified by the path name to the qtree.

You can determine the size of the target qtree.

- **Default quota**

Automatically applies a quota limit to a large set of quota targets without creating separate quotas for each target.

Default quotas can be applied to all three types of quota target (users, groups, and qtrees). The quota type is determined by the value of the type field.

Quota limits

You can apply a disk space limit or limit the number of files for each quota type. If you do not specify a limit for a quota, none is applied.

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

System Manager includes several features that help you to create, edit, or delete quotas. You can create a user, group, or tree quota and you can specify quota limits at the disk and file levels. All quotas are established on a per-volume basis.

After creating a quota, you can perform the following tasks:

- Enable and disable quotas
- Resize quotas

Quotas window

You can use the Quotas window to create, display, and manage information about quotas.

Tabs

- **User Defined Quotas**

You can use the **User Defined Quotas** tab to view details of the quotas that you create and to create, edit, or delete quotas.

- **Quota Report**

You can use the Quota Report tab to view the space and file usage and to edit the space and file limits of quotas that are active.

- **Quota Status on Volumes**

You can use the Quota Status on Volumes tab to view the status of a quota and to turn quotas on or off and to resize quotas.

Command buttons

- **Create**

Opens the Create Quota wizard, which enables you to create quotas.

- **Edit Limits**

Opens the Edit Limits dialog box, which enables you to edit settings of the selected quota.

- **Delete**

Deletes the selected quota from the quotas list.

- **Refresh**

Updates the information in the window.

User Defined Quotas list

The quotas list displays the name and storage information for each quota.

- **Volume**

Specifies the volume to which the quota is applied.

- **Qtree**

Specifies the qtree associated with the quota. "All Qtrees" indicates that the quota is associated with all the qtrees.

- **Type**

Specifies the quota type: user, or group, or tree.

- **User/Group**

Specifies a user or a group associated with the quota. "All Users" indicates that the quota is associated with all the users. "All groups" indicates that the quota is associated with all the groups.

- **Quota Target**

Specifies the type of target that the quota is assigned to. The target can be qtree, user, or group.

- **Space Hard Limit**

Specifies the disk space limit applied to hard quotas.

This field is hidden by default.

- **Space Soft Limit**

Specifies the disk space limit applied to soft quotas.

This field is hidden by default.

- **Threshold**

Specifies the disk space limit applied to threshold quotas.

This field is hidden by default.

- **File Hard Limit**

Specifies the maximum number of files in a hard quota.

This field is hidden by default.

- **File Soft Limit**

Specifies the maximum number of files in a soft quota.

This field is hidden by default.

Details area

The area below the quotas list displays quota details such as the quota error, space usage and limits, and file usage and limits.

Related information

[Creating quotas](#)

[Deleting quotas](#)

[Editing quota limits](#)

[Activating or deactivating quotas](#)

[Resizing quotas](#)

CIFS protocol

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

Related information

[SMB/CIFS management](#)

Setting up CIFS

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

Before you begin

- The CIFS license must be installed on your storage system.
- While configuring CIFS in the Active Directory domain, the following requirements must be met:
 - DNS must be enabled and configured correctly.
 - The storage system must be able to communicate with the domain controller by using the fully qualified domain name (FQDN).
 - The time difference (clock skew) between the cluster and the domain controller must not be more than five minutes.
- If CIFS is the only protocol that is configured on the storage virtual machine (SVM), the following requirements must be met:
 - The root volume security style must be NTFS.

By default, System Manager sets the security style as UNIX.

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

Steps

1. Click **Storage > SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **Configuration** tab, click **Set up**.
4. In the **General** tab of the **CIFS Server Setup** dialog box, specify the NetBIOS name and the Active Directory domain details.
5. Click the **Options** tab, and then perform the following actions:
 - In the SMB settings area, select or clear the SMB signing check box and the SMB encryption check box, as required.
 - Specify the default UNIX user.
 - In the WINS Servers area, add the required IP address.
6. Click **Set up**.

Related information

[Creating a CIFS share](#)

[CIFS window](#)

[Editing volume properties](#)

[Modifying export policy rules](#)

Editing the general properties for CIFS

You can modify the general properties for CIFS such as the default UNIX user and default Windows user by using System Manager. You can also enable or disable SMB signing for the CIFS server.

Steps

1. Click **Storage > SVMs**.

2. Select the SVM, and then click **SVM Settings**.
3. In the **Configuration** tab, click **Options**.
4. In the **CIFS Options** dialog box, modify the following CIFS server properties, as required:
 - UNIX user
 - Windows user
 - IP address
 - Enable or disable SMB signing

Enabling SMB signing prevents the data from being compromised. However, you might encounter performance degradation in the form of increased CPU usage on both the clients and the server, although the network traffic remains the same. You can disable SMB signing on any of your Windows clients that do not require protection against replay attacks.

For information about disabling SMB signing on Windows clients, see the Microsoft Windows documentation.

- Enable or disable SMB 3.0 encryption

You should enable SMB Multichannel to establish multiple channels between an SMB 3.0 session and transport connections.

5. Click either **Save** or **Save and Close**.

Related information

[CIFS window](#)

Add home directory paths

You can use System Manager to specify one or more paths that can be used by the storage system to resolve the location of the CIFS home directories of users.

Steps

1. Click **Storage > SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **Home Directories** area of the **Configuration** tab, click **Manage**.
4. In the **Manage Home Directories** dialog box, specify the paths that are to be used by the storage system to search for the CIFS home directories of users.
5. Click **Add**, and then click **Save and Close**.

Related information

[CIFS window](#)

Deleting home directory paths

You can use System Manager to delete a home directory path when you do not want the storage system to use the path to resolve the location of the CIFS home directories of users.

Steps

1. Click **Storage > SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **Home Directories** area of the **Configuration** tab, click **Manage**.
4. In the **Manage Home Directories** dialog box, select the home directory path that you want to delete, and then click **Delete**.
5. Click **Save and Close**.

Related information

[CIFS window](#)

Resetting CIFS domain controllers

You can use System Manager to reset the CIFS connection to domain controllers for the specified domain. Failure to reset the domain controller information can cause a connection failure.

About this task

You have to update the discovery information of the storage system's available domain controller after you add or delete a domain from the list of preferred domain controllers. You can update the storage system's available domain controller discovery information in ONTAP through the command-line interface (CLI).

Steps

1. Click **Storage > SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **Domain** tab, click **Reset**.

Related information

[CIFS window](#)

Updating the CIFS group policy configuration

You have to update the group policy after the policy configuration is changed through the command-line interface (CLI). You can use the CIFS window in System Manager to update the group policy.

Steps

1. Click **Storage > SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. Click the **Domain** tab.
4. In the **Group Policy** area, select the group policy configuration that you want to update, and then click **Update**.

Enabling or disabling a CIFS group policy configuration

You can enable or disable the CIFS group policy configuration from the CIFS window in System Manager.

Steps

1. Click **Storage > SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. Click the **Domain** tab.
4. In the **Group Policy** area, select the group policy configuration that you want to enable or disable, and then click **Enable** or **Disable**, as required.

Reloading CIFS group policy

You have to reload a CIFS group policy if the status of the policy is changed. You can use the CIFS window in System Manager to reload the group policy.

Steps

1. Click **Storage > SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. Click the **Domain** tab.
4. In the **Group Policy** area, select the group policy configuration that you want to reload, and then click **Reload**.

Setting up BranchCache

You can use System Manager to configure BranchCache on a CIFS-enabled storage virtual machine (SVM) to enable the caching of content on computers that are local to the requesting clients.

Before you begin

- CIFS must be licensed, and a CIFS server must be configured.
- For BranchCache version 1, SMB 2.1 or later must be enabled.
- For BranchCache version 2, SMB 3.0 must be enabled, and the remote Windows clients must support BranchCache 2.

About this task

- You can configure BranchCache on SVMs.
- You can create an all-shares BranchCache configuration if you want to offer caching services for all of the content that is contained within all of the SMB shares on the CIFS server.
- You can create a per-share BranchCache configuration if you want to offer caching services for the content that is contained within selected SMB shares on the CIFS server.

Steps

1. Click **Storage > SVMs**.
2. Select the SVM, and then click **SVM Settings**.

3. In the **BranchCache** tab, click **Set Up**.
4. In the **BranchCache Setup** dialog box, enter the following information:
 - a. Specify the path to the hash store.

The path can be to an existing directory where you want the hash data to be stored. The destination path must be read-writable. Read-only paths such as Snapshot directories are not allowed.

- b. Specify the maximum size (in KB, MB, GB, TB, or PB) for a hash data store.

If the hash data exceeds this value, older hashes are deleted to provide space for newer hashes. The default size for a hash store is 1 GB.

- c. Specify the operating mode for the BranchCache configuration.

The default operating mode is set to all shares.

- d. Specify a server key to prevent clients from impersonating the BranchCache server.

You can set the server key to a specific value so that if multiple servers are providing BranchCache data for the same files, clients can use hashes from any server using that same server key. If the server key contains any spaces, you must enclose the server key in quotation marks.

- e. Select the required BranchCache version.

By default, all of the versions that are supported by the client are selected.

5. Click **Set Up**.

Modifying the BranchCache settings

You can use the CIFS window in System Manager to modify the BranchCache settings that are configured for a CIFS-enabled storage virtual machine (SVM). You can change the hash store path, the hash store size, the operating mode, and the BranchCache versions that are supported.

Steps

1. Click **Storage > SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **BranchCache** tab, click **Edit**.
4. In the **Modify BranchCache Settings** dialog box, modify the required information:

- Hash store path

If you modify the hash store path, you are provided with an option to retain the cached hash data from the previous hash store.

- Hash store size
- Operating mode
- BranchCache version

5. Click **Modify**.

Deleting the BranchCache configuration

You can use System Manager to delete the BranchCache configuration if you no longer want to offer caching services on the storage virtual machine (SVM) that is configured for BranchCache.

Steps

1. Click **Storage > SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **BranchCache** tab, click **Delete**.
4. Select the confirmation check box, and then click **Delete**.

You can also remove existing hashes from the hash store.

Add preferred domain controllers

System Manager automatically discovers domain controllers through DNS. Optionally, you can add one or more domain controllers to the list of preferred domain controllers for a specific domain.

Steps

1. Click **Storage > SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **Domain** tab, click **Add** in the **Preferred Domain Controllers** area.
4. Enter the fully qualified domain name (FQDN) and the IP addresses of the domain controllers that you want to add.

You can add multiple domain controllers by entering the IP addresses of the domain controllers, separated by commas.

5. Click **Save**.
6. Verify that the domain controller that you added is displayed in the list of preferred domain controllers.

Editing preferred domain controllers

You can use System Manager to modify the IP address of the preferred domain controllers that are configured for a specific domain.

Steps

1. Click **Storage > SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **Preferred Domain Controllers** area of the **Domain** tab, double-click the domain controller that you want to edit.
4. Modify the IP addresses of the domain controller, and then click **Save**.

Deleting preferred domain controllers

You can use System Manager to delete a preferred domain controller to which the storage virtual machine (SVM) computer account is associated. You can do this when you no longer want to use a particular domain controller.

Steps

1. Click **Storage > SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **Domain** tab, select the domain that you want to delete from the **Preferred Domain Controllers** area, and then click **Delete**.
4. Select the confirmation check box, and then click **Delete**.

Viewing CIFS domain information

You can use System Manager to view information about the domain controllers and servers that are connected to the storage system.

Steps

1. Click **Storage > SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. Click the **Domain** tab.
4. Review the information about the connected domain controllers and servers.

CIFS window

You can use the CIFS window to configure the CIFS server, to manage domain controllers, to manage symbolic UNIX mappings, and to configure BranchCache.

Configuration tab

The Configuration tab enables you to create and manage the CIFS server.

