



Configure a cluster by using System Manager

System Manager Classic

NetApp
January 21, 2022

Table of Contents

- Configure a cluster by using System Manager 1
 - Accessing a cluster by using the ONTAP System Manager browser-based graphic interface 1
 - Setting up the cluster 2
 - Setting up the network 8
 - Setting up physical storage 9
 - Setting up logical storage 16

Configure a cluster by using System Manager

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

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

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

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

Before you begin

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

About this task

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

Steps

1. Point the web browser to the IP address of the cluster management LIF:
 - If you are using IPv4: `https://cluster-mgmt-LIF`
 - If you are using IPv6: `https://[cluster-mgmt-LIF]` Only HTTPS is supported for browser access of ONTAP System Manager.

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

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

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

- If you do not want to continue, click **Cancel**, and close the browser.
 - If you want to continue, click **OK** to navigate to the ONTAP System Manager login page.
3. Log in to ONTAP System Manager by using your cluster administrator credentials.

Related information

[Enabling SAML authentication](#)

Setting up the cluster

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

Updating the cluster name

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

Steps

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

Changing the cluster password

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

Steps

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

Editing DNS configurations

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

Steps

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

Create a cluster management logical interface

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

Steps

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

If you want to...	Then...
Specify the IP address by using a subnet	<ol style="list-style-type: none">a. Select Using a subnet.b. In the Add Details dialog box, select the subnet from which the IP address should be assigned. For an intercluster LIF, only the subnets that are associated with the selected IPspace are displayed.c. If you want to assign a specific IP address to the LIF, select Use a specific IP address, and then type the IP address. The IP address that you specify is added to the subnet if the IP address is not already present in the subnet range.d. Click OK.
Specify the IP address manually without using a subnet	<ol style="list-style-type: none">a. Select Without a subnet.b. In the Add Details dialog box, perform the following steps:<ol style="list-style-type: none">i. Specify the IP address and the network mask or prefix.ii. Optional: Specify the gateway.iii. If you do not want to use the default value for the Destination field, specify a new destination value. If you do not specify a value, the Destination field is populated with the default value based on the family of the IP address. If a route does not exist, a new route is automatically created based on the gateway and destination.c. Click OK.

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

Editing the node name

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

Steps

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

Create a node management logical interface

You can use System Manager to create a dedicated node management logical interface (LIF) for managing a particular node in a cluster. You can use this LIF to manage the system maintenance activities of the node.

Steps

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

If you want to...	Then...
Specify the IP address by using a subnet	<ol style="list-style-type: none">a. Select Using a subnet.b. In the Add Details dialog box, select the subnet from which the IP address should be assigned. For an intercluster LIF, only the subnets that are associated with the selected IPspace are displayed.c. If you want to assign a specific IP address to the LIF, select Use a specific IP address, and then type the IP address. The IP address that you specify is added to the subnet if the IP address is not already present in the subnet range.d. Click OK.

If you want to...	Then...
Specify the IP address manually without using a subnet	<p>a. Select Without a subnet.</p> <p>b. In the Add Details dialog box, perform the following steps:</p> <ul style="list-style-type: none"> i. Specify the IP address and the network mask or prefix. ii. Optional: Specify the gateway. iii. If you do not want to use the default value for the Destination field, specify a new destination value. <p>If you do not specify a value, the Destination field is populated with the default value based on the family of the IP address.</p> <p>If a route does not exist, a new route is automatically created based on the gateway and destination.</p> <p>c. Click OK.</p>

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

What to do next

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

Editing AutoSupport settings

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

Steps

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

You can add up to five email addresses for each host.

4. In the **Others** tab, select a transport protocol for delivering the email messages, and then specify the HTTP or HTTPS proxy server details.
5. Click **OK**.

Add licenses

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

Before you begin

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

About this task

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

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

- You can upload only capacity-based licenses.

The capacity-based licenses are of “json” type.

Steps

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

If you want to...	Do this...
Add a license for a specific ONTAP service	<ol style="list-style-type: none">a. Enter the software license key. You can add multiple licenses by entering the software license keys separated by commas.b. Click Add.
Add a capacity based license	<ol style="list-style-type: none">a. Click Browse, and then select the capacity based license file.b. Click Add.
Add a license for a specific ONTAP service and add a capacity-based license	<ol style="list-style-type: none">a. Enter the software license key. You can add multiple licenses by entering the software license keys separated by commas.b. Click Browse, and then select the capacity based license file.c. Click Add.

