Manage clusters

System Manager Classic

NetApp
January 21, 2022

Table of Contents

Managing clusters

You can use System Manager to manage clusters.

Related information

ONTAP concepts

Understanding quorum and epsilon

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

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

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

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

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

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

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

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

Dashboard window

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

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

- **Alerts and Notifications**

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

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

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

- **Cluster Overview**

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

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

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

  The Efficiency tab displays the storage efficiency savings for a cluster or node. You can view the total logical space used, total physical space used, and the overall savings. You can select a cluster or a specific node to view the storage efficiency savings. For System Manager 9.5, the space used for Snapshot copies is *not* included in the values for total logical space used, total physical space used, and overall savings. However, starting with System Manager 9.6, the space used for Snapshot copies is included in the values for total logical space used, total physical space used, and overall savings.

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

  The Protection tab displays information about cluster-wide volumes that do not have defined protection relationships. Only the FlexVol volumes and FlexGroup volumes that meet the following criteria are displayed:

  - The volumes are RW volumes and are online.

- The aggregate containing the volumes is online.
- The volumes have protection relationships and are not yet initialized. You can navigate to the Volumes window to view the volumes that do not have a defined protection relationship.

The Protection tab also displays the top five SVMs that have the highest number of volumes that do not have defined protection relationships.

- **Nodes**

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

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

  The refresh interval for the Nodes panel is 15 minutes.

- **Applications and Objects**

  You can use the Applications and Objects panel to display information about applications, clients, and files in a cluster.

  The Applications tab displays information about the top five applications of the cluster. You can view the top five applications based on either IOPS and latency (from low to high or from high to low) or capacity (from low to high or from high to low).

  You should click the specific bar chart to view more information about the application. The total space, used space, and available space are displayed for capacity, the IOPS details are displayed for IOPS, and the latency details are displayed for latency.

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

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

  This information is displayed only for CIFS and NFS protocols.

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

- **Performance**

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

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

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

Monitoring a cluster using the dashboard

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

Steps

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

Applications

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

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

> If someone adds new applications using the CLI or REST API while you are viewing the list of applications, then you will not be able to view those new applications when you scroll the list.

The following applications can be configured in System Manager:

General Applications

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

Databases

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

Virtual Infrastructure

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

Related information

ONTAP concepts

Provisioning a basic template

You can use System Manager to quickly provision basic templates for SAP HANA.

About this task

As the cluster administrator, you can provision applications by configuring a basic template. The example describes how to configure the **SAP HANA Server**.

Steps

1. Click **Applications & Tiers** > **Applications**
2. In the **Basic** tab, select the **SAP HANA Server** template.
3. In the **Database Details**section, specify the following:
   - Database name
   - Database size
   - Log size
   - Tempdb size
   - Number of server cores
   - Span HA Controller Notes
4. Click **Provision Storage**

Results

The SAP HANA Server application is provisioned.

Related information

Refer to Application Provisioning Settings for field descriptions

Storage service definitions

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

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

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

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

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

Media or node	Available storage service level
Disk	value

Media or node	Available storage service level
Virtual machine disk	value
FlexArray LUN	value
Hybrid	value
Capacity-optimized Flash	value
Solid-state drive (SSD) - non-AFF	value
Performance-optimized Flash - SSD (AFF)	extreme, performance, value

Add Microsoft SQL Server over SAN to System Manager

You can use the Enhanced tab to add an instance of Microsoft SQL Server over SAN to System Manager.

About this task

The following procedure describes how to add a **Microsoft SQL Server** instance over SAN to System Manager. You can choose SMB as the export protocol only if the cluster is licensed for CIFS, which must be configured on the storage virtual machine (SVM).

Steps

1. Click **Applications & Tiers** > **Applications**
2. In the **Enhanced** tab, click **Add**
3. Select **Microsoft SQL Server instance** from the menu.

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

   The Add Microsoft SQL Server Instance window is displayed.

4. Specify the following details:
   - Database name
   - Database size and the required ONTAP service level
   - Number of server cores
   - Log size and the required ONTAP service level
   - Provision for Tempdb

     Specify if the server should be provisioned for Tempdb.

   - Export Protocol (SMB or SAN)

     Specify SAN

   - Host operating system

- LUN format
- Host mapping

5. Click **Add Application**

Results

The Microsoft SQL Server instance over SAN is added to System Manager.

Application provisioning settings

When setting up a basic or enhanced template for a database, server, or virtual desktop, you must provide details to System Manager. After an application is provisioned, you can edit the details and specify a resizing (increased size only). This section describes the fields in each template. Only the fields that are required for provisioning or editing the settings of the specific application are displayed.

Details for Microsoft SQL Database Applications over SAN

You enter the following information to provision Microsoft SQL Database applications over SAN or edit the settings:

- **Database Name**

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

- **Database Size**

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

- **ONTAP Service Level for Database**

  Mandatory: The service level for the database.

- **Log Size**

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

- **ONTAP Service Level for Log**

  Mandatory: The service level for the log.

- **Tempdb**

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

- **Export Protocol**

  Mandatory: The export protocol is SAN

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

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

- **Span HA Controller Nodes**

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

Details for provisioning a SAP HANA database

- **Active SAP HANA Nodes**

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

- **Memory Size per HANA Node**

  The memory size of a single SAP HANA node.

- **Data Disk Size per HANA Node**

  The data disk size for each node.

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

Details for Microsoft SQL Database Applications over SMB

You enter the following information to provision Microsoft SQL Database applications over SMB or edit the settings:

- **Database Name**

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

- **Database Size**

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

- **Database Service Level**

  Mandatory: The service level for the database.

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

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

- **Log Size**

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

- **Log Service Level**

  Mandatory: The service level for the log.

- **Provision for Tempdb**

  Mandatory: Indicates whether tempdb is provisioned.

- **Export Protocol**

  Mandatory: The export protocol is SMB or SAN.

  SMB can be chosen only when the cluster is licensed for CIFS, which has been configured for the SVM.

- **Grant Access to User**

  Mandatory: The access level for the application.

- **Permission**

  Mandatory: The permission level for the application.

Details for a SQL Server Account

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

The installation account is granted `SeSecurityPrivilege`.

- **SQL Server Service Account**

  Mandatory: This is an existing domain account; specify as `domain\user`.

- **SQL Server Agent Service Account**

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

Details for Oracle Database Applications

You enter the following information to provision Oracle database applications or edit the settings:

- **Database Name**

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

- **Datafile Size**

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

- **ONTAP Service Level for Datafile**

  Mandatory: The service level for the datafile.

- **Redo Log Group Size**

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

- **ONTAP Service Level for Redo Log Group**

  Mandatory: The service level for the redo log group.

- **Archive Log Size**

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

- **ONTAP Service Level for the Archive Log**

  Mandatory: The service level for the archive group.

- **Export Protocol**

  The export protocol: SAN or NFS

- **Initiators**

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

- **Grant Access to Host**

  The host name to give the application access to.

Details for MongoDB Applications

You enter the following information to provision MongoDB applications or edit the settings:

- **Database Name**

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

- **Data Set Size**

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

- **ONTAP Service Level for Data Set**

  Mandatory: The service level for the datafile.

- **Replication Factor**

  Mandatory: The number of replications.

- **Mapping for Primary Host**

  Mandatory: The name of primary host.

- **Mapping for Replica Host 1**

  Mandatory: The name of first host replica.

- **Mapping for Replica Host 2**

  Mandatory: Name of second host replica.

Details for Virtual Desktop Applications

You enter the following information to provision virtual desktop infrastructures (VDI) or edit the settings:

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

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

- **Desktop Size**

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

- **ONTAP Service Level for Desktops**

  Mandatory: The service level for the datafile.

- **Number of Desktops**

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

  This is not used to provision the virtual machines.

- **Select Hypervisor**

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

- **Desktop Persistence**

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

  These policies can be modified manually after provisioning.