- **Server**

Specifies the status of the CIFS server, name of the server, authentication mode, name of the active directory domain, and status of SMB multichannel.

- **Home Directories**

Specifies home directory paths and the style for determining how PC user names are mapped to home directory entries.

- **Command buttons**

- **Setup**

Opens the CIFS Setup wizard, which enables you to set up CIFS on your storage virtual machine (SVM).

- **Options**

Displays the CIFS Options dialog box, which enables you to enable or disable SMB 3.0 signing, to enable or disable SMB 3.0 encryption, and to add Windows Internet Name Service (WINS) servers.

SMB signing prevents the network traffic between the CIFS server and the client from being compromised.

- **Delete**

Enables you to delete the CIFS server.

- **Refresh**

Updates the information in the window.

Domain tab

The Domain tab enables you to view and reset your CIFS domain controllers, and to add or delete preferred domain controllers. You can also use this tab to manage CIFS group policy configurations.

- **Servers**

Displays information about discovered authentication servers and your preferred domain controllers on the CIFS-enabled SVM.

You can also reset the information about the discovered servers, add a preferred domain controller, delete a domain controller, or refresh the list of domain controllers.

- **Group Policy**

Enables you to view, enable, or disable group policy configurations on the CIFS server. You can also reload a group policy if the status of the policy is changed.

Symlinks tab

The Symlinks tab enables you to manage the mappings of UNIX symbolic links for CIFS users.

- **Path Mappings**

Displays the list of symbolic link mappings for CIFS.

- **Command buttons**

- **Create**

Opens the Create New Symlink Path Mappings dialog box, which enables you to create a UNIX symbolic link mapping.

- **Edit**

Opens the Edit Symlink Path Mappings dialog box, which enables you to modify the CIFS share and path.

- **Delete**

Enables you to delete the symbolic link mapping.

- Refresh

Updates the information in the window.

BranchCache tab

The BranchCache tab enables you to set up and manage BranchCache settings on CIFS-enabled SVMs.

You can view the status of the BranchCache service, the path to the hash store, the size of the hash store, and the operating mode, server key, and version of BranchCache.

• Command buttons

- Setup

Opens the BranchCache Setup dialog box, which enables you to configure BranchCache for the CIFS server.

- Edit

Opens the Modify BranchCache Settings dialog box, which enables you to modify the properties of the BranchCache configuration.

- Delete

Enables you to delete the BranchCache configuration.

- Refresh

Updates the information in the window.

Related information

[Setting up CIFS](#)

[Editing the general properties for CIFS](#)

[Adding home directory paths](#)

[Deleting home directory paths](#)

[Resetting CIFS domain controllers](#)

NFS protocol

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

Related information

[NFS management](#)

Editing NFS settings

You can use System Manager to edit the NFS settings such as enabling or disabling NFSv3, NFSv4, and NFSv4.1, enabling or disabling read and write delegations for NFSv4 clients, and enabling NFSv4 ACLs. You can also edit the default Windows user.

Steps

1. Click **Storage > SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **Protocols** pane, click **NFS**.
4. In the **NFS** window, click **Edit**.
5. In the **Edit NFS Settings** dialog box, make the required changes.
6. Click **Save and Close**.

Related information

[NFS window](#)

NFS window

You can use the NFS window to display and configure your NFS settings.

- **Server Status**

Displays the status of the NFS service. The service is enabled if the NFS protocol is configured on the storage virtual machine (SVM).



If you have upgraded to ONTAP 8.3 or later from an NFS-enabled storage system running Data ONTAP 8.1.x, the NFS service is enabled in ONTAP 8.3 or later. However, you must enable support for NFSv3 or NFSv4 because NFSv2 is no longer supported.

Command buttons

- **Enable**

Enables the NFS service.

- **Disable**

Disables the NFS service.

- **Edit**

Opens the Edit NFS Settings dialog box, which enables you to edit NFS settings.

- **Refresh**

Updates the information in the window.

Related information

NVMe protocol


You can use System Manager to configure the NVMe protocol. The NVMe is a transport protocol that provides high speed access to flash-based network storage. Systems that use NVMe protocol have a subsystem consisting of specific NVME controllers, namespaces, nonvolatile storage medium, hosts, ports and interface between the controller and storage medium.

Setting up NVMe

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

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 System Manager.

The following illustration shows the workflow for setting up NVMe:



Create an NVMe namespace

You can use System Manager to create one or more NVMe namespaces and connect each to a host or set of hosts in a storage virtual machine (SVM). The NVMe namespace

is a quantity of memory that can be formatted into logical blocks. Each namespace can be mapped to an NVMe subsystem.

Before you begin

The SVM must already be configured with the NVMe protocol. To map a namespace, at least one LIF with the data protocol NVMe must exist in the node that owns the namespace.

Steps

1. Click **Storage > NVMe > NVMe namespaces**.
2. Select the SVM that will contain the namespace.
3. Ensure that at least one NVMe LIF is configured for each node of the HA pair. You can create a maximum of two NVMe LIFs per node.
4. Configure the size of the namespace (between 1MB and 16TB).
5. Enter the block size.

For System Manager 9.5, the block size defaults to 4 KB, and this field is not shown.

For System Manager 9.6, you can specify a block size of 4 KB or 512 Bytes.

6. Select the existing volume or create a new volume by choosing the aggregate.

Click on the + symbol to set up additional namespaces (max 250) within the SVM.

7. Select the NVMe subsystem that will be associated with this namespace.

You can choose from the following options:

- None: No subsystems are mapped.
- Use an existing subsystem: The subsystems listed are based on the selected SVM.
- Create a new subsystem: You can choose to create a new subsystem and map to all the new namespaces.

8. Select the host operating system.
9. Click **Submit**.

Related information

[NVMe namespaces window](#)

Editing an NVMe namespace

You can use System Manager to edit the namespace by changing the subsystem that the namespace is mapped to.

About this task

You can only modify the NVMe subsystem settings in this window, you cannot edit the other namespace details.

Steps

1. Click **NVMe > NVMe namespaces**.

2. In the **NVMe namespaces window**, select the namespace you want to edit.
3. Select a subsystem option:
 - None: Choosing this option unmaps the existing subsystem mapping for this namespace only. This option is preselected if no subsystem mapping is present for the selected namespace.
 - Use an existing subsystem: This option is preselected if subsystem-to-namespace mapping is present. Choosing a different subsystem maps the new subsystem by unmapping the previously mapped subsystem.

Cloning an NVMe namespace

You can use System Manager to quickly create another namespace of the same configuration by choosing to clone a namespace. You can map the newly cloned namespace to another host NQN.

Before you begin

You must have a FlexClone license to clone a namespace.

About this task

You can clone a namespace with the selected host mapping and associate it with another subsystem.

Steps

1. Click **NVMe > NVMe namespaces**.
2. In the **NVMe namespaces window**, select the namespace you want to clone.
3. You can rename the cloned namespace if you need a specific name but it is not required.

The dialog provides a default name of the namespace to-be-cloned.

4. Modify the subsystem mapping for the cloned namespace.
5. Click **OK**.

The online, mapped namespace is cloned inside the same SVM with a different name. Host mapping will not be cloned.

Starting and stopping the NVMe service

The NVMe service enables you to manage NVMe adapters for use with namespaces. You can use System Manager to start the NVMe service to bring the adapters online. You can stop the NVMe service to take the NVMe adapters offline and to disable access to the namespaces.

Before you begin

NVMe capable adapters must be present before you start the NVMe service.

Steps

1. Click **Storage > SVMs**.
2. Select the SVM, and then click **SVM settings**.
3. In the **Protocols** menu, click **NVMe**.

4. Click **Start** or **Stop** service as required.

What NVMe is

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

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 bottlenecks 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 vserver in ONTAP.

An NVMe subsystem can be created using System Manager. You can associate the NVMe subsystem with different hosts and namespaces within the vserver. Also, each vserver 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.

Create NVMe subsystems

You can use System Manager to create an NVMe subsystem.

Steps

1. Click **Create** in the **NVMe Subsystems** window.
2. Provide entries in the **NVMe Subsystems: Create** window for the following fields:

- **SVM**

From the drop-down menu, select the SVM on which you want to create the subsystem.

- **Name**

Enter a name for the subsystem. The subsystem name cannot already exist in the SVM. The name is case-sensitive and is limited to 96 characters. Special characters are allowed.

- **Host OS**

From the drop-down menu, select the type of Host OS of the subsystem.

- **Host NQN**

Enter the Host NQN attached to the controller. You can enter more than one Host NQN by separating them with commas.

3. Click **Save**.

The NVMe subsystem is created, and the NVMe Subsystemswindow is displayed.

Related information

[NVMe Subsystems window](#)

Editing NVMe subsystems details

You can use System Manager to edit the details of an NVMe subsystem.

Steps

1. Find the NVMe subsystem you want to edit in the **NVMe Subsystem** window.
2. Check the box to the left of the name of the subsystem you want to edit.
3. Click **Edit**.

The current details of the NVMe subsystem are displayed in the NVMe Subsystems: Editwindow.

4. You can modify only the information in the **Host NQN** field.

- **Host NQN**

Modify the Host NQN attached to the controller. You can enter more than one Host NQN by separating them with commas.

The **Associated NVMe Namespaces** table displays below the Host NQN field. For each namespace, that table lists the namespace path and namespace ID.

5. Click **Save**.

The NVMe subsystem details are updated, and the NVMe Subsystems window is displayed.

Related information

[NVMe Subsystems window](#)

Deleting an NVMe subsystem

You can use System Manager to delete an NVMe subsystem from a cluster.

About this task

The following actions occur when you delete an NVMe subsystem:

- If the NVMe subsystem has configured hosts, then mapped hosts will be removed.
- If the NVMe subsystem has mapped namespaces, then they will be unmapped.

Steps

1. Find the NVMe subsystem you want to delete on the **NVMe Subsystem** window.
2. Check the box to the left of the name of the subsystem you want to delete.
3. Click **Delete**.

A Warning message is displayed.

4. Click the **Delete the NVMe Subsystem** check box to confirm the deletion, then click **Yes**.

The NVMe subsystem is deleted from the cluster, and the NVMe Subsystems window is displayed.

Related information

[NVMe Subsystems window](#)

NVMe Subsystems window

The NVMe Subsystems window displays by default an inventory list of NVMe subsystems in a cluster. You can filter the list to display only subsystems that are specific to an SVM. The window also enables you to create, edit, or delete NVMe subsystems. You can access this window by selecting **Storage > NVMe > Subsystems**.

- [NVMe Subsystems table](#)
- [Toolbar](#)

NVMe Subsystems table

The NVMe Subsystems table lists the inventory of NVMe subsystems in a cluster. You can refine the list by using the drop-down menu in the **SVM** field to select an SVM to display only the NVMe subsystems associated with that SVM. The **Search** field and **Filtering** drop-down menu enable you to further customize the list.

The NVMe Subsystems table contains the following columns:

- **(check box)**

Enables you to specify on which subsystems you want to perform actions.

Click the check box to select the subsystem, then click the action in the toolbar that you want to perform.

- **Name**

Displays the name of the subsystem.

You can search for a subsystem by entering its name in the **Search** field.

- **Host OS**

Displays the name of the host OS associated with the subsystem.

- **Host NQN**

Displays the NVMe Qualified Name (NQN) attached to the controller. If multiple NQNs are displayed, they are separated by commas.

- **Associated NVMe Namespaces**

Displays the number of the NVM namespaces associated with the subsystem. You can hover over the number to display the associated namespaces paths. Click on a path to display the Namespace Details window.

Toolbar

The toolbar is located above the column header. You can use the fields and buttons in the toolbar to perform various actions.

- **Search**

Enables you to search on values that might be found in the **Name** column.

- **Filtering**

Allows you to select from a drop-down menu that lists various methods of filtering the list.