The new license is added.

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

4. Click **Close**.

Results

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

Related information

[Licenses window](#)

Setting the time zone for a cluster

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

About this task

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

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

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

Steps

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

Related information

[Date and Time window](#)

[Creating a Kerberos realm configuration](#)

Monitoring HA pairs

You can use System Manager to monitor the node status and interconnect status of all of the high-availability (HA) pairs in a cluster. You can also verify whether takeover or giveback is enabled or has occurred, and view the reasons why takeover or giveback is

not currently possible.

Steps

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

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

Related information

[High Availability window](#)

Setting up the network

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

Create IPspaces

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

About this task

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

Steps

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

Create broadcast domains

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

Steps

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

Related information

[Network window](#)

Create subnets

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

Before you begin

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

About this task

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

Steps

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

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

4. Click **Create**.

Related information

[Network window](#)

Setting up physical storage

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

Assigning disks to nodes

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

About this task

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

You must use the command-line interface instead.

Steps

1. Click **Storage > Aggregates & Disks > Disks**.
2. In the **Disks** window, select the **Inventory** tab.
3. Select the disks that you want to assign, and then click **Assign**.
4. In the **Assign Disks** dialog box, select the node to which you want to assign the disks.
5. Click **Assign**.

Zeroing spare disks

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

About this task

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

Steps

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

Related information

[Storage recommendations for creating aggregates](#)

Provisioning storage through aggregates

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

Before you begin

You must have enough spare disks to create an aggregate.

About this task

You cannot perform the following actions by using System Manager:

- Combine disks of different sizes even if there are enough spare disks of different sizes.

You can initially create an aggregate with disks of the same size and then add disks of a different size later.

- Combine disks with different checksum types.

You can initially create an aggregate with a single checksum type and add storage of a different checksum type later.

Related information

[Aggregates window](#)

[Storage Tiers window](#)

Provisioning storage by creating an aggregate based on storage recommendations

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

About this task

- You cannot create an aggregate based on storage recommendations in Cloud Volumes ONTAP, ONTAP Select, and MetroCluster configurations.
- Errors, if any, are displayed on the screen.

You can fix these errors and then create an aggregate based on the storage recommendations, or you can create an aggregate manually.

Steps

1. Create an aggregate by using one of the following methods:
 - Click **Applications & Tiers > Storage Tiers > Add Aggregate**.
 - Click **Storage > Aggregate & Disks > Aggregates > Create**.
2. Review the storage recommendations, and then click **Submit**.

The Information dialog box displays the status of the aggregates.

3. Click **Run in Background** to navigate to the **Aggregates** window.
4. Click **Refresh** to view the aggregates that are created.

Provisioning storage by creating an aggregate manually

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

Before you begin

All of the disks must be of the same size.

About this task

- If you are creating an aggregate on a four-node cluster in ONTAP Select, the mirrored aggregate option is selected by default.
- Starting with ONTAP 9.0, you can create aggregates with disk size equal to or larger than 10 TB.

- If the disk type of the aggregate disks is FSAS or MSATA, and the disk size is equal to or larger than 10 TB, then RAID-TEC is the only available RAID type.

Steps

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

- a. Specify the name of the aggregate, the disk type, and the number of disks or partitions to include in the aggregate.

The minimum hot spare rule is applied to the disk group that has the largest disk size.

- b. Modify the RAID configuration of the aggregate:
 - i. Click **Change**.
 - ii. In the Change RAID Configuration dialog box, specify the RAID type and the RAID group size.

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

- iii. Click **Save**.

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

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

4. Click **Create**.

Results

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

Provisioning storage by creating a Flash Pool aggregate manually

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

Before you begin

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

About this task

- You cannot use partitioned SSDs while creating a Flash Pool aggregate.

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

Steps

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

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

5. Click **Use Flash Pool Cache with this aggregate**.
6. Specify the cache source:

If you want to select the cache source as...	Then...
Storage pools	<ol style="list-style-type: none"> a. Select Storage pools as the Cache Source. b. Select the storage pool from which the cache can be obtained, and then specify the cache size. c. Modify the RAID type, if required.
Dedicated SSDs	<ol style="list-style-type: none"> a. Select Dedicated SSDs as the Cache Source. b. Select the SSD size and the number of SSDs to include in the aggregate. c. Modify the RAID configuration, if required: <ol style="list-style-type: none"> i. Click Change. ii. In the Change RAID Configuration dialog box, specify the RAID type and the RAID group size. iii. Click Save.