- **Datastore Prefix**

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

- **Export Protocol**

  The export protocol: SAN or NFS

- **Initiators**

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

- **Grant Access to Host**

  The host name to give the application access to.

Initiator Details

You enter the following information to set up the initiator:

- **Initiator Group**

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

- **Initiator Group Name**

  The name of the new initiator group.

- **Initiators**

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

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

- **Initiator OS Type**

  The operating system type of the new initiator group.

- **FCP Portset**

  The FCP portset that the initiator group is bound to.

Host Access Configuration

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

- **Volume Export Configuration**

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

  - Allow All

    This option implies that an export rule is created which permits read-write access to any clients.

  - Create Custom Policy

    This option allows you to specify a list of host IP addresses to receive read-write access.

  

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

- **Host IP Addresses**

  This is a comma-separated list of IP addresses.

  

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

Application Details

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

- **Type**

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

- **SVM**

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

- **Size**

  The total size of the volume.

- **Available**

  The amount of space currently available in the volume.

- **Protection**

  The type of data protection configured.

You can expand the **Components** and **Volumes** panes for performance details about space used, IOPs, and latency.

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

Editing an application

You can edit a provisioned application to increase to storage size or to manage the Snapshot copies of the application.

About this task

As the cluster administrator, after you provision an application, you can edit it to modify the storage size. You can also create, restore, or delete Snapshot copies of the application. The example procedure that follows describes how to edit a **NAS Container** application.

Steps

1. Click **Applications & Tiers** > **Applications**
2. Click on the name of the NAS container application.

   If someone adds new applications using the CLI or REST API while you are viewing the list of applications, then you will not be able to view those new applications when you scroll the list.

   The **Overview** tab of the Application Details: nas window displays the application settings.

3. Click **Edit**.

   The Edit NAS Container: nas displays the current storage size setting and the **NFS Access - Grant Access to Host** address.

4. Modify the **Storage Total Size** value.
5. In the size units field, select from the drop-down menu to specify the correct size units (Bytes, MB, GB, or TB).
6. In the **ONTAP Service Level** field, select from the drop-down menu to specify the value.
7. Click **Save**.
8. Navigate back to the **Application Details: nas** window, and select the **Snapshot Copies**tab.

   A list of Snapshot copies for this provisioned application is displayed. You can use the **Search** field to search for Snapshot copies by name.

9. Manage the Snapshot copies by performing the following tasks as necessary:

Task	Actions
Create	Click **Create** to create a new Snapshot copy.
Restore	Click the check boxes next to the Snapshot copies you want to restore, and then click **Restore**.
Delete	Click the check boxes next to the Snapshot copies you want to delete, and then click **Delete**.

Deleting an application

You can delete a provisioned application when it is no longer required.

About this task

As the cluster administrator, after you provision an application, you can delete it when you no longer require it. The example procedure that follows describes how to delete a **NAS Container** application.

Steps

1. Click **Applications & Tiers** > **Applications**
2. Click the name of the NAS container application.

   If someone adds new applications using the CLI or REST API while you are viewing the list of applications, then you will not be able to view those new applications when you scroll the list.

   The **Overview** tab of the Application Details: nas window displays the application settings.

3. Click **Delete**.

   A dialog box displays a warning message that asks you if you are sure you want to delete this application.

4. Click **Delete**.

> Any volume deleted using the Application delete operation is not placed in the recovery queue. The volume is deleted immediately.

Applications window

You can use System Manager to display a list of the applications in a storage virtual machine (SVM). The list includes detailed information about each application.

Tabs

Depending on the configuration of the cluster, System Manager displays information about applications using one of the following methods:

- **No tabs**

  Detailed information about the application, including the name, the type, storage usage, performance, and related information.

- **Two tabs**

  The display provides two tabs of information about the application.

  - **Enhanced**

    Detailed information about the application, including the name, the type, storage usage, performance, and related information.

  - **Basic**

    Basic information about the application.

List of applications

Applications for the selected SVM are displayed on the **Enhanced** tab in a list in the following ways:

- For System Manager 9.5 and earlier, up to a maximum of 32 applications are displayed in the list.
- For System Manager 9.6, the first 25 applications are displayed in the list. As you scroll to the bottom of the list, another 25 applications are added to the list. When you continue to scroll, you can continue to add 25 applications at a time to expand the list up to a maximum of 1000 applications.

List columns

The information about each application is listed on the **Enhanced** tab in the following columns.

- **Expand/collapse arrow** ▸

  Contains an arrow that you can click to expand the information to a show a detailed view or to collapse the information back to the summary view.

- **Name**

The name of the application.

- **Type**

  The application type.

- **Component**

  The component of the application.

- **ONTAP Service Level**

  The level of ONTAP service for the application.

- **Usage**

  A graphical bar that shows the percentage of usage.

- **Used**

  The amount of storage space used by the application.

- **Available**

  The amount of storage space still available for the application.

- **Size**

  The size of the application.

- **IOPs**

  The number of input and output operations per second (IOPs) for the application.

- **Latency**

  The amount of latency for the application.

Entry fields

The following fields can be used to modify the display of information:

- **SVM**

  Enables you to display a drop-down list of SVMs from which you can select the SVM that contains the applications you want to display.

- **Search field**

  Enables you to type all or part of an application name to initiate a search based on the criteria you type. Only the applications with names that match the criteria are then displayed in the list.

- **Sort by field**

  Enables you to sort the list of applications based on name, size, or type.

Action icons

The following icons on the **Enhanced** tab can be used to initiate actions:

- **Add icon** 

  Enables you to add an application to the selected SVM.

- **Filter icon** 

  Enables you to specify the type of application you want to display in your search results.

- **Display icon** 

  Enables you to switch between a list view and a card view of the application information.

Configuration update

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

Configure the administration details of an SVM

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

About this task

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

Steps

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

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

5. Specify the network details:

If you want to…	Then…
Specify the IP address by using a subnet	a. Select **Using a subnet**. 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 that IP address is not already present in the subnet range. d. Click **OK**.
Specify the IP address manually without using a subnet	a. Select **Without a subnet**. b. In the Add Details dialog box, perform the following steps: i. Specify the IP address and network mask or prefix. ii. Optional: Specify the gateway. The destination field is populated with the default value based on the family of the IP address. iii. If you do not want the default value, specify a new destination value. If a route does not exist, a new route is automatically created based on the gateway and destination. c. Click **OK**.

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

Configuration Updates window

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

Tabs

- **Nodes**

Enables you to configure details of the node.

- **SVMs**

  Enables you to configure details of the SVM.

Nodes tab

Command buttons

- **Edit Node Name**

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

- **Create Node-management LIF**

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

- **Edit AutoSupport**

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

SVMs tab

Command button

- **Configure Administration Details**

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

Related information

Creating a cluster

Setting up a network when an IP address range is disabled

Service Processors

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

Isolating management network traffic

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

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

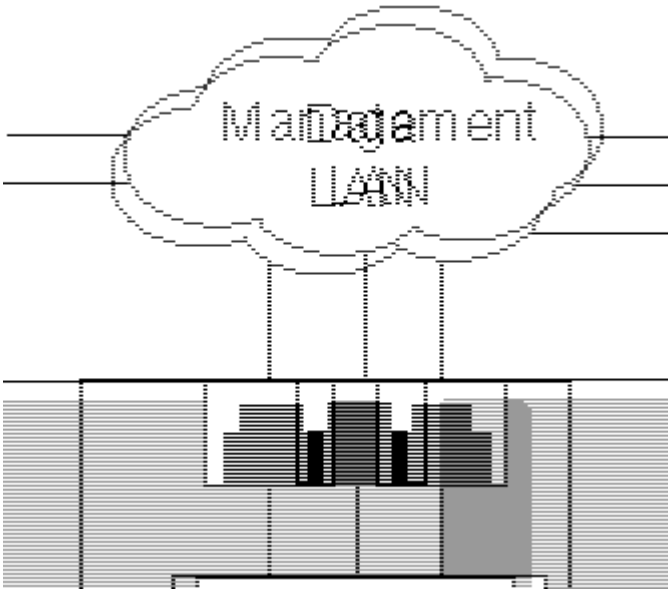

If you plan to use both the remote management device and e0M, you must configure them on the same IP subnet. Since these are low-bandwidth interfaces, the best practice is to configure SP/BMC and e0M on a subnet dedicated to management traffic.