- **Create**

Opens the Create NVMe Subsystem dialog box, which enables you to create an NVMe subsystem.

- **Edit**

Opens the Edit NVMe Subsystem dialog box, which enables you to edit an existing NVMe subsystem.

- **Delete**

Opens the Delete NVMe Subsystem confirmation dialog box, which enables you to delete an existing NVMe subsystem.

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.

NVMe namespaces window

You can use the NVMe namespaces window to set up and manage your namespaces and associated subsystems for the NVMe protocol. You can search for an existing namespace using the namespace path.

Command Buttons

- **Create**

Opens the NVMe namespace create dialog box, which allows you to set up a new namespace and map it to an NVMe subsystem.

- **Edit**

Enables you to edit the namespace mapping.

- **Delete**

Deletes the selected namespace.

- **More Actions**

Allows you to create a clone of the selected namespace, which can be associated with an existing subsystem, or you can choose not to map it to a subsystem.

- **Refresh**

Updates the information in the window.

NVMe List

- **Status**

Displays if the namespace is online or offline.

- **Namespace Path**

The path to the new namespace in the `/vol/volume'/file` format. The namespace path is a clickable link. Clicking the link takes you to the namespace detailspage.

- **NVMe Subsystem**

The name of the subsystem attached to a namespace. If no subsystems are attached, the value of this column is shown as `None`. You can see the list of unmapped namespaces by filtering this column for NVMe subsystem contains `None`.

- **SVMs**

The SVM name on which the namespace is created. The SVM name is a clickable link. Clicking the link takes you to the existing SVM dashboard page.

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.

- **Namespace ID**

A unique identifier used by the controller to provide access to a namespace. This is not a user input; it is generated by the system when the new namespace is created.

- **Total Space**

Displays the total size of the namespace.

- **Used Space**

Displays the amount of used space in the namespace.

- **%Used**

Displays the amount of space (in percentage) that is used in the namespace. The value for this field is calculated using total and used space.

Details Area

You can select a namespace to view information about the selected namespace. From this area, you can also edit, delete or clone the namespace.

- **Overview tab**

Displays general information about the selected namespace, and displays a pictorial representation of the space allocation of the namespace and the performance of the namespace.

In the Overview tab, the SVM and volume names are clickable links. Clicking the link takes you to the SVM and volume pages, respectively. The number of hosts can be one or more; by default two host names are shown. If more than two host names are shown, you can click a link to access the additional hosts.

The Overview tab also displays a space chart that shows the total and used space details for the namespace and a performance chart that shows details such as latency, IOPS, and throughput.

- **Status**

The status of the namespace; the value can be online or offline.

- **Host NQN**

The host NVMe Qualified Names (NQNs) uniquely describes the host for the purposes of identification and authentication. This field can accept comma separated NVMe qualification name (NQN) values. The host NQN starts with `nqn` and rest of the validation is the same as the initiator qualification name (IQN).

- **Host OS**

The host operating system for the namespace: Hyper-V, Linux, VMware, Windows or Xen.

- **Volume**

Displays the volume name on which the namespace is hosted.

- **Read-Only**

Displays whether the namespace is read-only or not.

- **Node**

The node that owns the namespace.

- **Block Size**

The size of the storage block.

- **Restore Inaccessible**

If unmapping a subsystem fails and partial data remains, unmapped namespaces cannot be restored.

iSCSI protocol

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

Related information

[SAN administration](#)

Create iSCSI aliases

An iSCSI alias is a user-friendly identifier that you assign to an iSCSI target device (in this case, the storage system) to make it easier to identify the target device in user interfaces. You can use System Manager to create an iSCSI alias.

About this task

An iSCSI alias is a string of 1 to 128 printable characters. An iSCSI alias must not include spaces.

Steps

1. Click **Storage > SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **Protocols** pane, click **iSCSI**.
4. In the **Service** tab of the **iSCSI** window, click **Edit**.
5. In the **Edit iSCSI Service Configuration** dialog box, enter an iSCSI alias in the **Target Alias** field, and then click **OK**.

Related information

[iSCSI window](#)

Enabling or disabling the iSCSI service on storage system interfaces

You can use System Manager to control which network interfaces are used for iSCSI communication by enabling or disabling the interfaces. When the iSCSI service is enabled, iSCSI connections and requests are accepted over those network interfaces that are enabled for iSCSI, but not over disabled interfaces.

Before you begin

You must have terminated any outstanding iSCSI connections and sessions that are currently using the interface. By default, the iSCSI service is enabled on all of the Ethernet interfaces after you enable the iSCSI license.

Steps

1. Click **Storage > SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **Protocols** pane, click **iSCSI**.
4. In the **iSCSI Interfaces** area, select the interface on which you want to enable or disable the iSCSI service.
5. Click **Enable** or **Disable**, as required.

Related information

[iSCSI window](#)

[Configuring iSCSI protocol on SVMs](#)

Add the security method for iSCSI initiators

You can use System Manager to add an initiator and to specify the security method that is used to authenticate the initiator.

Steps

1. Click **Storage > SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **Protocols** pane, click **iSCSI**.
4. In the **iSCSI** window, click the **Initiator Security** tab.
5. Click **Add** in the **Initiator Security** area.
6. Specify the initiator name and the security method for authenticating the initiator.

For CHAP authentication, you must provide the user name and password, and then confirm your password for inbound settings. For outbound settings, this login information is optional.

7. Click **OK**.

Related information

[iSCSI window](#)

Editing default security settings

You can use the Edit Default Security dialog box in System Manager to edit the default security settings for the iSCSI initiators that are connected to the storage system.

Steps

1. Click **Storage > SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **Protocols** pane, click **iSCSI**.
4. In the **Default Security** area of the **Initiator Security** tab, click **Edit**.
5. In the **Edit Default Security** dialog box, change the security type.

For CHAP authentication, you must provide the user name and password, and then confirm your password for inbound settings. For outbound settings, this login information is optional.

6. Click **OK**.

Related information

[iSCSI window](#)

Editing initiator security

The security style that is configured for an initiator specifies how authentication is done for that initiator during the iSCSI connection login phase. You can use System Manager to change the security for selected iSCSI initiators by changing the authentication method.

Steps

1. Click **Storage > SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **Protocols** pane, click **iSCSI**.
4. In the **Initiator Security** tab, select one or more initiators from the initiator list, and then click **Edit** in the **Initiator Security** area.
5. Change the security type.

For CHAP authentication, you must provide the user name and password, and then confirm your password for inbound settings. For outbound settings, this login information is optional.

6. Click **OK**.
7. Verify the changes that you made in the **Initiator Security** tab.

Related information

[iSCSI window](#)

Changing the default iSCSI initiator authentication method

You can use System Manager to change the default iSCSI authentication method, which

is the authentication method that is used for any initiator that is not configured with a specific authentication method.

Steps

1. Click **Storage > SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **Protocols** pane, click **iSCSI**.
4. In the **Initiator Security** tab, click **Edit** in the **Default Security** area.
5. Change the security type.

For CHAP authentication, you must provide the user name and password, and then confirm your password for inbound settings. For outbound settings, this login information is optional.

6. Click **OK**.

Related information

[iSCSI window](#)

Setting the default security for iSCSI initiators

You can use System Manager to remove the authentication settings for an initiator and to use the default security method to authenticate the initiator.

Steps

1. Click **Storage > SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **Protocols** pane, click **iSCSI**.
4. In the **Initiator Security** tab, select the initiator for which you want to change the security setting.
5. Click **Set Default** in the **Initiator Security** area, and then click **Set Default** in the confirmation dialog box.

Related information

[iSCSI window](#)

Starting or stopping the iSCSI service

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

Steps

1. Click **Storage > SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **Protocols** pane, click **iSCSI**.
4. Click **Start** or **Stop**, as required.

Related information

[iSCSI window](#)

Viewing initiator security information

You can use System Manager to view the default authentication information and all the initiator-specific authentication information.

Steps

1. Click **Storage > SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **Protocols** pane, click **iSCSI**.
4. In the **Initiator Security** tab of the **iSCSI** window, review the details.

iSCSI window

You can use the iSCSI window to start or stop the iSCSI service, change a storage system iSCSI node name, and create or change the iSCSI alias of a storage system. You can also add or change the initiator security setting for an iSCSI initiator that is connected to your storage system.

Tabs

- **Service**

You can use the **Service** tab to start or stop the iSCSI service, change a storage system iSCSI node name, and create or change the iSCSI alias of a storage system.

- **Initiator Security**

You can use the **Initiator Security** tab to add or change the initiator security setting for an iSCSI initiator that is connected to your storage system.

Command buttons

- **Edit**

Opens Edit iSCSI Service Configurations dialog box, which enables you to change iSCSI node name and iSCSI alias of the storage system.

- **Start**

Starts the iSCSI service.

- **Stop**

Stops the iSCSI service.

- **Refresh**

Updates the information in the window.

Details area

The details area displays information about the status of the iSCSI service, iSCSI target node name, and iSCSI target alias. You can use this area to enable or disable the iSCSI service on a network interface.

Related information

[Creating iSCSI aliases](#)

[Enabling or disabling the iSCSI service on storage system interfaces](#)

[Adding the security method for iSCSI initiators](#)

[Editing default security settings](#)

[Editing initiator security](#)

[Changing the default iSCSI initiator authentication method](#)

[Setting the default security for iSCSI initiators](#)

[Starting or stopping the iSCSI service](#)

FC/FCoE protocol

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

Related information

[SAN administration](#)

Starting or stopping the FC or FCoE service

The FC service enables you to manage FC target adapters for use with LUNs. You can use System Manager to start the FC service to bring the adapters online and to enable access to the LUNs on the storage system. You can stop the FC service to take the FC adapters offline and to disable access to the LUNs.

Before you begin

- The FC license must be installed.
- An FC adapter must be present in the target storage system.

Steps

1. Click **Storage > SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **Protocols** pane, click **FC/FCoE**.
4. Click **Start** or **Stop**, as required.

Related information

[FC/FCoE window](#)

Changing an FC or FCoE node name

If you replace a storage system chassis and reuse it in the same Fibre Channel SAN, the node name of the replaced storage system might be duplicated in certain cases. You can change the node name of the storage system by using System Manager.

Steps

1. Click **Storage > SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **Protocols** pane, click **FC/FCoE**.
4. Click **Edit**.
5. Type the new name, and then click **OK**.

Related information

[FC/FCoE window](#)

The FCoE protocol

Fibre Channel over Ethernet (FCoE) is a new model for connecting hosts to storage systems. Like the traditional FC protocol, FCoE maintains existing FC management and controls, but it uses a 10-gigabit Ethernet network as the hardware transport.

Setting up an FCoE connection requires one or more supported converged network adapters (CNAs) in the host, connected to a supported data center bridging (DCB) Ethernet switch. The CNA is a consolidation point and effectively serves as both an HBA and an Ethernet adapter.

In general, you can configure and use FCoE connections the same way you use traditional FC connections.

FC/FCoE window

You can use the FC/FCoE window to start or stop the FC service.

Command buttons

- **Edit**

Opens the Edit Node Name dialog box, which enables you to change the FC or FCoE node name.

- **Start**

Starts the FC/FCoE service.

- **Stop**

Stops the FC/FCoE service.

- **Refresh**

Updates the information in the window.

FC/FCoE details

The details area displays information about the status of FC/FCoE service, the node name, and the FC/FCoE adapters.

Related information

[Starting or stopping the FC or FCoE service](#)

[Changing an FC or FCoE node name](#)

[Configuring FC protocol and FCoE protocol on SVMs](#)

Export policies

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

Create an export policy

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

Steps

1. Click **Storage > SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **Policies** pane, click **Export Policies**.
4. Click **Create**.
5. In the **Create Export Policy** dialog box, specify a name for the export policy.
6. If you want to create an export policy by copying the rules from an existing export policy, select the **Copy Rules from** check box, and then select the storage virtual machine (SVM) and the export policy.