7. Click **Create**.

Results

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

Related information

Provisioning storage by creating a SnapLock aggregate manually

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

Before you begin

The SnapLock license must have been added.

About this task

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

Steps

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

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

The minimum hot spare rule is applied to the disk group that has the largest disk size.

- b. Modify the RAID configuration of the aggregate:
 - i. Click **Change**.
 - ii. In the Change RAID Configuration dialog box, specify the RAID type and the RAID group size.

Shared disks support two RAID types: RAID-DP and RAID-TEC.
 - iii. Click **Save**.
- c. Specify the SnapLock type.
- d. If you have not initialized the system ComplianceClock, select the **Initialize ComplianceClock** check box.

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



You must ensure that the current system time is correct. The ComplianceClock is set based on the system clock. Once the ComplianceClock is set, you cannot modify or stop the ComplianceClock.

e. If you want to mirror the aggregate, select the **Mirror this aggregate** check box.

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

By default, the mirroring option is disabled for SnapLock Compliance aggregates.

4. Click **Create**.

Provisioning storage by creating a FabricPool-enabled aggregate manually

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

Before you begin

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

About this task

The following object stores can be used as cloud tiers:

- StorageGRID
- Alibaba Cloud (Starting with System Manager 9.6)
- Amazon Web Services (AWS) Simple Storage Service (S3)
- Amazon Web Services (AWS) Commercial Cloud Service (C2S)
- Microsoft Azure Blob storage
- IBM Cloud
- Google Cloud



- Azure Stack, which is an on-premises Azure services, is not supported.
- If you want to use any object store as a cloud tier, other than StorageGRID, you must have the FabricPool capacity license.

Steps

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

- a. Specify the name of the aggregate, the disk type, and the number of disks or partitions to include in the aggregate.



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

The minimum hot spare rule is applied to the disk group that has the largest disk size.

- b. Modify the RAID configuration of the aggregate:
 - i. Click **Change**.
 - ii. In the Change RAID Configuration dialog box, specify the RAID type and the RAID group size.

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

- iii. Click **Save**.

4. Select the **FabricPool** checkbox, and then select a cloud tier from the list.

5. Click **Create**.

Setting up logical storage

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

Create SVMs

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

Before you begin

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

About this task

- While creating SVMs, you can perform the following tasks:
 - Create and fully configure SVMs.
 - Configure the volume type that is allowed on SVMs.
 - Create and configure SVMs with minimal network configuration.
 - Delegate the administration to SVM administrators.
- To name the SVM, you can use alphanumeric characters and the following special characters: “.” (period), “-” (hyphen), and “_” (underscore).

The SVM name should start with an alphabet or “_” (underscore) and must not contain more than 47 characters.



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

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

The language of the SVM determines the character set that is used to display file names and data for all NAS volumes in the SVM.

- You cannot use a SnapLock aggregate as the root aggregate of SVMs.

Steps

1. Click **Storage > SVMs**.
2. Click **Create**.
3. In the **Storage Virtual Machine (SVM) Setup** window, specify the following details:
 - SVM name
 - IPspace allocated to the SVM
 - Volume type allowed
 - Protocols allowed
 - SVM language
 - Security style of the root volume
 - Root aggregate The default language setting for any SVM is C.UTF-8.

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

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

+



NVMe does not allow the combination of protocols.

+ In a MetroCluster configuration, only the aggregates that are contained in the cluster are displayed.

4. Specify the DNS domain names and the name server IP addresses to configure the DNS services.

The default values are selected from the existing SVM configurations.

5. When configuring a data LIF to access data using a protocol, specify the target alias, subnets, and the number of LIFs per node.

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

You can edit the details of the portset in a particular node by selecting the node from the nodes list in the details area.

6. Enable host-side applications such as SnapDrive and SnapManager for the SVM administrator by providing the SVM credentials.

7. For protocols other than NVMe, create a new LIF for SVM management by clicking **Create a new LIF for SVM management**, and then specify the portsets and the IP address with or without a subnet for the new management LIF.

For CIFS and NFS protocols, data LIFs have management access by default. You must create a new management LIF only if required. For iSCSI and FC, a SVM management LIF is required because data protocols and management protocols cannot share the same LIF.

8. For NVMe protocol, starting with ONTAP 9.5, configure a minimum of one LIF for each node on the second page of the SVM Setup wizard: **Configure NVMe Protocol**.