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

Some storage controllers, such as the AFF A800, have two external ports, one for BMC and the other for e0M. For these controllers, there is no requirement to configure BMC and e0M on the same IP subnet.

Assigning IP addresses to Service Processors

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

Steps

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

If you want to…	Then…
Assign IP addresses automatically from a DHCP server	Select **DHCP**.

If you want to…	Then…
Assign IP addresses from a subnet	Select **Subnet**.
Manually provide IP addresses	Select **Manual Assignment**.

4. Click **Save**.

Editing Service Processor settings

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

About this task

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

Steps

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

Understanding the Service Processor

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

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

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

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

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

Service Processors window

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

Command buttons

- **Edit**

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

- **Global Settings**

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

- **Refresh**

  Updates the information in the window.

Service processors list

- **Node**

  Specifies the node on which the Service Processor is located.

- **IP Address**

  Specifies the IP addresses of the Service Processor.

- **Status**

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

- **MAC Address**

  Specifies the MAC address of the Service Processor.

Details area

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

Related information

Setting up a network when an IP address range is disabled

Cluster peers

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

Generating a peering passphrase

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

Steps

1. Click **Configuration** > **Cluster Peers**.
2. Click **Generate Peering Passphrase**.

   The Generate Peering Passphrase dialog window displays.

3. Complete the following fields:
   - **IPspace**: Select the IPspace from the pull-down menu.
   - **Passphrase Validity**: Select from the drop-down menu the duration for which you want the passphrase to be valid.
   - **SVM Permissions**: Select one of the following:
     - **All SVMs** to indicate all SVMs are permitted to access the cluster.
     - **Selected SVMs** to indicate specific SVMs that are permitted to access the cluster. Highlight the SVM names in the field that you want to specify.
4. Select the checkbox if the effective cluster version of the remote cluster is earlier than ONTAP 9.6. Otherwise, the cluster peering fails to generate.
5. Click **Generate** to generate the passphrase.

   For a successful generation, a message displays that identifies your passphrase.

6. If you want to email or copy the passphrase, perform one of the following actions:
   - Click **Email passphrase details**.
   - Click **Copy passphrase**.

Modifying the cluster peer passphrase

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

Steps

1. Click **Configuration** > **Cluster Peers**.
2. Select the peered cluster, and click **Edit**

   The drop-down menu displays.

3. Click **Local Cluster Passphrase**.

   The Edit Local Cluster Passphrase dialog window displays.

4. In the **Enter Passphrase** field, enter a new passphrase, and then click **Apply**.

   The minimum required length of the passphrase is eight characters.

   The passphrase is modified immediately. However, there might be a delay before the correct authentication status is displayed.

5. Log in to the remote cluster, and perform Steps 1 through 4 to modify the passphrase in the remote cluster.

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

Modifying LIFs that are configured for the remote cluster

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

Before you begin

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

Steps

1. Click **Configuration** > **Cluster** > **Configuration Updates**.
2. Select the peered cluster, and click **Edit**

   The drop-down menu displays.

3. Click **Peer Cluster Network Parameters**.

   The Edit Peer Network Parameters dialog window displays.

4. If required, modify the following fields:
   - **IPspace**: Select the IPspace from the pull-down menu.
   - **Intercluster LIFs**: Add or remove intercluster IP addresses. You can add multiple IP addresses by separating them with commas.
5. Click **Modify**.
6. Verify the changes that you made in the **Cluster Peers** window.

Changing the peering encryption status

You can use System Manager to change the peering encryption status for the selected cluster.

About this task

The encryption status can be enabled or disabled. You can change the status from enabled to disabled or from

disabled to enabled by selecting **Change Encryption**.

Steps

1. Click **Configuration** > **Cluster Peers**.
2. Select the peered cluster, and click **Edit**

   The drop-down menu displays.

3. Click **Change Encryption**.

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

   The Change Encryption dialog window displays. The toggle button indicates the current encryption status.

4. Slide the toggle button to change the peering encryption status and proceed.
   - If the current encryption status is “none”, you can enable encryption by sliding the toggle button to change the status to “tls_psk”.
   - If the current encryption status is “tls_psk”, you can disable the encryption by sliding the toggle button to change the status to “none”.
5. After you enable or disable peering encryption, you can either generate a new passphrase and provide it at the peered cluster or you can apply an existing passphrase that was already generated at the peered cluster.

   If the passphrase used on the local site does not match the passphrase used on the remote site, the cluster peering relationship will not function properly.

   Select one of the following:

   - **Generate a passphrase**: Proceed to Step #STEP_1ABAF15926174E709CA59192E200ABE3.
   - **Already have a passphrase**: Proceed to Step #STEP_2EFD822431974811AD2260C3F31DC977.
6. If you chose **Generate a passphrase**, complete the necessary fields:
   - **IPspace**: Select the IPspace from the drop-down menu.
   - **Passphrase Validity**: Select from the drop-down menu the duration for which you want the passphrase to be valid.
   - **SVM Permissions**: Select one of the following:
     - **All SVMs** to indicate that all SVMs are permitted to access the cluster.
     - **Selected SVMs** to indicate specific SVMs that are permitted to access the cluster. Highlight the SVM names in the field that you want to specify.
7. Select the checkbox if the effective cluster version of the remote cluster is earlier than ONTAP 9.6. Otherwise, the passphrase fails to generate.
8. Click **Apply**.

   The passphrase is generated for the relationship and displayed. You can either copy the passphrase or email it.

   The authentication status for the local cluster is displayed as `ok_and_offer` for the selected passphrase validity period until you provide the passphrase at the remote cluster.

9. If you already generated a new passphrase in the remote cluster, then perform the following substeps:
    a. Click **Already have a passphrase**.
    b. Enter in the **Passphrase** field the same passphrase that was generated in the remote cluster.
    c. Click **Apply**.

Deleting cluster peer relationships

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

Steps

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

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

Cluster Peers window

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

Command buttons

- **Create**

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

- **Edit**

    Displays a drop-down menu with the following choices:

    - **Local Cluster Passphrase**

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

    - **Peer Cluster Network Parameters**

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

        You can add multiple IP addresses, separated by commas.

    - **Change Encryption**

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

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

- **Delete**

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

- **Refresh**

  Updates the information in the window.

- **Manage SVM Permissions**

  Enables SVMs to automatically accept SVM peering requests.

- **Generate Peering Passphrase**

  Enables you to generate a passphrase for the local cluster IPspace by specifying the IPspace, setting the passphrase validity duration, and specifying which SVMs are given permission.

  You use the same passphrase in the remote cluster for peering.

Peer cluster list

- **Peer Cluster**

  Specifies the name of the peer cluster in the relationship.

- **Availability**

  Specifies whether the peer cluster is available for communication.

- **Authentication Status**

  Specifies whether the peer cluster is authenticated or not.

- **Local Cluster IPspace**

  Displays IPspace associated with the local cluster peer relationship.

- **Peer Cluster Intercluster IP Addresses**

  Displays IP addresses associated with the intercluster peer relationship.

- **Last Updated Time**

  Displays the time at which peer cluster was last modified.

- **Encryption**

  Displays the status of the encryption of the peering relationship.

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

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

High availability

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

Related information

ONTAP concepts

High Availability window

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

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

The color indicates the HA pair status:

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

  Green also indicates the takeover in progress state, giveback in progress state, and waiting for giveback state.