You should not select the destination SVM for disaster recovery from the drop-down menu to create an export policy.

7. In the **Export Rules** area, click **Add** to add rules to the export policy.
8. Click **Create**.
9. Verify that the export policy that you created is displayed in the **Export Policies** window.

Renaming export policies

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

Steps

1. Click **Storage > SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **Policies** pane, click **Export Policies**.
4. Select the export policy that you want to rename, and then click **Rename Policy**.

5. In the **Rename Policy** dialog box, specify a new policy name, and then click **Modify**.
6. Verify the changes that you made in the **Export Policies** window.

Deleting export policies

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

Steps

1. Click **Storage > SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **Policies** pane, click **Export Policies**.
4. Select the export policy that you want to delete, and then click **Delete Policy**.
5. Select the confirmation check box, and then click **Delete**.

Add rules to an export policy

You can use System Manager to add rules to an export policy, which enables you to define client access to data.

Before you begin

You must have created the export policy to which you want to add the export rules.

Steps

1. Click **Storage > SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **Policies** pane, click **Export Policies**.
4. Select the export policy to which you want to add rules, and from the **Export Rules** tab, click **Add**.
5. In the **Create Export Rule** dialog box, perform the following steps:

- a. Specify the client that requires access to the data.

You can specify multiple clients as comma-separated values.

You can specify the client in any of the following formats:

- As a host name; for instance, host1
- As an IPv4 address; for instance, 10.1.12.24
- As an IPv4 address with a network mask; for instance, 10.1.16.0/255.255.255.0
- As an IPv6 address; for instance, FE80::0202:B3FF:FE1E:8329
- As an IPv6 address with a network mask; for instance, 2001:db8::/32
- As a netgroup, with the netgroup name preceded by an at symbol (@); for instance, @netgroup
- As a domain name preceded by a period (.); for instance, .example.com



You must not enter an IP address range, such as 10.1.12.10 through 10.1.12.70. Entries in this format are interpreted as a text string and are treated as a host name.

+ You can enter the IPv4 address 0.0.0.0/0 to provide access to all of the hosts.

- a. If you want to modify the rule index number, select the appropriate rule index number.
- b. Select one or more access protocols.

If you do not select any access protocol, the default value “Any” is assigned to the export rule.

- c. Select one or more security types and access rules.

6. Click **OK**.

7. Verify that the export rule that you added is displayed in the **Export Rules** tab for the selected export policy.

Modifying export policy rules

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

Steps

1. Click **Storage > SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **Policies** pane, click **Export Policies**.
4. In the **Export Policies** window, select the export policy for which you want to edit the export rule, and in the **Export Rules** tab, select the rule that you want to edit, and then click **Edit**.
5. Modify the following parameters as required:
 - Client specification
 - Access protocols
 - Access details
6. Click **OK**.
7. Verify that the updated changes for the export rule are displayed in the **Export Rules** tab.

Related information

[Setting up CIFS](#)

Deleting export policy rules

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

Steps

1. Click **Storage > SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **Policies** pane, click **Export Policies**.
4. Select the export policy for which you want to delete the export rule.
5. In the **Export Rules** tab, select the export rule that you want to delete, and then click **Delete**.
6. In the confirmation box, click **Delete**.

How export policies control client access to volumes or qtrees

Export policies contain one or more *export rules* that process each client access request. The result of the process determines whether the client is denied or granted access and what level of access. An export policy with export rules must exist on the storage virtual machine (SVM) for clients to access data.

You associate exactly one export policy with each volume or qtree to configure client access to the volume or qtree. The SVM can contain multiple export policies. This enables you to do the following for SVMs with multiple volumes or qtrees:

- Assign different export policies to each volume or qtree of the SVM for individual client access control to each volume or qtree in the SVM.
- Assign the same export policy to multiple volumes or qtrees of the SVM for identical client access control without having to create a new export policy for each volume or qtree.

If a client makes an access request that is not permitted by the applicable export policy, the request fails with a permission-denied message. If a client does not match any rule in the export policy, then access is denied. If an export policy is empty, then all accesses are implicitly denied.

You can modify an export policy dynamically on a system running ONTAP.

Export Policies window

You can use the Export Policies window to create, view, and manage information about export policies and its related export rules.

Export Policies

The Export Policies window enables you to view and manage the export policies created for the storage virtual machine (SVM).

- **Command buttons**

- Create

Opens the Create Export Policy dialog box, which enables you to create an export policy and add export rules. You can also copy export rules from an existing SVM.

- Rename

Opens the Rename Policy dialog box, which enables you to rename the selected export policy.

- Delete

Opens the Delete Export Policy dialog box, which enables you to delete the selected export policy.

- Refresh

Updates the information in the window.

Export Rules tab

The Export Rules tab enables you to view information about the export rules created for a particular export policy. You can also add, edit, and delete rules.

- **Command buttons**

- Add

Opens the Create Export Rule dialog box, which enables you to add an export rule to the selected export policy.

- Edit

Opens the Modify Export Rule dialog box, which enables you to modify the attributes of the selected export rule.

- Delete

Opens the Delete Export Rule dialog box, which enables you to delete the selected export rule.

- Move Up

Moves up the rule index of the selected export rule.

- Move Down

Moves down the rule index of the selected export rule.

- Refresh

Updates the information in the window.

- **Export rules list**

- Rule Index

Specifies the priority based on which the export rules are processed. You can use the Move Up and Move Down buttons to choose the priority.

- Client

Specifies the client to which the rule applies.

- Access Protocols

Displays the access protocol that is specified for the export rule.

If you have not specified any access protocol, the default value "Any" is considered.

- Read-Only Rule

Specifies one or more security types for read-only access.

- Read/Write Rule

Specifies one or more security types for read/write access.

- Superuser Access

Specifies the security type or types for superuser access.

Assigned Objects tab

The Assigned Objects tab enables you to view the volumes and qtrees that are assigned to the selected export policy. You can also view whether the volume is encrypted or not.

Efficiency policies

You can use System Manager to create, edit, and delete efficiency policies.

Add efficiency policies

You can use System Manager to add efficiency policies for running the deduplication operation on a volume on a specified schedule or when the change in volume data reaches a specified threshold value.

Steps

1. Click **Storage > SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **Policies** pane, click **Efficiency Policies**.
4. Click **Add**, and then specify the policy name.
5. Specify how the storage efficiency policy should be run:
 - Select **Schedule**, and specify the schedule name and the schedule details.

You can specify the maximum run-time duration of the efficiency policy, if required.
 - Select **ChangeLog Threshold**, and specify the threshold value (in percent) for the change in volume data.
6. Select the **Set QoS policy to background** check box to reduce performance impact on client operations.
7. Click **Add**.

Editing efficiency policies

You can use System Manager to modify the attributes of an efficiency policy such as the policy name, schedule name, and maximum runtime.

Steps

1. Click **Storage > SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **Policies** pane, click **Efficiency Policies**.
4. In the **Efficiency Policies** window, select the policy that you want to edit, and then click **Edit**.
5. In the **Edit Efficiency Policy** dialog box, make the required changes.

6. Click **Save**.

Deleting efficiency policies

You can use System Manager to delete an efficiency policy that is no longer required.

Before you begin

The efficiency policy must be disabled.

Steps

1. Click **Storage > SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **Policies** pane, click **Efficiency Policies**.
4. Select the efficiency policy that you want to delete, and then click **Delete**.
5. Select the confirmation check box, and then click **Delete**.

Enabling or disabling efficiency policies

You can use System Manager to enable or disable an efficiency policy.

Steps

1. Click **Storage > SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **Policies** pane, click **Efficiency Policies**.
4. Select one or more efficiency policies that you want to enable or disable.
5. Click **Status > Enable** or **Status > Disable**, as required.
6. If you are disabling an efficiency policy, select the confirmation check box, and then click **OK**.

What an efficiency policy is

An efficiency policy is a job schedule for a deduplication operation on a FlexVol volume.

You can run deduplication on a FlexVol volume either by scheduling the operations to start at a specific time or by specifying that the operations are triggered if a threshold percentage is exceeded. You can schedule a deduplication operation by creating job schedules that are enclosed within the efficiency policies. The volume efficiency policies support only job schedules that are of type cron. Alternately, you can specify a threshold percentage. When new data exceeds the specified percentage, the deduplication operation is started.

Understanding predefined efficiency policies

You can configure a volume with efficiency policies to achieve additional space savings. You can configure a volume to run inline compression without a scheduled or manually started background efficiency operation configured on the volume.

When you create an SVM, the following efficiency policies are created automatically and cannot be deleted:

- Default

You can configure a volume with the efficiency policy to run the scheduled deduplication operations on the volume.

- **Inline-only**

You can configure a volume with the inline-only efficiency policy and enable inline compression, to run inline compression on the volume without any scheduled or manually started background efficiency operations.

For more information about the inline-only and default efficiency policies, see the man pages.

Efficiency Policies window

You can use the Efficiency Policies window to create, display, and manage information about efficiency policies.

Command buttons

- **Add**

Opens the Add Efficiency Policy dialog box, which enables you to run a deduplication operation on a volume for a specified duration (schedule-based) or when the change in volume data reaches a specified threshold value (threshold-based).

- **Edit**

Opens the Edit Efficiency Policy dialog box, which enables you to modify the schedule, threshold value, QoS type, and maximum run time for a deduplication operation.

- **Delete**

Opens the Delete Efficiency Policy dialog box, which enables you to delete the selected efficiency policy.

- **Status**

Open a drop-down menu, which provides options to enable or disable the selected efficiency policy.

- **Refresh**

Updates the information in the window.

Efficiency policies list

- **Policy**

Specifies the name of an efficiency policy.

- **Status**

Specifies the status of an efficiency policy. The status can be one of the following:

- Enabled

Specifies that the efficiency policy can be assigned to a deduplication operation.

- Disabled

Specifies that the efficiency policy is disabled. You can enable the policy by using the status drop-down menu and assign it later to a deduplication operation.

- **Run By**

Specifies whether the storage efficiency policy is run based on a schedule or based on a threshold value (change log threshold).

- **QoS Policy**

Specifies the QoS type for the storage efficiency policy. The QoS type can be one of the following:

- Background

Specifies that the QoS policy is running in the background, which reduces potential performance impact on the client operations.

- Best-effort

Specifies that the QoS policy is running on a best-effort basis, which enables you to maximize the utilization of system resources.

- **Maximum Runtime**

Specifies the maximum run-time duration of an efficiency policy. If this value is not specified, the efficiency policy is run till the operation is complete.

Details area

The area below the efficiency policy list displays additional information about the selected efficiency policy, including the schedule name and the schedule details for a schedule-based policy, and the threshold value for a threshold-based policy.

Protection policies

You can use System Manager to create, edit, and delete protection policies.

Create protection policies

You can use System Manager to create asynchronous mirror policies, vault policies, or mirror and vault policies, and to apply these policies to a data protection relationship.

Steps

1. Click **Storage > SVMs**.
2. Select the storage virtual machine (SVM) for which you want to create a protection policy, and then click **SVM Settings**.
3. In the **Policies** pane, click **Protection Policies**.
4. Click **Create**.
5. In the **Create Policy** dialog box, select the policy type that you want to create.

6. Specify the policy name and transfer priority.

Low indicates that the transfer has the least priority, and the transfer is usually scheduled after normal priority transfers. By default, the priority is set to Normal.

7. For a policy of type asynchronous mirror, select the **Transfer All Source Snapshot Copies** check box to include the “all_source_snapshots” rule to the mirror policy, which backs up all of the Snapshot copies from the source volume.

8. Select the **Enable Network Compression** check box to compress the data that is being transferred during a data transfer.