You must configure at least one LIF for each node in the HA pair. You can also specify two LIFs per node. Click the settings icon to toggle between one or two LIFs configurations.

9. Click **Submit & Continue**.

The SVM is created with the specified configuration.

Results

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

What to do next

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

Configure CIFS and NFS protocols on SVMs

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

Before you begin

- The protocols that you want to configure or enable on the SVM must be licensed.

If the protocol that you want to configure is not enabled on the SVM, you can use the Edit Storage Virtual Machine window to enable the protocol for the SVM.

- You must have the Active Directory, organizational unit, and administrative account credentials for configuring the CIFS protocol.

About this task

SnapLock aggregates are not considered for automatically creating volumes.

Steps

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

If you do not retain the same data LIF configuration for both CIFS and NFS, you must specify the IP address and ports separately for CIFS and NFS.

5. Specify the IP address by choosing one of the following options:

If you want to...	Then...
Specify the IP address by using a subnet	<p>a. Select Using a subnet.</p> <p>b. In the Add Details dialog box, select the subnet from which the IP address must be assigned.</p> <p>For intercluster LIFs, only the subnets that are associated with the selected IPspace are displayed.</p> <p>c. If you want to assign a specific IP address to the interface, select Use a specific IP address, and then type the IP address.</p> <p>The IP address that you specify is added to the subnet if the IP address is not already present in the subnet range.</p> <p>d. Click OK.</p>
Specify the IP address manually without using a subnet	<p>a. Select Without a subnet.</p> <p>b. In the Add Details dialog box, perform the following steps:</p> <ul style="list-style-type: none"> i. Specify the IP address and the network mask or prefix. ii. Optional: Specify the gateway. iii. If you do not want to use the default value for the Destination field, specify a new destination value. <p>If you do not specify a destination value, the Destination field is populated with the default value based on the family of the IP address.</p> <p>If a route does not exist, a new route is automatically created based on the gateway and destination.</p> <p>c. Click OK.</p>

6. Specify a port to create a data LIF:

- a. Click **Browse**.
- b. In the **Select Network Port or Adapter** dialog box, select a port.

- c. Click **OK**.
7. Configure the CIFS server by performing the following steps:
 - a. Specify the following information to create a CIFS server:
 - CIFS server name
 - Active Directory to associate with the CIFS server
 - Organizational unit (OU) within the Active Directory domain to associate with the CIFS server

By default, this parameter is set to CN=Computers.

 - Credentials of an administrative account that has sufficient privileges to add the CIFS server to the OU
 - b. Select **Encrypt Data while accessing all shares of this SVM** to enable SMB 3.0 encryption for all of the shares of the SVM.
 - c. Provision a volume for CIFS storage when configuring the protocol by specifying the share name, size of the share, and access permissions.
 - d. Select **Encrypt Data while accessing this share** to enable SMB 3.0 encryption for a particular share.
 8. Configure NIS services:
 - a. Specify the IP addresses of the NIS servers and NIS domain name to configure NIS services on the SVM.
 - b. Select the appropriate database type for which you want to add the “nis” name service source.
 - c. Provision a volume for NFS storage by specifying the export name, size, and permission.
 9. Click **Submit & Continue**.

Results

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

Configure iSCSI protocol on SVMs

You can configure the iSCSI protocol on a storage virtual machine (SVM) to provide block-level data access by using System Manager. You can create iSCSI LIFs and portsets and then add the LIFs to the portsets. LIFs are created on the most suitable adapters and are assigned to portsets to ensure data path redundancy.

Before you begin

- The iSCSI license must be enabled on the cluster.

If the iSCSI protocol is not enabled on the SVM, you can use the Edit Storage Virtual Machine window to enable the protocol for the SVM.

- All of the nodes in the cluster must be healthy.
- Each node must have at least two data ports, and the port state must be `up`.

About this task

- You can configure the iSCSI protocol while creating the SVM or you can do so at a later time.

- SnapLock aggregates are not considered for automatically creating volumes.

Steps

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

The maximum number of characters for an alias name is 128. If you do not specify a target alias, the SVM name is used as an alias.

5. Specify the number of iSCSI LIFs that can be assigned to a single node.

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

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

If the number of LIFs that you want to assign to the node is more than two, you must assign at least one portset to each LIF.