- Red: Indicates a downgraded state such as a takeover failure.
- Yellow: Indicates that the interconnect status is down.

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

Actions

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

- Takeover `node_name`

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

- Giveback `node_name`

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

- Enable or Disable automatic giveback

  Enables or disables the automatic giveback operation.

  Automatic giveback is enabled by default.

Command buttons

- **Refresh**

  Updates the information in the window.

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

Related information

Monitoring HA pairs

Licenses

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

Related information

System administration

Deleting licenses

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

Before you begin

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

Steps

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

If you want to…	Do this…
Delete a specific license package on a node or a master license	Click the **Details** tab.

If you want to…	Do this…
Delete a specific license package across all of the nodes in the cluster	Click the **Packages** tab.

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

   You can delete only one license package at a time.

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

Results

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

Related information

Licenses window

License types and entitlement risk

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

License types

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

- Node-locked license or standard license

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

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

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

- Master or site license

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

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

- Demo or temporary license

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

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

- Capacity license (ONTAP Select and FabricPool only)

  An ONTAP Select instance is licensed according to the amount of data that the user wants to manage. For example, the user might buy a 10 TB capacity license to enable ONTAP Select to manage up to 10 TB of data. If more storage capacity is attached to the system than ONTAP Select is licensed to manage, ONTAP Select will not operate. By default, the maximum storage capacity that can be attached to an ONTAP Select instance is 2 TB until a capacity license (for example, a 5 TB capacity license, a 10 TB capacity license, and so on) is purchased and installed.

  Starting with ONTAP 9.2, FabricPool-enabled aggregates require a capacity license to be used with a third-party storage tier (for example, AWS). The FabricPool capacity license defines the amount of data that can be stored in the cloud tier storage.

Entitlement risk

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

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

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

  If a site license is installed on a cluster, the entitlement risk is never high.

- Medium risk

  If a site license is not installed, and the node-locked license is non-uniformly installed on the nodes in a cluster

- No risk

  There is no entitlement risk if a node-locked license is installed on all of the nodes, or a site license is installed on the cluster, irrespective of usage.

- Unknown

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

Licenses window

Your storage system arrives from the factory with preinstalled software. If you want to add

or remove a software license after you receive the storage system, you can use the Licenses window.

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

Command buttons

- **Add**

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

- **Delete**

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

- **Refresh**

  Updates the information in the window.

Packages tab

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

- **Package**

  Displays the name of the license package.

- **Entitlement Risk**

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

- **Description**

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

License Package details area

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

Details tab

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

- **Package**

  Displays the name of the license package.

- **Cluster/Node**

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

- **Serial Number**

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

- **Type**

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

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

- **State**

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

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

- **Legacy**

  Displays whether the license is a legacy license.

- **Maximum Capacity**
  - For ONTAP Select, displays the maximum amount of storage that can be attached to the ONTAP Select instance.
  - For FabricPool, displays the maximum amount of third-party object store storage that can be used as cloud tier storage.

- **Current Capacity**
  - For ONTAP Select, displays the total amount of storage that is currently attached to the ONTAP Select instance.
  - For FabricPool, displays the total amount of third-party object store storage that is currently used as cloud tier storage.
- **Expiration Date**

  Displays the expiration date of the software license package.

Related information

Adding licenses

Deleting licenses

Creating a cluster

Cluster Expansion

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

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

Add nodes to a cluster

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

Before you begin

- New compatible nodes must be cabled to the cluster.