9. Click **Add Comments** to add additional comments for the policy.

10. For a policy of type vault or mirror vault, specify a SnapMirror label and a destination retention count.

11. Click **Create**.

Deleting protection policies

You can use System Manager to delete a protection policy if you no longer want to use the policy.

About this task

The cluster-level mirror policies or vault policies are not displayed.

Steps

1. Click **Storage > SVMs**.
2. Select the storage virtual machine (SVM), and then click **SVM Settings**.
3. In the **Protection Policies** window, select the policy that you want to delete, and then click **Delete**.
4. In the **Delete Policy** dialog box, click **Delete**.

Editing protection policies

You can use System Manager to modify a protection policy and to apply the policy to a data protection relationship.

About this task

The protection policies are not displayed at the cluster level.

Steps

1. Click **Storage > SVMs**.
2. Select the storage virtual machine (SVM), and then click **SVM Settings**.
3. In the **Policies** pane, click **Protection Policies**.
4. Select the protection policy that you want to edit, and then click **Edit**.
5. Modify the transfer priority, and then enable or disable network compression.
6. For an asynchronous mirror policy, back up all of the source Snapshot copies.
7. For a vault policy or mirror vault policy, modify the SnapMirror label and retention count.

You cannot remove the sm_created label for a mirror vault policy.

8. Click **Save**.

Protection Policies window

You can use the Protection Policies window to create, manage, and display information about mirror, vault, and mirror vault policies.

Command buttons

- **Create**

Opens the Create Policy dialog box, which enables you to create a mirror, vault, or mirror vault policy.

- **Edit**

Opens the Edit Policy dialog box, which enables you to edit a policy.

- **Delete**

Opens the Delete Policy dialog box, which enables you to delete a policy.

- **Refresh**

Updates the information in the window.

Protection policies list

- **Name**

Displays the name of the protection policy.

- **Type**

Displays the policy type, which can be Vault, Mirror Vault, or Asynchronous Mirror.

- **Comment**

Displays the description specified for the policy.

- **Transfer Priority**

Displays the data transfer priority, such as Normal or Low.

Details area

- **Policy Details tab**

Displays details of the protection policy, such as the user who created the policy, number of rules, retention count, and status of network compression.

- **Policy Rules tab**

Displays details of the rules that are applied to the policy. The Policy Rules tab is displayed only if the selected policy contains rules.

QoS policy groups

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

Create QoS policy groups

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

Steps

1. Click **Storage > SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **Policies** pane, click **QoS Policy Groups**.
4. In the **QoS Policy Groups** window, click **Create**.
5. In the **Create Policy Group** dialog box, specify a group name for the policy.
6. Specify the minimum throughput limit.
 - In System Manager 9.5, you can set the minimum throughput limit only on a performance-based All Flash Optimized personality. In System Manager 9.6, you can also set the minimum throughput limit for ONTAP Select Premium systems.
 - You cannot set the minimum throughput limit for volumes on a FabricPool-enabled aggregate.
 - If you do not specify the minimum throughput value or if the minimum throughput value is set to 0, the system automatically displays “None” as the value.

This value is case-sensitive.

7. Specify the maximum throughput limit.
 - The minimum throughput limit and the maximum throughput limit must be of the same unit type.
 - If you do not specify the minimum throughput limit, you can set the maximum throughput limit in IOPS and B/s, KB/s, MB/s, and so on.
 - If you do not specify the maximum throughput limit, the system automatically displays “Unlimited” as the value.

This value is case-sensitive. The unit that you specify does not affect the maximum throughput.

8. Click **OK**.

Deleting QoS policy groups

You can use System Manager to delete a Storage Quality of Service (QoS) policy group that is no longer required.

Before you begin

You must have unassigned all of the storage objects that are assigned to the policy group.

Steps

1. Click **Storage > SVMs**.

2. Select the SVM, and then click **SVM Settings**.
3. In the **Policies** pane, click **QoS Policy Groups**.
4. In the **QoS Policy Groups** window, select the policy group that you want to delete, and then click **Delete**.
5. In the confirmation dialog box, click **Delete**.

Editing QoS policy groups

You can use the Edit Policy Group dialog box in System Manager to modify the name and maximum throughput of an existing storage Quality of Service (QoS) policy group.

About this task

- In System Manager 9.5, you can set the minimum throughput limit only on a performance-based All Flash Optimized personality. In System Manager 9.6, you can also set the minimum throughput limit for ONTAP Select Premium systems.
- You cannot set the minimum throughput limit for volumes on a FabricPool-enabled aggregate.

Steps

1. Click **Storage > SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **Policies** pane, click **QoS Policy Groups**.
4. Select the QoS policy group that you want to edit, and then click **Edit**.
 - The minimum throughput limit and the maximum throughput limit must be of the same unit type.
 - If you do not specify the minimum throughput limit, you can set the maximum throughput limit in IOPS and B/s, KB/s, MB/s, and so on.
 - If you do not specify the maximum throughput limit, the value is set to unlimited, and the unit that you specify does not affect the maximum throughput.
5. In the **Edit Policy Group** dialog box, edit the QoS policy group details, and then click **Save**.

Managing workload performance by using Storage QoS

Storage Quality of Service (QoS) can help you manage risks around meeting your performance objectives. You can use Storage QoS to limit the throughput to workloads and to monitor workload performance. You can reactively limit workloads to address performance problems, and you can proactively limit workloads to prevent performance problems.


A workload represents the input/output (I/O) operations to one of the following kinds of storage objects:

- FlexVol volumes
- LUNs
- FlexGroup volumes

You can assign a storage object to a policy group to control and monitor a workload. You can monitor workloads without controlling them.

The following illustration shows a sample environment before and after using Storage QoS. On the left, the


workloads compete for cluster resources to transmit I/O. These workloads get “best effort” performance, which means that you have less performance predictability (for example, a workload might get such good performance that it negatively impacts other workloads). On the right, the same workloads are assigned to policy groups. The policy groups enforce a maximum throughput limit.



How Storage QoS works

Storage QoS controls workloads that are assigned to policy groups by throttling and prioritizing client operations (SAN and NAS data requests) and system operations.

The following illustration shows a sample environment before and after using Storage QoS. On the left, workloads compete for cluster resources to transmit I/O. These workloads get “best effort” performance, which means that you have less performance predictability (for example, a workload might get such good performance that it negatively impacts other workloads). On the right, the same workloads are assigned to policy groups that enforce maximum throughput limits.



The `-max-throughput` parameter specifies the maximum throughput limit for the policy group that the policy group must not exceed. The value of this parameter is specified in terms of IOPS or MB/s, or a combination of comma-separated IOPS and MB/s values, and the range is zero to infinity.

The units are base 10. There should be no space between the number and the unit. The default value for the `-max-throughput` parameter is `infinity`, which is specified by the special value `INF`.



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.



The combined throughput to the workloads in a policy group might exceed the specified limit by up to 10 percent. A deviation might occur if you have a workload that experiences rapid changes in throughput (sometimes called a *bursty workload*).

Note the following about specifying a maximum throughput:

- You must not set the limit too low because you might underutilize the cluster.
- You must consider the minimum amount of throughput that you want to reserve for workloads that do not have limits.

For example, you can ensure that your critical workloads get the throughput that they need by limiting noncritical workloads.

- You might want to provide room for growth.

For example, if you see an average utilization of 500 IOPS, you might specify a limit of 1,000 IOPS.

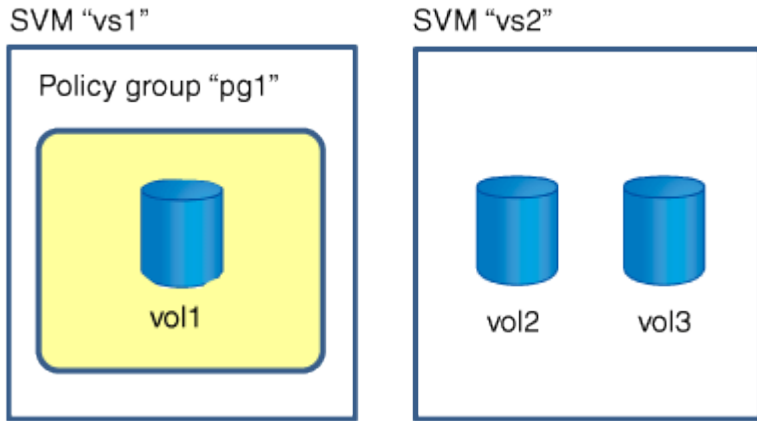
Rules for assigning storage objects to policy groups

You should be aware of rules that dictate how you can assign storage objects to Storage QoS policy groups.

Storage objects and policy groups must belong to the same SVM

A storage object must be contained by the SVM to which the policy group belongs. You specify the SVM to which the policy group belongs when you create the policy group. Multiple policy groups can belong to the same SVM.

In the following illustration, the policy group pg1 belongs to SVM vs1. You cannot assign volumes vol2 or vol3 to policy group pg1 because those volumes are contained by a different SVM.




Nested storage objects cannot belong to policy groups

You cannot assign a storage object to a policy group if its containing object or its child objects belong to a policy group. The following table lists the restrictions.

If you assign the...	Then you cannot assign...
SVM to a policy group	Any storage objects contained by the SVM to a policy group
Volume to a policy group	The volume's containing SVM or any child LUNs to a policy group
LUN to a policy group	The LUN's containing volume or SVM to a policy group
File to a policy group	The file's containing volume or SVM to a policy group

In the following illustration, the SVM vs3 is assigned to policy group pg2. You cannot assign volumes vol4 or vol5 to a policy group because an object in the storage hierarchy (SVM vs3) is assigned to a policy group.



QoS Policy Groups window

Storage QoS (Quality of Service) can help you manage risks related to meeting your performance objectives. Storage QoS enables you to limit the throughput of workloads and to monitor workload performance. You can use the QoS Policy groups window to manage your policy groups and view information about them.

Command buttons

- **Create**

Opens the Create QoS Policy Group dialog box, which enables you to create new policy groups.

- **Edit**

Opens the Edit QoS Policy Group dialog box, which enables you to modify the selected policy group.

- **Delete**

Deletes the selected policy groups.

- **Refresh**

Updates the information in the window.

QoS Policy Groups list

The QoS Policy Groups list displays the policy group name and the maximum throughput for each policy group.

- **Name**

Displays the name of the QoS policy group.

- **Minimum Throughput**

Displays the minimum throughput limit specified for the policy group.

If you have not specified any minimum throughput value, the system automatically displays “None” as the value and this value is case-sensitive.

- **Maximum Throughput**

Displays the maximum throughput limit specified for the policy group.

If you have not specified any maximum throughput value, the system automatically displays “Unlimited” as the value and this value is case-sensitive.

- **Storage Objects Count**

Displays the number of storage objects assigned to the policy group.

Details area

The area below the QoS Policy Groups list displays detailed information about the selected policy group.

- **Assigned Storage Objects tab**

Displays the name and type of the storage object that is assigned to the selected policy group.

NIS services

You can use System Manager to add, edit, and manage Network Information Service (NIS) domains.

Related information

[NFS configuration](#)

Add NIS domains

You can maintain host information centrally by using NIS. You can use System Manager to add the NIS domain name of your storage system. Only one NIS domain can be active on a storage virtual machine (SVM) at any given time.

Steps

1. Click **Storage > SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **Services** pane, click **NIS**.
4. Click **Create**.
5. Type the NIS domain name, and then add one or more NIS servers.
6. Click **Create**.

Editing NIS domains

You can use System Manager to modify NIS domains based on the requirement for storage virtual machine (SVM) authentication and authorization.

Steps

1. Click **Storage > SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **Services** pane, click **NIS**.
4. Select the NIS domain, and then click **Edit**.
5. Make the required changes, and then click **Edit**.

NIS window

The NIS window enables you to view the current NIS settings for your storage system.

Command buttons

- **Create**

Opens the Create NIS Domain dialog box, which enables you to create NIS domains.