6. Specify the network details, including the subnet details, to create iSCSI LIFs:

If you want to...	Then...
Specify the IP address by using a subnet	<ol style="list-style-type: none"> a. Select Using a subnet. b. In the Add Details dialog box, select the subnet from which the IP address must be assigned. For intercluster LIFs, only the subnets that are associated with the selected IPspace are displayed. c. If you want to assign a specific IP address to the interface, select Use a specific IP address, and then type the IP address. The IP address that you specify is added to the subnet if the IP address is not already present in the subnet range. d. Click OK.

If you want to...	Then...
Specify the IP address manually without using a subnet	<p>a. Select Without a subnet.</p> <p>b. In the Add Details dialog box, perform the following steps:</p> <ul style="list-style-type: none"> i. Specify the IP address and the network mask or prefix. ii. Optional: Specify the gateway. iii. If you do not want to use the default value for the Destination field, specify a new destination value. <p>If you do not specify a destination value, the Destination field is populated with the default value based on the family of the IP address.</p> <p>If a route does not exist, a new route is automatically created based on the gateway and destination.</p> <p>c. Click OK.</p>

7. Select the broadcast domain.

8. Select the adapter type.

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

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

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



The ifgrp port must be added in the broadcast domain.

9. Provision a LUN for iSCSI storage when configuring the iSCSI protocol by specifying the LUN size, OS type for the LUN, and host initiator details.

10. If you want to verify or modify the configuration of the automatically generated iSCSI LIFs, select **Review or Modify LIFs configuration (Advanced Settings)**.

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

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

11. Click **Submit & Continue**.

Results

The iSCSI LIFs and portsets are created with the specified configuration. The LIFs are distributed among the

portsets based on the selected portset. The iSCSI service is started if all of the LIFs are successfully created.

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

Configure FC protocol and FCoE protocol on SVMs

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

Before you begin

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

About this task

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

If the protocols are not allowed on the SVM, you can use the Edit Storage Virtual Machine window to enable the protocols for the SVM.

- SnapLock aggregates are not considered for automatically creating volumes.

Steps

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

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

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

If the number of LIFs that you want to assign to the node is more than two, you must assign at least one portset to each LIF.

6. If you want to verify or modify the automatically generated LIFs configuration, select **Review or Edit the Interface Association**.

You can modify only the LIF name and home port. You must ensure that you do not specify duplicate

entries.

7. Provision a LUN for the FC storage or FCoE storage when configuring the protocol by providing the LUN size, OS type for the LUN, and host initiator details.
8. Click **Submit & Continue**.

Results

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

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

Related information

[ONTAP 9 Documentation Center](#)

Configure NVMe protocol on SVMs

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

About this task

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

Steps

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



If there are any other protocols enabled, you must deselect these to make NVMe available to select. NVMe cannot be combined with any other protocol.

5. In the **Edit Storage Virtual Machine** pane, click on **Resource Allocation**.
6. In the **Resource Allocation** tab, you can choose not to delegate volume creation or you can select an aggregate to provision the volumes automatically.
7. Click on the **Services** tab to configure the Name Service Switch details.
8. Click **Save and Close**

The NVMe protocol is configured on the SVM. After the protocol has been configured, you can start or stop the service using **SVM Settings**

Related information

[Setting up NVMe](#)

Delegating administration to SVM administrators

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

About this task

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

Steps

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

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

3. Specify the network details, including subnet details, for creating iSCSI LIFs:

If you want to...	Then...
Specify the IP address by using a subnet	<ol style="list-style-type: none">a. Select Using a subnet.b. In the Add Details dialog box, select the subnet from which the IP address must be assigned. For intercluster LIFs, only the subnets that are associated with the selected IPspace are displayed.c. If you want to assign a specific IP address to the interface, select Use a specific IP address, and then type the IP address. The IP address that you specify is added to the subnet if the IP address is not already present in the subnet range.d. Click OK.

If you want to...	Then...
Specify the IP address manually without using a subnet	<p>a. Select Without a subnet.</p> <p>b. In the Add Details dialog box, perform the following steps:</p> <ul style="list-style-type: none"> i. Specify the IP address and the network mask or prefix. ii. Optional: Specify the gateway. iii. If you do not want to use the default value for the Destination field, specify a new destination value. <p>If you do not specify a custom value, the Destination field is populated with the default value based on the family of the IP address.</p> <p>If a route does not exist, a new route is automatically created based on the gateway and destination.</p> <p>c. Click OK.</p>

4. Specify a port for creating a data LIF:

- a. Click **Browse**.
- b. Select a port from the Select Network Port or Adapter dialog box.
- c. Click **OK**.