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

Steps

1. Add the new compatible nodes to the cluster:

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

Configure the network details of the nodes

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

Before you begin

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

Steps

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

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

What to do next

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

Related information

Creating network interfaces

Updating clusters

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

Updating clusters in a non MetroCluster configuration

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

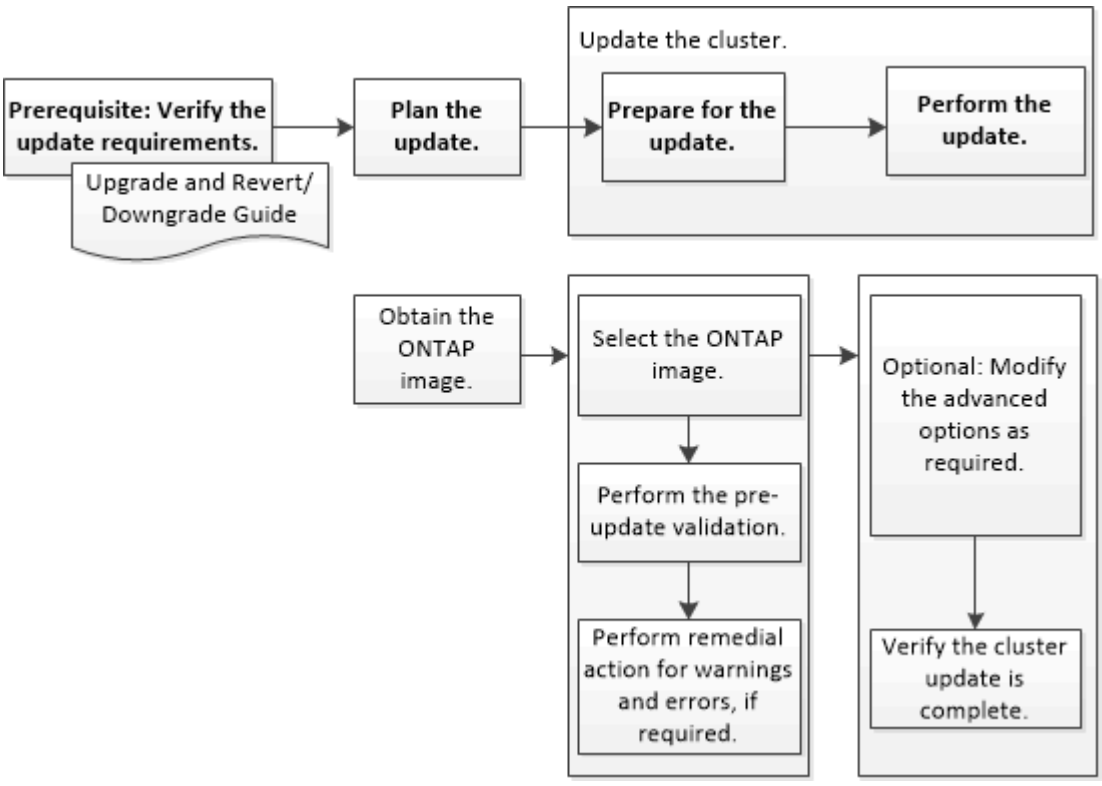

Related information

Upgrade ONTAP

Updating clusters in a MetroCluster configuration

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

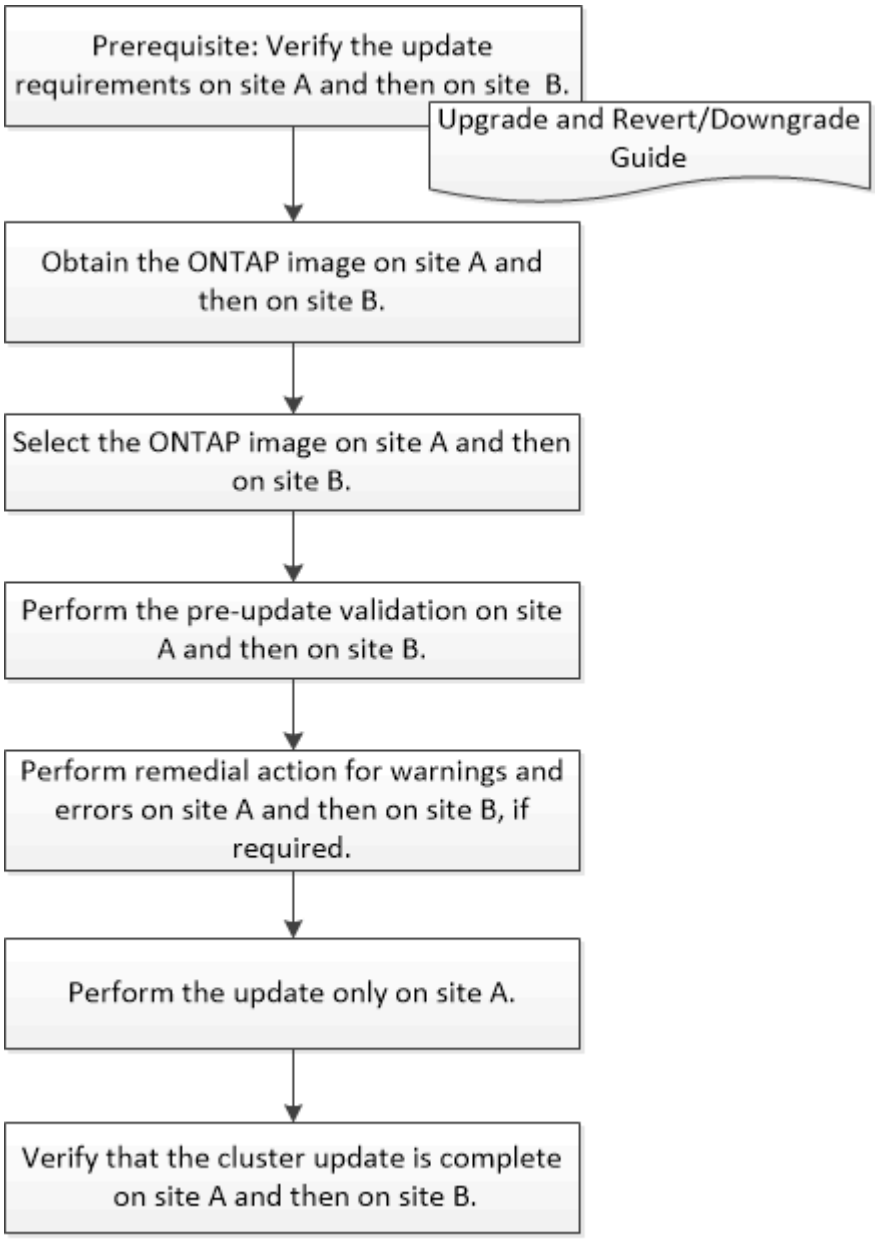

Updating site A automatically updates site B.

Related information

Upgrade ONTAP

Obtaining ONTAP software images

For ONTAP 9.4 and later, you can copy the ONTAP software image from the NetApp Support Site to a local folder. For upgrades from ONTAP 9.3 or earlier, you must copy the ONTAP software image to an HTTP server or FTP server on your network.

About this task

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

- Software images are specific to platform models.

  You must obtain the correct image for your cluster.

- Software images include the latest version of system firmware that was available when a given version of ONTAP was released.
- If you are upgrading a system with NetApp Volume Encryption to ONTAP 9.5 or later, you must download the ONTAP software image for non-restricted countries, which includes NetApp Volume Encryption.

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

- If you are upgrading from ONTAP 9.3 to 9.7, you must copy the software image for ONTAP 9.5 and 9.7.

Steps

1. Locate the target ONTAP software in the **Software Downloads** area of the NetApp Support Site.
2. Copy the software image.
   - For ONTAP 9.3 or earlier, copy the software image (for example, `93_q_image.tgz`) from the NetApp Support Site to the directory on the HTTP server or FTP server from which the image will be served
   - For ONTAP 9.4 or later, copy the software image (for example, `97_q_image.tgz`) from the NetApp Support Site to the directory on the HTTP server or FTP server from which the image will be served or to a local folder.

Updating single-node clusters

You can use System Manager to update single-node clusters.

Before you begin

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

  Obtaining ONTAP software images

About this task

- Starting with System Manager 9.5, you can update single-node clusters in two-pack MetroCluster configurations.

  You must perform this operation on both the sites.

- Updating single-node clusters in MetroCluster configurations is not disruptive.

  The System Manager user interface is not available while the cluster is rebooting.

- In System Manager 9.4 and later, you can update single-node clusters in non-MetroCluster configurations.

Updating single-node clusters in non-MetroCluster configurations is disruptive. The client data is not available while the update is in progress.

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

  You must wait for the update to finish before performing any operations.

- If the NVMe protocol is configured in System Manager 9.4 and you perform an update from System Manager 9.4 to System Manager 9.5, then the NVMe protocol is available for a grace period of 90 days without a license.

  This feature is not available in MetroCluster configurations.

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

  This feature is not available in MetroCluster configurations.

Steps

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

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

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

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

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

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

   Validation is performed again.

   - When the validation is complete, a table displays any errors and warnings, along with any required remedial actions to be taken before proceeding.
   - If the validation is completed with warnings, you can choose to select the **Continue update with warnings** checkbox, and then click **Continue**. When the validation is complete and the update is in progress, the update might be paused because of errors. You can click the error message to view the details, and then perform the remedial actions before resuming the update.

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

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

Updating a cluster nondisruptively

You can use System Manager to update a cluster or individual nodes in high-availability (HA) pairs that are running ONTAP 8.3.1 or later to a specific version of ONTAP software without disrupting access to client data.

Before you begin

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

  Obtaining ONTAP software images

About this task

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

  You must wait for the update to finish before performing any operations.

- A rolling update is performed for clusters with fewer than eight nodes, and a batch update is performed for clusters with more than eight nodes.

  In a rolling update, the nodes in the cluster are updated one at a time. In a batch update, multiple nodes are updated in parallel.

- You can nondisruptively update ONTAP software from one long-term service (LTS) release to the next LTS release (LTS+1).

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

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

  This feature is not available in MetroCluster configurations.

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

  This feature is not available in MetroCluster configurations.

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

Steps

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

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

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

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

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

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

   Validation is performed again.

   - When the validation is complete, a table displays any errors and warnings, along with any required remedial actions to be taken before proceeding.
   - If the validation is completed with warnings, you can choose to select the **Continue update with warnings** checkbox, and then click **Continue**. When the validation is complete and the update is in progress, the update might be paused because of errors. You can click the error message to view the details, and then perform the remedial actions before resuming the update.

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

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

Related information

How to update a cluster nondisruptively

Update a cluster nondisruptively

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

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

Planning and preparing for the update

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

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

Performing an update

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

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

- **Rolling update**

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

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

- **Batch update**

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

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

Related information

Updating a cluster nondisruptively

Cluster Update window

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

Tabs

- **Cluster Update**

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

- **Update History**

  Displays the details of previous cluster updates.

Cluster Update tab

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

Command buttons

- **Refresh**

  Updates the information in the window.

- **Select**

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

  - Cluster Version Details: Displays the current cluster version in use and the version details of the nodes or high-availability (HA) pairs.

- Available Software Images: Enables you to select an existing software image for the update.

  Alternatively, you can download a software image from the NetApp Support Site and add the image for the update.

- **Validate**

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

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

- **Update**

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

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

Update History tab

Displays details about the cluster update history.

Update History list

- **Image Version**

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

- **Software Updates Installed on**

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

- **Status**

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

- **Start Time**

  Specifies the time when the update was started.

- **Completion Time**

  Specifies the time when the update was completed.

  This field is hidden by default.

- **Time Taken for the Update**

Specifies the time taken for the update to finish.

- **Previous Version**

  Specifies the ONTAP version of the node before the update.

- **Updated Version**

  Specifies the ONTAP version of the node after the update.

MetroCluster switchover and switchback

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

A switchover operation allows one cluster (Site A) to take over the tasks that another cluster (Site B) usually performs. After the switchover, the cluster that has been taken over (Site B) can be brought down for maintenance and repairs. After the maintenance is completed, Site B can come up and healing tasks are completed, then you can initiate a switchback operation that allows the repaired cluster (Site B) to resume the tasks it usually performs.