- **Edit**

Opens the Edit NIS Domain dialog box, which enables you to add, delete, or modify NIS servers.

- **Delete**

Deletes the selected NIS domain.

- **Refresh**

Updates the information in the window.

LDAP client services

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



Add an LDAP client configuration

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

About this task

At the SVM level, you can add an LDAP client only for a selected SVM.

Steps

1. Add an LDAP client configuration by using one of the following methods:
 - Cluster level: click  > **LDAP**.
 - SVM level: click **SVM** > **SVM Settings** > **LDAP Client**.
2. Click **Add**.
3. Type the name of the LDAP client.
4. Add either the Active Directory domain or the LDAP server.
5. Click  (advanced options), select the **Schema**, and click **Apply**.
6. Specify the **Base DN** and **TCP Port**.
7. Click **Binding**, and then specify the authentication details.
8. Click **Save and Close**.
9. Verify that the LDAP client that you added is displayed.

Related information

[LDAP](#)


Deleting an LDAP client configuration

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

About this task

At the SVM level, you can delete an LDAP client only for a selected SVM.

Steps

1. To delete an LDAP client configuration:
 - Cluster level: Click  > **LDAP**.
 - SVM level: Click **SVM** > **SVM Settings** > **LDAP Client**.
2. Select the LDAP client that you want to delete, and then click **Delete**.
3. Select the confirmation check box, and then click **Delete**.
4. Verify that the LDAP client that you deleted is no longer displayed.

Related information

[LDAP](#)


Editing an LDAP client configuration

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

About this task

At the SVM level, you can edit an LDAP client only for a selected SVM.

Steps

1. To edit an LDAP client configuration:
 - Cluster level: Click  > **LDAP**.
 - SVM level: Click **SVM** > **SVM Settings** > **LDAP Client**.
2. Select the LDAP client that you want to modify, and then click **Edit**.
3. In the **Edit LDAP Client** dialog box, edit the LDAP client configuration as required.
4. Click **Save and Close**.
5. Verify that the changes that you made to the LDAP client configuration are displayed.

Related information

[LDAP](#)

LDAP Client window

You can use the LDAP Client window to create LDAP clients for user authentication, file access authorization, user search, and mapping services between NFS and CIFS at the storage virtual machine (SVM) level.

Command buttons

- **Add**

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

- **Edit**

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

- **Delete**

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

- **Refresh**

Updates the information in the window.

LDAP client list

Displays (in tabular format) details about LDAP clients.

- **LDAP Client Configuration**

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

- **Storage Virtual Machine**

Displays the name of the SVM for each LDAP client configuration.

- **Schema**

Displays the schema for each LDAP client.

- **Minimum Bind Level**

Displays the minimum bind level for each LDAP client.

- **Active Directory Domain**

Displays the Active Directory domain for each LDAP client configuration.

- **LDAP Servers**

Displays the LDAP server for each LDAP client configuration.

- **Preferred Active Directory Servers**

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

LDAP configuration services

You can use System Manager to manage LDAP configurations.

Editing active LDAP clients

You can use System Manager to associate an active LDAP client with a storage virtual machine (SVM), which enables you to use LDAP as a name service or for name mapping.

Steps

1. Click **Storage > SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **Services** pane, click **LDAP Configuration**.
4. In the **LDAP Configuration** window, click **Edit**.
5. In the **Active LDAP Client** dialog box, select the LDAP client that you want to edit, and perform the following actions:
 - Modify the Active Directory domain servers.
 - Modify the preferred Active Directory servers.
6. Click **OK**.
7. Verify that the changes that you made are updated in the **LDAP Configuration** window.

Deleting active LDAP clients

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

Steps

1. Click **Storage > SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. Click the **SVM Settings** tab.
4. In the **Services** pane, click **LDAP Configuration**.
5. Click **Delete**.
6. Select the confirmation check box, and then click **Delete**.

LDAP Configuration window

You can use the LDAP Configuration window to edit or delete active LDAP clients at the storage virtual machine (SVM) level.

Command buttons

- **Edit**

Opens the Active LDAP Client dialog box, which enables you to edit the properties of the active LDAP

client, such as Active Directory domain servers and preferred Active Directory servers.

- **Delete**

Opens the Delete Active LDAP Client dialog box, which enables you to delete the active LDAP client.

- **Refresh**

Updates the information in the window.

LDAP Configuration area

Displays the details about the active LDAP client.

- **LDAP client name**

Displays the name of the active LDAP client.

- **Active Directory Domain Servers**

Displays the Active Directory domain for the active LDAP client.

- **Preferred Active Directory Servers**

Displays the preferred Active Directory server for the active LDAP client.

Kerberos realm services

You can use System Manager to create and manage Kerberos realm services.

Related information

[NFS management](#)

Create a Kerberos realm configuration

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

Before you begin

- The CIFS license must be installed if CIFS shares are used, and the NFS license must be installed if an LDAP server is used.
- Active Directory (Windows 2003 or Windows 2008) with DES MD5 encryption capability must be available.
- You must have set the time zone and synchronized the time across the cluster by configuring NTP.

This prevents authentication errors, and ensures that the timestamps in log files are consistent across the cluster.

About this task

While creating a Kerberos realm, you must set the following attributes in the Create Kerberos Realm wizard:

- Kerberos realm
- KDC IP address and port number

The default port number is 88.

- Kerberos Key Distribution Center (KDC) vendor
- Administrative server IP address if the KDC vendor is not Microsoft
- Password server IP address
- Active Directory server name and IP address if the KDC vendor is Microsoft

Steps

1. Click **Storage > SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **Services** pane, click **Kerberos Realm**.
4. In the **Kerberos Realm** window, click **Create**.
5. Type or select information as prompted by the wizard.
6. Confirm the details, and then click **Finish** to complete the wizard.

Related information

[Setting the time zone for a cluster](#)

[NetApp Technical Report 4067: NFS in NetApp ONTAP](#)

[NetApp Technical Report 4616: NFS Kerberos in ONTAP with Microsoft Active Directory](#)

[NetApp Technical Report 4835: How to Configure LDAP in ONTAP](#)

Editing a Kerberos realm configuration

You can use System Manager to edit a Kerberos realm configuration at the storage virtual machine (SVM) level.

About this task

You can modify the following attributes by using the Kerberos Realm Edit wizard:

- The KDC IP address and port number
- The IP address of the administrative server if the KDC vendor is not Microsoft
- The IP address of the password server
- The Active Directory server name and IP address if the KDC vendor is Microsoft

Steps

1. Click **Storage > SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **Services** pane, click **Kerberos Realm**.

4. In the **Kerberos Realm** window, select the Kerberos realm configuration that you want to modify, and then click **Edit**.
5. Type or select information as prompted by the wizard.
6. Confirm the details, and then click **Finish** to complete the wizard.

Deleting Kerberos realm configurations

You can use System Manager to delete a Kerberos realm configuration.

Steps

1. Click **Storage > SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **Services** pane, click **Kerberos Realm**.
4. In the **Kerberos Realm** window, select one or more Kerberos realm configurations that you want to delete, and then click **Delete**.
5. Select the confirmation check box, and then click **Delete**.

Using Kerberos with NFS for strong security

You can use Kerberos to provide strong authentication between SVMs and NFS clients to provide secure NFS communication. Configuring NFS with Kerberos increases the integrity and security of NFS client communications with the storage system.

Kerberos authentication for CIFS

With Kerberos authentication, upon connection to your CIFS server, the client negotiates the highest possible security level. However, if the client cannot use Kerberos authentication, Microsoft NTLM or NTLM V2 is used to authenticate with the CIFS server.

Kerberos Realm window

You can use the Kerberos Realm window to provide authentication between storage virtual machines (SVMs) and NFS clients to ensure secure NFS communication.

Command buttons

- **Create**

Opens the Kerberos Realm Create wizard, which enables you to configure a Kerberos realm to retrieve user information.

- **Edit**

Opens the Kerberos Realm Edit wizard, which enables you to edit a Kerberos realm configuration based on the requirement for SVM authentication and authorization.

- **Delete**

Opens the Delete Kerberos Realm(s) dialog box, which enables you to delete Kerberos realm

configuration.

- **Refresh**

Updates the information in the window.

Kerberos Realm list

Provides details about the Kerberos realms, in tabular format.

- **Realm**

Specifies the name of the Kerberos realm.

- **KDC Vendor**

Specifies the name of the Kerberos Distribution Center (KDC) vendor.

- **KDC IP Address**

Specifies the KDC IP address used by the configuration.

Details area

The details area displays information such as the KDC IP address and port number, KDC vendor, administrative server IP address and port number, Active Directory server and server IP address of the selected Kerberos realm configuration.

Kerberos interface services

You can use System Manager to manage Kerberos interface services.

Editing Kerberos configuration

You can use System Manager to enable Kerberos and to edit a Kerberos configuration that is associated with a storage virtual machine (SVM), which enables the SVM to use Kerberos security services for NFS.

Before you begin

- You must have at least one Kerberos realm configured at the SVM level.
- You must have a minimum of two data LIFs on the SVM.

One data LIF is used by the Service Principal Name (SPN) for both the UNIX and CIFS-related Kerberos traffic. The other data LIF is used for accessing non-Kerberos traffic.



A CIFS server is not required for basic NFS Kerberos access. A CIFS server is required for multiprotocol access or when using Active Directory as an LDAP server for name mapping purposes.

About this task

If you are using Microsoft Active Directory Kerberos, the first 15 characters of any SPNs that are used in the

domain must be unique. Microsoft Active Directory has a limitation for SPNs of 15 characters maximum and does not allow duplicate SPNs.

Steps

1. Click **Storage > SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **Services** pane, click **Kerberos Interface**.
4. In the **Kerberos Interface** window, select the interface, and then click **Edit**.
5. In the **Edit Kerberos Configuration** dialog box, make the required changes, and then click **OK**.

Related information

[NetApp Technical Report 4067: NFS in NetApp ONTAP](#)

[NetApp Technical Report 4616: NFS Kerberos in ONTAP with Microsoft Active Directory](#)

[NetApp Technical Report 4835: How to Configure LDAP in ONTAP](#)

Kerberos Interface window

You can use the Kerberos Interface window to enable Kerberos and to edit the Kerberos configuration for storage virtual machines (SVMs).

Command buttons

- **Edit**

Opens the Edit Kerberos Configuration dialog box, which you can use to enable Kerberos and to edit the Kerberos configuration associated with the SVM.

- **Refresh**

Updates the information in the window.

Kerberos Interface list

Provides details about the Kerberos configuration.

- **Interface Name**

Specifies the logical interfaces associated with the Kerberos configuration for SVMs.

- **Service Principal Name**

Specifies the Service Principal Name (SPN) that matches the Kerberos configuration.

- **Realm**

Specifies the name of the Kerberos realm associated with the Kerberos configuration.

- **Kerberos Status**

Specifies whether Kerberos is enabled.

DNS/DDNS Services

You can use System Manager to manage DNS/DDNS services.

Enabling or disabling DDNS

You can use System Manager to enable or disable DDNS on a storage system.

About this task

- DNS is enabled by default.
- DDNS is disabled by default.
- System Manager does not perform any validation checks for the DNS and DDNS settings.

Steps

1. Click **Storage > SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **Services** pane, click **DNS/DDNS**.
4. In the **DNS/DDNS Services** window, click **Edit**.
5. In the **Edit DNS/DDNS Settings** dialog box, enable DDNS by selecting the **DDNS service** check box.

You can disable DDNS by clearing the **DDNS service** check box.

6. Click **OK**.

Related information

[DNS/DDNS Services window](#)

Editing DNS and DDNS settings

You can maintain host information centrally by using DNS. You can use System Manager to add or modify the DNS domain name of your storage system. You can also enable DDNS on your storage system to update the name server automatically in the DNS server.