Results

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

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

What to do next

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



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

Create FlexVol volumes

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

Before you begin

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

If you have not configured the protocols but have installed any one of these licenses, you can create only data protection (DP) volumes.

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

Steps

1. Click **Storage > Volumes**.
2. Click **Create > Create FlexVol**.
3. Browse and select the SVM in which you want to create the volume.

The Create Volume dialog box is displayed. The dialogue box includes the following tabs:

- General
- Storage Efficiency
- SnapLock
- Quality of Service
- Protection

4. On the **General** tab, perform the following steps:
 - a. Specify a name for the FlexVol volume.
 - b. Click the **FabricPool** button to specify that the volume is a FabricPool volume.
 - c. Click **Choose** to select an aggregate.

You can select only FabricPool-enabled aggregates if the volume is a FabricPool FlexVol volume, and

you can select only non-FabricPool-enabled aggregates if the volume is a non-FabricPool FlexVol volume. If you choose an encrypted aggregate (NAE), the volume you are creating will inherit the encryption of the aggregate.

- d. Select a storage type.
- e. Specify the volume size and measurement units.
- f. Indicate how much space should be reserved for Snapshot copies.
- g. Select a space reserve option from the **Space Reserve** drop-down menu.
- h. Select the **Volume Encryption** checkbox to enable encryption for the volume. This option is available only if you have enabled the Volume Encryption license and if the corresponding platform is capable of supporting encryption.

5. On the **Storage Efficiency** tab, perform the following steps:

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

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

- b. Specify the tiering policy for the volume.
- c. Specify the size of the volume and the percentage of the total volume size that you want to reserve for Snapshot copies.

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

- d. Select **Default**, **Thin provisioned**, or **Thick provisioned** for the volume.

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



- For AFF storage systems, the value of thin provisioning is “Default”, and for other storage systems, the value of thick provisioning is “Default”.
- For FabricPool-enabled aggregates, the value of thin provisioning is “Default”.

- e. Specify whether you want to enable deduplication on the volume.

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

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

6. On the **Quality of Service** tab, perform the following steps:

- a. Select the **Manage Storage Quality of Service** checkbox if you want to enable storage QoS for the FlexVol volume to manage workload performance.
- b. 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 set the minimum throughput limit for the policy group. ◦ 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, 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...
<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 set the minimum throughput limit for the policy group. ◦ 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, 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="margin-left: 40px;">This value is case-sensitive. 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>

7. On the **Protection** tab, perform the following steps:

a. Specify whether you want to enable **Volume Protection**.

A non-FabricPool FlexGroup volume can be protected with a FabricPool FlexGroup volume.

A FabricPool FlexGroup volume can be protected with a non-FabricPool FlexGroup volume.

b. 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 for the destination volume.</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 for the destination volume.</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>

8. Click **Create**.

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

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

Related information

[Volumes window](#)

Create SnapLock volumes

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

Before you begin

- The SnapLock license must have been installed.
- The SnapLock aggregate must be online.
- For creating an encrypted volume, you must have installed the volume encryption license by using System Manager, and 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 delete a complete SnapLock Enterprise volume or a file in a SnapLock Enterprise volume; however, you cannot delete only the data within a file in a SnapLock Enterprise volume.
- You cannot delete a SnapLock Compliance volume if data is committed to the volume.
- You cannot encrypt a volume in Cloud Volumes ONTAP.
- If encryption is enabled on the source volume and if the destination cluster is running a version of ONTAP software earlier than ONTAP 9.3, then encryption is disabled on the destination volume by default.

Steps

1. Click **Storage > Volumes**.
2. Click **Create > Create FlexVol**.
3. Browse and select the storage virtual machine (SVM) in which you want to create the volume.
4. In the **Create Volume** dialog box, specify a new name if you want to change the default name of the volume.

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

5. Select the container aggregate for the volume.

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

6. Select the **Volume Encryption** checkbox to enable encryption for the volume.

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

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

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

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

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

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

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

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

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

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

- a. Specify the autocommit period.

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

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

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

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

- c. Select the default retention period.

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

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

13. Create a 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 storage QoS 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, 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...
<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, 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>If the policy group is assigned to more than one object, the maximum throughput that you specify is shared among the objects.</p>

14. Enable **Volume Protection** in the **Protection** tab to protect the volume:

15. In the **Protection** tab, 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 for the destination volume.</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 for the destination volume.</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>

16. Click **Create**.

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

Results

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

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.