System Manager supports two kinds of switchover operations, based on the status of the remote cluster site:

- A negotiated (planned) switchover: You initiate this operation when you need to do planned maintenance on a cluster or test your disaster recovery procedures.
- An unplanned switchover: You initiate this operation when a disaster has occurred on a cluster (Site B) and you want another site or cluster (Site A) to take over the tasks of the cluster affected by the disaster (Site B) while you perform repairs and maintenance.

You perform the same steps in System Manager for both switchover operations. When you initiate a switchover, System Manager determines whether the operation is feasible and aligns the workload accordingly.

MetroCluster switchover and switchback workflow

Starting with System Manager 9.6, you can use MetroCluster switchover and switchback operations after a disaster that renders all the nodes in the source cluster unreachable and powered off. You can also use the switchover workflow for a negotiated (planned) switchover in cases such as disaster recovery testing or a site going offline for maintenance.

The overall process for switchover and switchback workflow includes the following three phases:

1. **Switchover**: The switchover process allows you to transfer control of the storage and client access from a source cluster site (Site B) to another cluster site (Site A). This operation helps you provide nondisruptive operations during testing and maintenance. In addition, this process also enables you to recover from a site failure. For disaster recovery testing or planned site maintenance, you can perform a MetroCluster switchover to transfer control to a disaster recovery (DR) site (Site A). Before you start the process, at least one of the surviving site nodes must be up and running before you perform the switchover. If a switchover operation previously failed on certain nodes on the DR site, the operation can be retried on all of those nodes.

2. **Site B Operations**: After switchover is completed, System Manager completes the healing process for the MetroCluster IP configuration. Healing is a planned event, which gives you full control of each step to minimize downtime. Healing is a two-phase process that occurs on the storage and controller components to prepare the nodes at the repaired site for the switchback process. During the first phase, the process heals the aggregates by resynchronizing the mirrored plexes and then heals the root aggregates by switching them back to the disaster site.

   In the second phase, the site is made ready for the switchback process.

3. **Switchback**: After maintenance and repairs are performed on Site B, you initiate the switchback operation to return control of the storage and client access from Site A to Site B. For a successful switchback, the following conditions must exist:
   - The home nodes and storage shelves must be powered on and reachable by nodes in Site A.
   - System Manager must have successfully completed the healing phase before you can initiate the switchback operation.
   - All the aggregates in Site A should be in mirrored status and cannot be in degraded or resyncing status.
   - All previous configuration changes must be complete before performing a switchback operation. This prevents those changes from competing with the negotiated switchover or switchback operation.

MetroCluster switchover and switchback workflow flowchart

The following flowchart illustrates the phases and processes that occur when you initiate switchover and switchback operations.

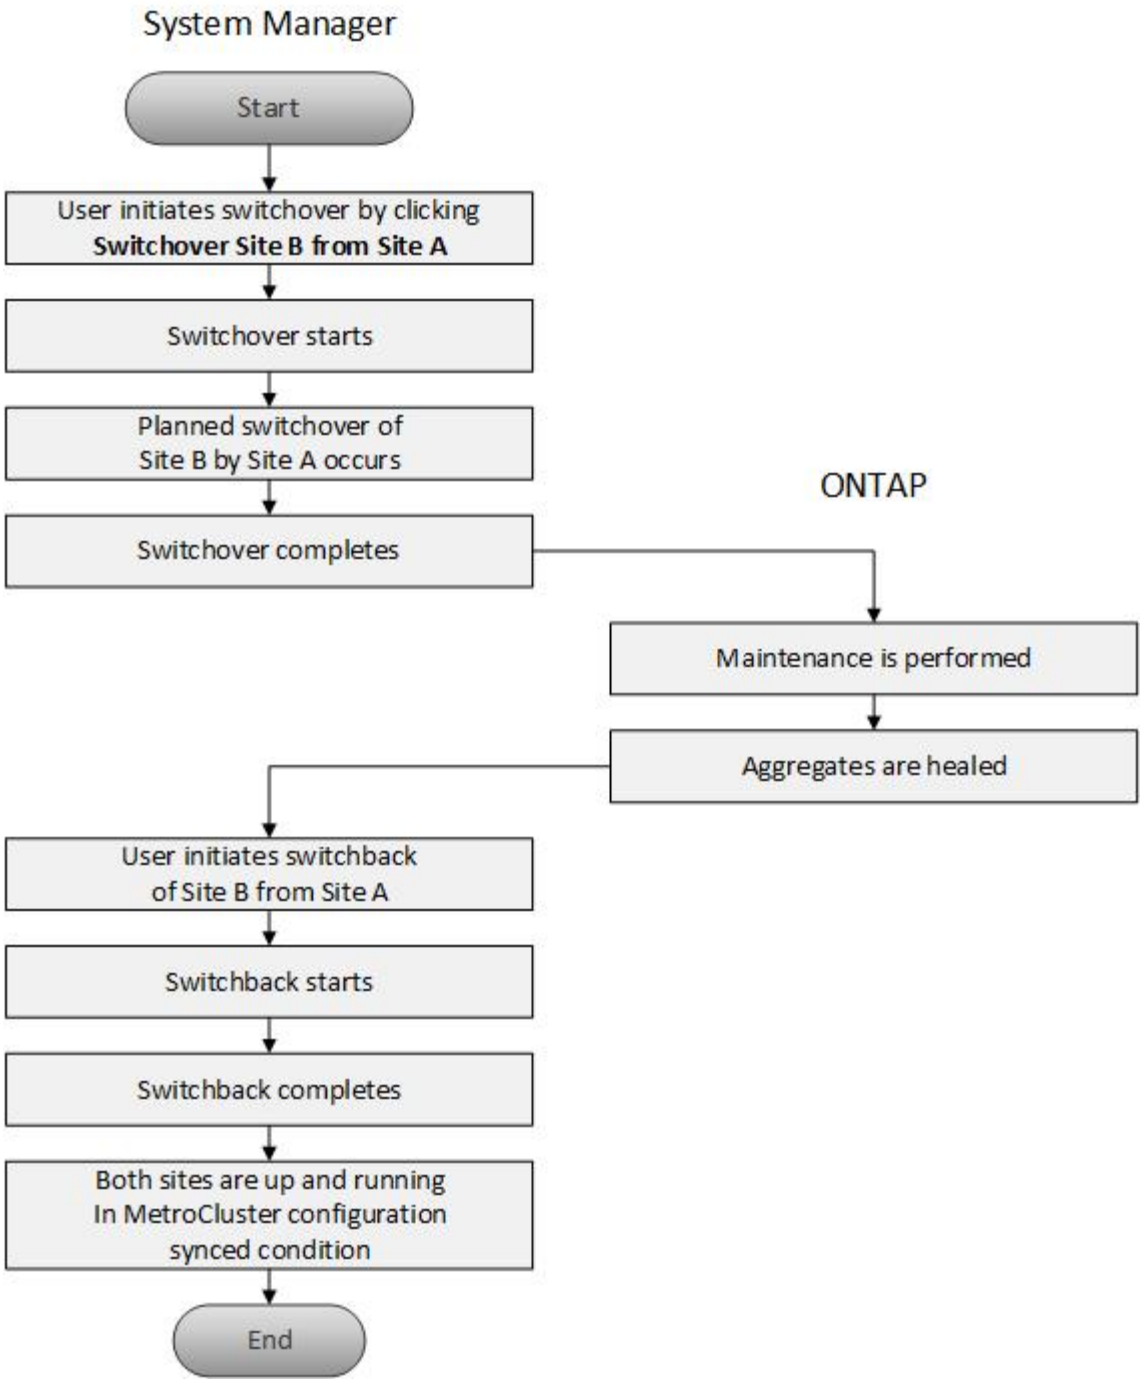

Preparing for switchover and switchback operations

Before you perform switchover operations using System Manager 9.6, you should verify that the necessary steps have been performed on the affected site.