Before you begin

You must have set up a CIFS server or an Active Directory account for the storage virtual machine (SVM) for secure DDNS to work.

About this task

System Manager does not perform any validation checks for the DNS and DDNS settings.

Steps

1. Click **Storage > SVMs**.
2. Select the SVM, and then click **SVM Settings**.

3. In the **Services** pane, click **DNS/DDNS**.
4. Click **Edit**.
5. In the **DNS Domains and Name Servers** area, add or modify the DNS domain names and the IP addresses.
6. Select the **DDNS service** check box to enable DDNS.
 - a. Select the **Enable Secure DDNS** check box to enable secure DDNS.
 - b. Specify the fully qualified domain name (FQDN) and the time to live value for the DDNS service.

By default, time to live is set to 24 hours and FQDN is set to `SVM name . domain name`.
7. Click **OK** to save the changes that you made.

Related information

[DNS/DDNS Services window](#)

DNS/DDNS Services window

The DNS/DDNS Services window enables you to view and edit the current DNS and DDNS settings for your system.

Command buttons

- **Edit**

Opens the Edit DNS/DDNS Settings dialog box, which you can use to add or modify DNS or DDNS details. You can also enable or disable DDNS.

- **Refresh**

Updates the information in the window.

Related information

[Enabling or disabling DDNS](#)

[Editing DNS and DDNS settings](#)

Users

You can use System Manager to create and manage storage virtual machine (SVM) user accounts.

Add SVM user accounts

You can use System Manager to add a storage virtual machine (SVM) user account and to specify a user login method for accessing the storage system.

Steps

1. Click **Storage > SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **SVM User Details** pane, click **Users**.
4. Click **Add**.
5. Specify a user name and password for connecting to the storage system, and confirm the password.
6. Add one or more user login methods, and then click **Add**.

A login method for the new vsadmin account is automatically included that uses HTTP as the application and is authenticated with a certificate.

Changing the password for SVM user accounts

You can use System Manager to reset the password for a storage virtual machine (SVM) user account.

Steps

1. Click **Storage > SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **SVM User Details** pane, click **Users**.
4. Select the user account for which you want to modify the password, and then click **Reset Password**.
5. In the **Reset Password** dialog box, type the new password, confirm the new password, and then click **Change**.

Editing SVM user accounts

You can use System Manager to edit a storage virtual machine (SVM) user account by modifying the user login methods for accessing the storage system.

Steps

1. Click **Storage > SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **SVM User Details** pane, click **Users**.
4. Select the user account that you want to edit, and then click **Edit**.
5. Modify one or more user login methods, and then click **Modify**.

Locking or unlocking SVM user accounts

You can use System Manager to lock or unlock storage virtual machine (SVM) user accounts.

Steps

1. Click **Storage > SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **SVM User Details** pane, click **Users**.

4. In the **Users** window, select the user account for which you want to modify the account status, and then click either **Lock** or **Unlock**, as required.

Users window

You can use the Users window to manage user accounts, to reset the password of a user, and to view information about all of the user accounts.

Command buttons

- **Add**

Opens the Add User dialog box, which enables you to add user accounts.

- **Edit**

Opens the Modify User dialog box, which enables you to modify user login methods.



It is a best practice to use a single role for all of the access and authentication methods of a user account.

- **Delete**

Enables you to delete a selected user account.

- **Change Password**

Opens the Change Password dialog box, which enables you to reset a selected user's password.

- **Lock**

Locks the user account.

- **Refresh**

Updates the information in the window.

Users list

The area below the users list displays detailed information about the selected user.

- **User**

Displays the name of the user account.

- **Account Locked**

Displays whether the user account is locked.

User Login Methods area

- **Application**

Displays the access method that a user can use to access the storage system. The supported access methods include the following:

- System console (console)
- HTTP(S) (http)
- ONTAP API (ontapi)
- Service Processor (service-processor)
- SSH (ssh)

- **Authentication**

Displays the default supported authentication method, which is “password”.

- **Role**

Displays the role of a selected user.

Roles

You can use System Manager to create and manage roles.

Related information

[Administrator authentication and RBAC](#)

Add roles

You can use System Manager to add an access-control role and to specify the command or command directory that the users of the role can access. You can also control the level of access the role has to the command or command directory, and you can specify a query that applies to the command or command directory.

Steps

1. Click **Storage > SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **SVM User Details** pane, click **Roles**.
4. Click **Add**.
5. In the **Add Role** dialog box, specify the role name, and then add the role attributes.
6. Click **Add**.

Editing roles

You can use System Manager to modify the access of an access-control role to a command or command directory and to restrict a user’s access to only a specified set of commands.

Steps

1. Click **Storage > SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **SVM User Details** pane, click **Roles**.
4. Select the role that you want to modify, and then click **Edit**.
5. Modify the role attributes, and then click **Modify**.

Roles window

You can use the Roles window to manage the roles that are associated with user accounts.

Command buttons

- **Add**

Opens the Add Role dialog box, which enables you to create an access-control role and specify the command or command directory that the role's users can access.

- **Edit**

Opens the Edit Role dialog box, which enables you to add or modify role attributes.

- **Refresh**

Updates the information in the window.

Roles list

The roles list provides a list of roles that are available to be assigned to users.

Role Attributes area

The details area displays the role attributes, such as the command or command directory that the selected role can access, the access level, and the query that applies to the command or command directory.

UNIX

You can use System Manager to maintain a list of local UNIX users and groups for each storage virtual machine (SVM).

UNIX window

You can use the UNIX window to maintain a list of local UNIX users and groups for each storage virtual machine (SVM). You can use local UNIX users and groups for authentication and name mappings.

Groups tab

You can use the Groups tab to add, edit, or delete UNIX groups that are local to an SVM.

Command buttons

- **Add Group**

Opens the Add Group dialog box, which enables you to create UNIX groups that are local to SVMs. Local UNIX groups are used with local UNIX users.

- **Edit**

Opens the Edit Group dialog box, which enables you to edit a group ID.

- **Delete**

Deletes the selected group.

- **Refresh**

Updates the information in the window.

Groups list

- **Group Name**

Displays the name of the group.

- **Group ID**

Displays the ID of the group.

Users tab

You can use the **Users** tab to add, edit, and delete UNIX users that are local to SVMs.

Command buttons

- **Add User**

Opens the Add User dialog box, which enables you to create UNIX users that are local to SVMs.

- **Edit**

Opens the Edit User dialog box, which enables you to edit the User ID, UNIX group to which the user belongs, and the full name of the user.

- **Delete**

Deletes the selected user.

- **Refresh**

Updates the information in the window.

Users list

- **User Name**

Displays the name of the user.

- **User ID**

Displays the ID of the user.

- **Full Name**

Displays the full name of the user.

- **Primary Group ID**

Displays the ID of the group to which the user belongs.

- **Primary Group Name**

Displays the name of the group to which the user belongs.

Windows

You can use System Manager to create and manage Windows groups and user accounts.

Related information

[SMB/CIFS management](#)

Create a local Windows group

You can use System Manager to create local Windows groups that can be used for authorizing access to the data contained in the storage virtual machine (SVM) over an SMB connection. You can also assign the privileges that define the user rights or capabilities that a member of the group has when performing administrative activities.

Before you begin

CIFS server must be configured for the SVM.

About this task

- You can specify a group name with or without the local domain name.

The local domain is the name of the CIFS server for the SVM. For example, if the CIFS server name of the SVM is “CIFS_SERVER” and you want to create an “engineering” group, you can specify either “engineering” or “CIFS_SERVER\engineering” as the group name.

The following rules apply when using a local domain as part of the group name:

- You can specify only the local domain name for the SVM to which the group is applied.

For example, if the local CIFS server name is “CIFS_SERVER”, you cannot specify

“CORP_SERVER\group1” as the group name.

- You cannot use “BUILTIN” as a local domain in the group name.

For example, you cannot create a group with “BUILTIN\group1” as the name.

- You cannot use an Active Directory domain as a local domain in the group name.

For example, you cannot create a group named “AD_DOM\group1”, where “AD_DOM” is the name of an Active Directory domain.

- You cannot use a group name that already exists.
- The group name that you specify must meet the following requirements:
 - Must not exceed 256 characters
 - Must not end in a period
 - Must not include commas
 - Must not include any of the following printable characters: " / \ [] : | < > + = ; ? * @
 - Must not include characters in the ASCII range 1 through 31, which are non-printable

Steps

1. Click **Storage > SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **Host Users and Groups** pane, click **Windows**.
4. In the **Groups** tab, click **Create**.
5. In the **Create Group** dialog box, specify a name for the group and a description that helps you to identify the new group.
6. Assign a set of privileges to the group.

You can select the privileges from the predefined set of supported privileges.

7. Click **Add** to add users to the group.
8. In the **Add Members to Group** dialog box, perform one of the following actions:
 - Specify the Active Directory user or Active Directory group to be added to a particular local group.
 - Select the users from the list of available local users in the SVM.
 - Click **OK**.
9. Click **Create**.

Results

The local Windows group is created and is listed in the Groups window.

Related information

[Windows window](#)

Editing local Windows group properties

You can manage local group memberships by adding and removing a local user, an Active

Directory user, or an Active Directory group by using System Manager. You can modify the privileges that are assigned to a group and the description of a group to easily identify the group.

About this task

You must keep the following in mind when adding members to or removing members from a local Windows group:

- You cannot add users to or remove users from the special *Everyone* group.
- You cannot add a local Windows group to another local Windows group.

Steps

1. Click **Storage > SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **Host Users and Groups** pane, click **Windows**.
4. In the **Groups** tab, click **Edit**.
5. Specify a name for the group and a description to identify the new group.
6. Assign a set of privileges to the group.

You can select the privileges from the predefined set of supported privileges.

7. Click **Add** to add users to the group.
8. In the **Add Members** window, perform one of the following actions:
 - Specify the Active Directory user or Active Directory group to be added to a particular local group.
 - Select the users from the list of available local users in the storage virtual machine (SVM).
9. Click **Edit**.

Results

The local Windows group settings are modified, and the changes are displayed in the **Groups** tab.

Related information

[Windows window](#)

Add user accounts to a Windows local group

You can add a local user, an Active Directory user, or an Active Directory group (if you want users to have the privileges that are associated with that group) to a Windows local group by using System Manager.

Before you begin

- The group must exist before you can add a user to the group.
- The user must exist before you can add the user to a group.

About this task

You must keep the following in mind when adding members to a local Windows group:

- You cannot add users to the special *Everyone* group.
- You cannot add a local Windows group to another local Windows group.
- You cannot add a user account that contains a space in the user name by using System Manager.

You can either rename the user account or add the user account by using the command-line interface (CLI).

Steps

1. Click **Storage > SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **Host Users and Groups** pane, click **Windows**.
4. In the **Groups** tab, select the group to which you want to add a user, and then click **Add Members**.
5. In the **Add Members** window, perform one of the following actions:
 - Specify the Active Directory user or Active Directory group to be added to a particular local group.
 - Select the users from the list of available local users in the storage virtual machine (SVM).
6. Click **OK**.

Results

The user that you added is listed in the Userstab of the **Groups** tab.

Related information

[Windows window](#)

Renaming a local Windows group

You can use System Manager to rename a local Windows group to identify the group more easily.

About this task

- The new group name must be created in the same domain as the old group name.
- The group name must meet the following requirements:
 - Must not exceed 256 characters
 - Must not end in a period
 - Must not include commas
 - Must not include any of the following printable characters: " / \ [] : | < > + = ; ? * @
 - Must not include characters in the ASCII range 1 through 31, which are non-printable

Steps

1. Click **Storage > SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **Host Users and Groups** pane, click **Windows**.
4. In the **Groups** tab, select the group that you want to rename, and then click **Rename**.
5. In the **Rename Group** window, specify a new name for the group.