Steps

1. If you are recovering from a disaster on Site B, you must perform the following steps:
    a. Repair or replace any damaged disks or hardware.
    b. Restore power.
    c. Rectify error issues that occur.

d. Bring up the disaster site.

2. Ensure the following conditions exist in your cluster:
   - Both sites are in Active state if you are performing a planned switchover.
   - The MetroCluster system uses configuration type “IP_Fabric”.
   - Both sites are operating with a two-node configuration (two nodes in each cluster). Sites with a single-node or four-node configuration are not supported for switchover and switchback operations using System Manager.
3. If you are launching the remote site (Site B) from the local site (Site A), ensure that Site B is running System Manager 9.6 or a later version.

Renaming the MetroCluster local site (Site A)

You can use System Manager to rename the MetroCluster local site (Site A) in a cluster.

Steps

1. Click **Configuration** > **Configuration Updates**.
2. Click **Update cluster name**.
3. Update the name in the text box, then click **Submit**.

   You can view the updated name when the MetroCluster Site A status is displayed.

4. To display the updated name of MetroCluster Site A when viewing it from the remote site (Site B), execute the following command within the CLI on the remote site (Site B): `cluster peer modify-local-name`

Performing a negotiated switchover

Starting with System Manager 9.6, you can initiate a negotiated (planned) switchover of a MetroCluster site. This operation is useful when you want to perform disaster recovery testing or planned maintenance on the site.

Steps

1. In System Manager, use the cluster administrator credentials to log on to the local MetroCluster site (Site A).
2. Click **Configuration** > **MetroCluster**

   TheMetroCluster Switchover/Switchback Operations window displays.

3. Click **Next**.

   The MetroCluster Switchover and Switchback Operations window displays the status of the operations, and System Manager verifies whether a negotiated switchover is possible.

4. Perform one of the following substeps when the validation process has completed:
   - If validation is successful, proceed to Step 5.
   - If validation fails, but Site B is up, then an error has occurred, such as a problem with a subsystem or NVram mirroring is not synchronized. You can perform either of the following processes:
     - Fix the issue that is causing the error, click **Close**, and then start again at Step 1.

        - Halt the Site B nodes, click **Close**, and then perform the steps in Performing an unplanned switchover.
    - If validation fails, and Site B is down, then most likely there is a connection problem. Verify that Site B is really down, then perform the steps in Performing an unplanned switchover.
5. Click **Switchover from Site B to Site A** to initiate the switchover process.

   A warning message displays, warning you that the switchover operation stops all data SVMs on Site B and restarts them on Site A.

6. If you want to proceed, click **Yes**.

   The switchover process begins. The states of Site A and Site B are displayed above the graphic representations of their configurations. If the switchover operation fails, an error message displays. Click **Close**. Correct any errors and start again at Step1

7. Wait until System Manager shows that healing has been completed.

   When healing is completed, Site B is operational, and systems prepare for the switchback process.

   When the preparations for the switchback process are complete, the **Switchback from Site A to Site B** button is active at the bottom of the window.

8. To proceed with the switchback operation, perform the steps in Performing a switchback.

Performing a unplanned switchover

Starting with System Manager 9.6, you can initiate an unplanned switchover of a MetroCluster site. This operation is useful after an outage event or disaster event.

Before you begin

Your MetroCluster is running in normal operating condition; however, the nodes in the local cluster (Site A) are up, but the nodes in the remote cluster (Site B) are down.

Steps

1. Verify that Site B is actually down.

   A connection error might make Site B appear to be down.

   Starting the switchover process with Site B up could cause disastrous results.

2. In System Manager, log on to the local MetroCluster site (Site A) using the cluster administrator credentials.
3. Click **Configuration** > **MetroCluster**

   TheMetroCluster Switchover/Switchback Operations window displays.

4. Click **Next**.

   The MetroCluster Switchover/Switchback Operations window displays the status of the operations, and System Manager verifies whether a negotiated switchover is feasible.

5. When the validation process is complete, click **Switchover Site B to Site A** to initiate the switchover

process.

A warning message displays, warning you that the switchover operation switches control from Site B to Site A. The status of Site B should be “UNREACHABLE”, and all Site B nodes are shown in red text.

> As noted in Step 1, Site B must actually be down and not just unconnected. Also, you should be aware that the switchover operation might cause data loss.

6. If you want to proceed, ensure the check box is checked, and then click **Yes**.

   The switchover process begins. The states of Site A and Site B are displayed above the graphic representations of their configurations. If the switchover operation fails, an error message displays. Click **Close**. Correct any errors and start again at Step1

7. Perform all required maintenance activities for Site B.
8. Ensure Site B is up.

   The healing process begins. When System Manager shows the healing is completed, Site B is operational and the systems prepare for the switchback process. The **Switchback from Site A to Site B** button appears at the bottom of the window.

9. Proceed to Performing a switchback to initiate the switchback operation.

Performing a switchback

Starting with System Manager 9.6, you can perform a switchback operation that restores control to the original MetroCluster site (Site B) after the system has completed a successful switchover operation.

Before you begin

Before you perform a switchback operation, you must complete the following tasks:

- You must prepare the MetroCluster sites by Performing a negotiated (planned) switchover or Performing an unplanned switchover.
- If errors occurred during the healing operation, you must follow the displayed instructions to fix them.
- If the state of the remote site is displayed as “Getting ready for switchback”, then the aggregates are still resynchronizing. You should wait until the status of the remote site indicates that it is ready for switchback.

About this task

If a switchover operation is successful, the MetroCluster Switchover and Switchback Operations window displays. The window shows the status of both sites and provides a message that tells you the operation was successful.

Steps

1. Click **Switchback from Site A to Site B** to initiate the switchback operation.

   A warning message tells you that the switchback operation is returning MetroCluster control to Site B and that the process might take some time.

2. If you want to proceed, click **Yes**.
3. Perform one of the following substeps when the switchback process has completed:

- If the switchback operation is successful, click **Done** to acknowledge the completion of MetroCluster operations.

  Until you acknowledge the completion of the switchback operation, System Manager continues to display a message that the operation has completed. You cannot initiate another operation or monitor subsequent switchover or switchback operations until you acknowledge the completion of the switchback operation.

- If the switchback operation is not successful, error messages display at the top of the status area. Make corrections if needed, and click **Switchback from Site A to Site B** to retry the process.

MetroCluster Switchover and Switchback Operations window

Starting with System Manager 9.6, you can use the MetroCluster Switchover and Switchback Operations window to initiate a negotiated (planned) switchover or an unplanned switchover from one site or cluster (Site B) to another site or cluster (Site A). After you perform maintenance or repairs on Site B, you can initiate a switchback from Site A to Site B and view the status of the operation in this window.

Command Buttons

- **Switchover Site B to Site A**

  Initiates the process that switches Site B over to Site A.

- **Switchback Site A to Site B**

  Initiates the process that switches Site A back to Site B.

Other actions

- **Navigate to Site B cluster**

  Enter the cluster management IP address of Site B.

- **Checkbox for unplanned switchover**

  If you want to initiate an unplanned switchover, check the box labeled **Continue with unplanned switchover**.

Status areas

As the system progresses through the process of switching over or switching back, System Manager displays status with the following methods:

- **Progress line graphic**

  Displays phases of the operations and indicates the phases that have been completed. The phases are Switchover, Site B Operations, and Switchback.

- **Show Details**

Displays a list of time-stamped system events as the MetroCluster operations progress.

- **Local: Site A**

  Displays a graphic of the configuration of the cluster at Site A, including the status of that site as it progresses through the phases of the operation.

- **Remote: Site B**

  Displays a graphic of the configuration of the cluster at Site B, including the status of that site as it progresses through the phases of the operation.