Results

The local group name is changed, and the group is listed with the new name in the Groups window.

Related information

[Windows window](#)

Deleting a local Windows group

You can use System Manager to delete a local Windows group from a storage virtual machine (SVM) if the group is no longer required for determining access rights to the data contained on the SVM or for assigning SVM user rights (privileges) to group members.

About this task

- Removing a local group removes the membership records of the group.
- The file system is not altered.

Windows Security Descriptors on files and directories that refer to this group are not adjusted.

- The special “Everyone” group cannot be deleted.
- Built-in groups such as BUILTIN\Administrators and BUILTIN\Users cannot be deleted.

Steps

1. Click **Storage > SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **Host Users and Groups** pane, click **Windows**.
4. In the **Groups** tab, select the group that you want to delete, and then click **Delete**.
5. Click **Delete**.

Results

The local group is deleted along with its membership records.

Related information

[Windows window](#)

Create a local Windows user account

You can use System Manager to create a local Windows user account that can be used to authorize access to the data contained in the storage virtual machine (SVM) over an SMB connection. You can also use local Windows user accounts for authentication when creating a CIFS session.

Before you begin

- The CIFS server must be configured for the SVM.

About this task

A local Windows user name must meet the following requirements:

- Must not exceed 20 characters
- Must not end in a period
- Must not include commas
- Must not include any of the following printable characters: " / \ [] : | < > + = ; ? * @
- Must not include characters in the ASCII range 1 through 31, which are non-printable

The password must meet the following criteria:

- Must be at least six characters in length
- Must not contain the user account name
- Must contain characters from at least three of the following four categories:
 - English uppercase characters (A through Z)
 - English lowercase characters (a through z)
 - Base 10 digits (0 through 9)
 - Special characters: ~ ! @ # 0 ^ & * _ - + = ` \ | () [] : ; " ' < > , . ? /

Steps

1. Click **Storage > SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **Host Users and Groups** pane, click **Windows**.
4. In the **Users** tab, click **Create**.
5. Specify a name for the local user.
6. Specify the full name of the local user and a description that helps you to identify this new user.
7. Enter a password for the local user, and confirm the password.

The password must meet the password requirements.

8. Click **Add** to assign group memberships to the user.
9. In the **Add Groups** window, select the groups from the list of available groups in the SVM.
10. Select **Disable this account** to disable this account after the user is created.
11. Click **Create**.

Results

The local Windows user account is created and is assigned membership to the selected groups. The user account is listed in the **Users** tab.

Related information

[Windows window](#)

Editing the local Windows user properties

You can use System Manager to modify a local Windows user account if you want to change an existing user's full name or description, or if you want to enable or disable the user account. You can also modify the group memberships that are assigned to the user

account.

Steps

1. Click **Storage > SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **Host Users and Groups** pane, click **Windows**.
4. In the **Users** tab, click **Edit**.
5. In the **Modify User** window, make the required changes.
6. Click **Modify**.

Results

The attributes of the local Windows user account are modified and are displayed in the **Users** tab.

Related information

[Windows window](#)

Assigning group memberships to a user account

You can use System Manager to assign group membership to a user account if you want a user to have the privileges that are associated with a particular group.

Before you begin

- The group must exist before you can add a user to the group.
- The user must exist before you can add the user to a group.

About this task

You cannot add users to the special *Everyone* group.

Steps

1. Click **Storage > SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **Host Users and Groups** pane, click **Windows**.
4. In the **Users** tab, select the user account to which you want to assign group memberships, and then click **Add to Group**.
5. In the **Add Groups** window, select the groups to which you want to add the user account.
6. Click **OK**.

Results

The user account is assigned membership to all of the selected groups, and the user has the privileges that are associated with these groups.

Related information

[Windows window](#)

Renaming a local Windows user

You can use System Manager to rename a local Windows user account to identify the local user more easily.

About this task

- The new user name must be created in the same domain as the previous user name.
- The user name that you specify must meet the following requirements:
 - Must not exceed 20 characters
 - Must not end in a period
 - Must not include commas
 - Must not include any of the following printable characters: " / \ [] : | < > + = ; ? * @
 - Must not include characters in the ASCII range 1 through 31, which are non-printable

Steps

1. Click **Storage > SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **Host Users and Groups** pane, click **Windows**.
4. In the **Users** tab, select the user that you want to rename, and then click **Rename**.
5. In the **Rename User** window, specify a new name for the user.
6. Confirm the new name, and then click **Rename**.

Results

The user name is changed, and the new name is listed in the **Users** tab.

Related information

[Windows window](#)

Resetting the password of a Windows local user

You can use System Manager to reset the password of a Windows local user. For example, you might want to reset the password if the current password is compromised or if the user has forgotten the password.

About this task

The password that you set must meet the following criteria:

- Must be at least six characters in length
- Must not contain the user account name
- Must contain characters from at least three of the following four categories:
 - English uppercase characters (A through Z)
 - English lowercase characters (a through z)
 - Base 10 digits (0 through 9)

- Special characters: ~ ! @ # 0 ^ & * _ - + = ` \ | () [] ; : " ' < > , . ? /

Steps

1. Click **Storage > SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **Host Users and Groups** pane, click **Windows**.
4. In the **Users** tab, select the user whose password you want to reset, and then click **Set Password**.
5. In the **Reset Password** dialog box, set a new password for the user.
6. Confirm the new password, and then click **Reset**.

Related information

[Windows window](#)

Deleting a local Windows user account

You can use System Manager to delete a local Windows user account from a storage virtual machine (SVM) if the user account is no longer required for local CIFS authentication to the CIFS server of the SVM or for determining access rights to the data contained in the SVM.

About this task

- Standard users such as Administrator cannot be deleted.
- ONTAP removes references to the deleted local user from the local-group database, from the local-user-membership, and from the user-rights database.

Steps

1. Click **Storage > SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. In the **Host Users and Groups** pane, click **Windows**.
4. In the **Users** tab, select the user account that you want to delete, and then click **Delete**.
5. Click **Delete**.

Results

The local user account is deleted along with its group membership entries.

Related information

[Windows window](#)

Windows window

You can use the Windows window to maintain a list of local Windows users and groups for each storage virtual machine (SVM) on the cluster. You can use the local Windows users and groups for authentication and name mappings.

Users tab

You can use the Users tab to view the Windows users that are local to an SVM.

Command buttons

- **Create**

Opens the Create User dialog box, which enables you to create a local Windows user account that can be used to authorize access to data contained in the SVM over an SMB connection.

- **Edit**

Opens the Edit User dialog box, which enables you to edit local Windows user properties, such as group memberships and the full name. You can also enable or disable the user account.

- **Delete**

Opens the Delete User dialog box, which enables you to delete a local Windows user account from an SVM if it is no longer required.

- **Add to Group**

Opens the Add Groups dialog box, which enables you to assign group membership to a user account if you want the user to have privileges associated with that group.

- **Set Password**

Opens the Reset Password dialog box, which enables you to reset the password of a Windows local user. For example, you might want to reset the password if the password is compromised or if the user has forgotten the password.

- **Rename**

Opens the Rename User dialog box, which enables you to rename a local Windows user account to more easily identify it.

- **Refresh**

Updates the information in the window.

Users list

- **Name**

Displays the name of the local user.

- **Full Name**

Displays the full name of the local user.

- **Account Disabled**

Displays whether the local user account is enabled or disabled.

- **Description**

Displays the description for this local user.

Users Details Area

- **Group**

Displays the list of groups in which the user is a member.

Groups tab

You can use the Groups tab to add, edit, or delete Windows groups that are local to an SVM.

Command buttons

- **Create**

Opens the Create Group dialog box, which enables you to create local Windows groups that can be used for authorizing access to data contained in SVMs over an SMB connection.

- **Edit**

Opens the Edit Group dialog box, which enables you to edit the local Windows group properties, such as privileges assigned to the group and the description of the group.

- **Delete**

Opens the Delete Group dialog box, which enables you to delete a local Windows group from an SVM if it is no longer required.

- **Add Members**

Opens the Add Members dialog box, which enables you to add local or Active Directory users, or Active Directory groups to the local Windows group.

- **Rename**

Opens the Rename Group dialog box, which enables you to rename a local Windows group to more easily identify it.

- **Refresh**

Updates the information in the window.

Groups list

- **Name**

Displays the name of the local group.

- **Description**

Displays the description for this local group.

Groups Details Area

- **Privileges**

Displays the list of privileges associated with the selected group.

- **Users**

Displays the list of local users associated with the selected group.

Related information

[Creating a local Windows group](#)

[Editing local Windows group properties](#)

[Adding user accounts to a Windows local group](#)

[Renaming a local Windows group](#)

[Deleting a local Windows group](#)

[Creating a local Windows user account](#)

[Editing the local Windows user properties](#)

[Assigning group memberships to a user account](#)

[Renaming a local Windows user](#)

[Resetting the password of a Windows local user](#)

[Deleting a local Windows user account](#)

Name mapping

You can use System Manager to specify name mapping entries to map users from different platforms.

Related information

[SMB/CIFS management](#)

Name mapping conversion rules

An ONTAP system keeps a set of conversion rules for each SVM. Each rule consists of two pieces: a *pattern* and a *replacement*. Conversions start at the beginning of the appropriate list and perform a substitution based on the first matching rule. The pattern is a UNIX-style regular expression. The replacement is a string containing escape sequences representing subexpressions from the pattern, as in the UNIX `sed` program.

Name Mapping window

You can use the Name Mapping window to specify the name mapping entries to map users from different platforms.

Name Mappings

You can create and use name mappings to map your UNIX users to Windows users, Windows users to UNIX users, or Kerberos users to UNIX users.

Command buttons

- **Add**

Opens the Add Name Mapping Entry dialog box, which enables you to create a name mapping on storage virtual machines (SVMs).

- **Edit**

Opens the Edit Name Mapping Entry dialog box, which enables you to edit a name mapping on SVMs.

- **Delete**

Opens the Delete Name Mapping Entries dialog box, which enables you to delete a name mapping entry.

- **Swap**

Opens the Swap Name Mapping Entries dialog box, which enables you to interchange positions of the two selected name mapping entries.

- **Refresh**

Updates the information in the window.

Name mappings list

- **Position**

Specifies the name mapping's position in the priority list. Name mappings are applied in the order in which they occur in the priority list.

- **Pattern**

Specifies the user name pattern that must be matched.

- **Replacement**

Specifies the replacement pattern for the user name.

- **Direction**

Specifies the direction of the name mapping. Possible values are `krb_unix` for a Kerberos-to-UNIX name mapping, `win_unix` for a Windows-to-UNIX name mapping, and `unix_win` for a UNIX-to-Windows name mapping.

Command buttons

- **Add**

Opens the Add Group Mapping Entry dialog box, which enables you to create a group mapping on SVMs.

- **Edit**

Opens the Edit Group Mapping Entry dialog box, which enables you to edit the group mapping on SVMs.

- **Delete**

Opens the Delete Group Mapping Entries dialog box, which enables you to delete a group mapping entry.

- **Swap**

Opens the Swap Group Mapping Entries dialog box, which enables you to interchange positions of the two selected group mapping entries.

- **Refresh**

Updates the information in the window.

Group mappings list

- **Position**

Specifies the group mapping's position in the priority list. Group mappings are applied in the order in which they occur in the priority list.

- **Pattern**

Specifies the user name pattern that must be matched.

- **Replacement**

Specifies the replacement pattern for the user names.

- **Direction**

Specifies the direction of the group mapping. Possible values are `krb_unix` for a Kerberos-to-UNIX group mapping, `win_unix` for a Windows-to-UNIX group mapping, and `unix_win` for a UNIX-to-Windows group mapping.

Copyright Information

Copyright © 2022 NetApp, Inc. All rights reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means-graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system- without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

Trademark Information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.