If you log in to Site B and view the MetroCluster Switchover and Switchback Operations window, then the status of Site A is shown as “INACTIVE” and the status of Site B is shown as “SWITCHOVER MODE”.

Date and time settings of a cluster

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

Related information

System administration

Date and Time window

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

Command buttons

- **Edit**

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

- **Refresh**

  Updates the information in the window.

Details area

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

Related information

Setting the time zone for a cluster

Setting up a network when an IP address range is disabled

SNMP

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

Related information

Network management

Enabling or disabling SNMP

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

Steps

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

Editing SNMP information

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

About this task

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

Steps

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

   The community name can be of 32 characters and must not contain the following special characters: , / : " ' |.

5. In the **SNMPv3**tab, do the following:
   a. Click **Add** to add an SNMPv3 user.
   b. Specify the username and modify the engine ID, if required.
   c. Select the **Authentication Protocol** and enter your credentials.
   d. Select the **Privacy Protocol** and enter your credentials.
   e. Click **OK** to save the changes.
6. Click **OK**.
7. Verify the changes that you made to the SNMP settings in the **SNMP** window.

Related information

SNMP window

Enabling or disabling SNMP traps

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

About this task

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

Steps

1. Click .
2. In the **Setup** pane, click **SNMP**.
3. In the **SNMP** window, click **Edit**.
4. In the **Edit SNMP Settings** dialog box, select the **Trap hosts** tab, and then select or clear the **Enable traps** check box to enable or disable SNMP traps, respectively.
5. If you enable SNMP traps, add the host name or IP address of the hosts to which the traps are sent.
6. Click **OK**.

Related information

SNMP window

Testing the trap host configuration

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

Steps

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

SNMP window

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

Command buttons

- **Enable/Disable**

  Enables or disables SNMP.

- **Edit**

  Opens the Edit SNMP Settings dialog box, which enables you to specify the SNMP communities for your

storage system and enable or disable traps.

- **Test Trap Host**

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

- **Refresh**

  Updates the information in the window.

Details

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

- **SNMP**

  Displays whether SNMP is enabled or not.

- **Traps**

  Displays if SNMP traps are enabled or not.

- **Location**

  Displays the address of the SNMP server.

- **Contact**

  Displays the contact details for the SNMP server.

- **Trap host IP Address**

  Displays the IP addresses of the trap host.

- **Community Names**

  Displays the community name of the SNMP server.

- **Security Names**

  Displays the security style for the SNMP server.

Related information

Editing SNMP information

Enabling or disabling SNMP traps

LDAP

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

Related information

Adding an LDAP client configuration

Deleting an LDAP client configuration

Editing an LDAP client configuration

Viewing the LDAP client configuration

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

Steps

1. Click ⚙.
2. In the **Setup** pane, click **LDAP**.

   The list of LDAP clients are displayed in the LDAP window.

Using LDAP services

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

About this task

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

LDAP window

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

Command buttons

- **Add**

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

- **Edit**

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

- **Delete**

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

- **Refresh**

  Updates the information in the window.

LDAP client list

Displays (in tabular format) details about LDAP clients.

- **LDAP Client Configuration**

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

- **Storage Virtual Machine**

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

- **Schema**

  Displays the schema for each LDAP client.

- **Minimum Bind Level**

  Displays the minimum bind level for each LDAP client.

- **Active Directory Domain**

  Displays the Active Directory domain for each LDAP client configuration.

- **LDAP Servers**

  Displays the LDAP server for each LDAP client configuration.

- **Preferred Active Directory Servers**

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

Users

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

Add a cluster user account

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

About this task

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

Steps

1. Click .

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

Editing a cluster user account

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

Steps

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

Changing passwords for cluster user accounts

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

Steps

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

Locking or unlocking cluster user accounts

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

Steps

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

User accounts (cluster administrators only)

You can create, modify, lock, unlock, or delete a cluster user account, reset a user's password, or display information about all user accounts.

You can manage cluster user accounts in the following ways:

- Creating a login method for a user by specifying the user's account name, the access method, the authentication method, and, optionally, the access-control role that the user is assigned

- Displaying users' login information, such as the account name, allowed access method, authentication method, access-control role, and account status
- Modifying the access-control role that is associated with a user's login method

> It is best to use a single role for all the access and authentication methods of a user account.

- Deleting a user's login method, such as the access method or the authentication method
- Changing the password for a user account
- Locking a user account to prevent the user from accessing the system
- Unlocking a previously locked user account to enable the user to access the system again

Roles

You can use an access-control role to control the level of access a user has to the system. In addition to using the predefined roles, you can create new access-control roles, modify them, delete them, or specify account restrictions for the users of a role.

Users window

You can use the Users window to manage user accounts, to reset the password of a user, and to view information about all of the user accounts.

Command buttons

- **Add**

  Opens the Add User dialog box, which enables you to add user accounts.

- **Edit**

  Opens the Modify User dialog box, which enables you to modify user login methods.

  > It is a best practice to use a single role for all of the access and authentication methods of a user account.

- **Delete**

  Enables you to delete a selected user account.

- **Change Password**

  Opens the Change Password dialog box, which enables you to reset a selected user's password.

- **Lock**

  Locks the user account.

- **Refresh**

  Updates the information in the window.

Users list

The area below the users list displays detailed information about the selected user.

- **User**

  Displays the name of the user account.

- **Account Locked**

  Displays whether the user account is locked.

User Login Methods area

- **Application**

  Displays the access method that a user can use to access the storage system. The supported access methods include the following:

  - System console (console)
  - HTTP(S) (http)
  - ONTAP API (ontapi)
  - Service Processor (service-processor)
  - SSH (ssh)

- **Authentication**

  Displays the default supported authentication method, which is “password”.

- **Role**

  Displays the role of a selected user.

Roles

You can use System Manager to create access-controlled user roles.

Related information

Administrator authentication and RBAC

Add roles

You can use System Manager to add an access-control role and to specify the command or command directory that users of the role can access. You can also control the level of access that the role has to the command or command directory, and you can specify a query that applies to the command or command directory.

Steps

1. Click ⚙.

2. In the **Management** pane, click **Roles**.
3. In the **Roles** window, click **Add**.
4. In the **Add Role** dialog box, type the role name and add the role attributes.
5. Click **Add**.

Editing roles

You can use System Manager to modify an access-control role's access to a command or command directory and to restrict a user's access to only a specified set of commands. You can also remove a role's access to the default command directory.

Steps

1. Click ⚙.
2. In the **Management** pane, click **Roles**.
3. In the **Roles** window, select the role that you want to modify, and then click **Edit**.
4. In the **Edit Role** dialog box, modify the role attributes, and then click **Modify**.
5. Verify the changes that you made in the **Roles** window.

Roles and permissions

The cluster administrator can restrict a user's access to only a specified set of commands by creating a restricted access-control role and then assigning the role to a user.

You can manage access-control roles in the following ways:

- By creating an access-control role, and then specifying the command or command directory that the role's users can access.
- By controlling the level of access that the role has for the command or command directory, and then specifying a query that applies to the command or command directory.
- By modifying an access-control role's access to a command or command directory.
- By displaying information about access-control roles, such as the role name, the command or command directory that a role can access, the access level, and the query.
- By deleting an access-control role.
- By restricting a user's access to only a specified set of commands.
- By displaying ONTAP APIs and their corresponding command-line interface (CLI) commands.

Roles window

You can use the Roles window to manage the roles that are associated with user accounts.

Command buttons

- **Add**

  Opens the Add Role dialog box, which enables you to create an access-control role and specify the

command or command directory that the role's users can access.

- **Edit**

  Opens the Edit Role dialog box, which enables you to add or modify role attributes.

- **Refresh**

  Updates the information in the window.

Roles list

The roles list provides a list of roles that are available to be assigned to users.

Role Attributes area

The details area displays the role attributes, such as the command or command directory that the selected role can access, the access level, and the query that applies to the command or command directory.