



Clustered Data ONTAP® 8.3

# Cluster Management Workflows for OnCommand® System Manager

February 2016 | 215-10887\_A0  
[doccomments@netapp.com](mailto:doccomments@netapp.com)

Updated for 8.3.2





# Contents

<b>Contents of the Workflow Guide .....</b>	<b>5</b>
<b>Performing a nondisruptive cluster update .....</b>	<b>7</b>
Obtaining Data ONTAP software images .....	7
Updating the cluster nondisruptively .....	8
<b>Setting up the network .....</b>	<b>11</b>
Creating VLAN interfaces .....	12
Creating interface groups .....	12
Creating broadcast domains .....	13
Creating subnets .....	13
Creating network interfaces .....	14
<b>Assigning disks or array LUNs to nodes .....</b>	<b>16</b>
Assigning disks to nodes .....	16
Zeroing spare disks .....	17
Assigning array LUNs .....	17
Zeroing spare array LUNs .....	18
<b>Assigning disks to existing aggregates .....</b>	<b>19</b>
Adding capacity disks .....	21
Provisioning cache to aggregates by adding SSDs .....	22
Creating a storage pool .....	23
<b>Provisioning cache through Flash Pool aggregates .....</b>	<b>24</b>
Provisioning cache to aggregates by adding SSDs .....	25
Increasing the cache for Flash Pool aggregates by adding SSDs .....	26
Creating a storage pool .....	26
<b>Providing data resiliency by mirroring aggregates .....</b>	<b>28</b>
Provisioning storage through aggregates .....	28
Provisioning storage by creating an aggregate .....	29
Provisioning storage by creating a Flash Pool aggregate .....	29
Mirroring aggregates .....	31
<b>Providing data access to qtrees using export policies .....</b>	<b>32</b>
Creating qtrees .....	32
Assigning export policies to qtrees .....	33
Creating an export policy .....	34
Adding rules to an export policy .....	34
<b>Configuring a Service Processor .....</b>	<b>36</b>
Assigning IP addresses to Service Processors .....	36
Editing Service Processor settings .....	37
<b>Setting up a BranchCache configuration .....</b>	<b>38</b>
Setting up CIFS .....	38
Setting up BranchCache .....	39
Editing share settings .....	40
<b>Protecting data using mirror and vault relationship .....</b>	<b>42</b>

Creating a mirror and vault relationship from a source volume .....	42
Creating a mirror and vault relationship from a destination SVM .....	44
<b>Copyright information .....</b>	<b>48</b>
<b>Trademark information .....</b>	<b>49</b>
<b>How to send comments about documentation and receive update notifications .....</b>	<b>50</b>
<b>Index .....</b>	<b>51</b>

## Contents of the Workflow Guide

---

The workflow guide provides workflows for some of the new functionality introduced in OnCommand System Manager 8.3 and later. These workflows include assigning disks or array LUNs to nodes, mirroring aggregates, protecting data using mirror and vault relationships, and updating the cluster nondisruptively.

Workflows are provided for the following functionalities that were introduced or enhanced in System Manager 8.3.2:

- [Assigning disks or array LUNs to nodes](#) on page 16  
Provides information about assigning disks and array LUNs to nodes.
- [Providing data resiliency by mirroring aggregates](#) on page 28  
Provides information about mirroring aggregates using the SyncMirror feature.
- [Protecting data using mirror and vault relationship](#) on page 42  
Provides information about protecting data using mirror and vault relationships.

Workflows are provided for the following functionalities that were introduced or enhanced in System Manager 8.3.1:

- [Performing a nondisruptive cluster update](#) on page 7  
Provides information about performing a nondisruptive cluster update.
- [Setting up the network](#) on page 11  
Provides information about configuring network components during and after cluster setup to manage and access data from your SVMs.
- [Assigning disks to existing aggregates](#) on page 19  
Provides information about adding disks (capacity or cache) to increase the size of aggregates to provide storage for one or more volumes.

Workflows are provided for the following functionalities that were introduced or enhanced in System Manager 8.3:

- [Provisioning cache through Flash Pool aggregates](#) on page 24  
Provides information about provisioning cache by creating Flash Pool aggregates or converting existing aggregates to Flash Pool aggregates by adding SSDs.
- [Providing data access to qtrees using export policies](#) on page 32  
Provides information about exporting a specific qtree on a volume by assigning export policies to it.
- [Configuring a Service Processor](#) on page 36  
Provides information about configuring a Service Processor (SP) to monitor and manage various system parameters of your storage systems.
- [Setting up the BranchCache configuration](#) on page 38  
Provides information about setting up the BranchCache configuration on a CIFS-enabled SVM and to enable BranchCache on shares to enable caching of content on computers that are local to the requesting clients.

System Manager 8.3 and later also includes the following functionalities that are covered in the workflows in *Express Guides*:

- Creating cluster peer relationships

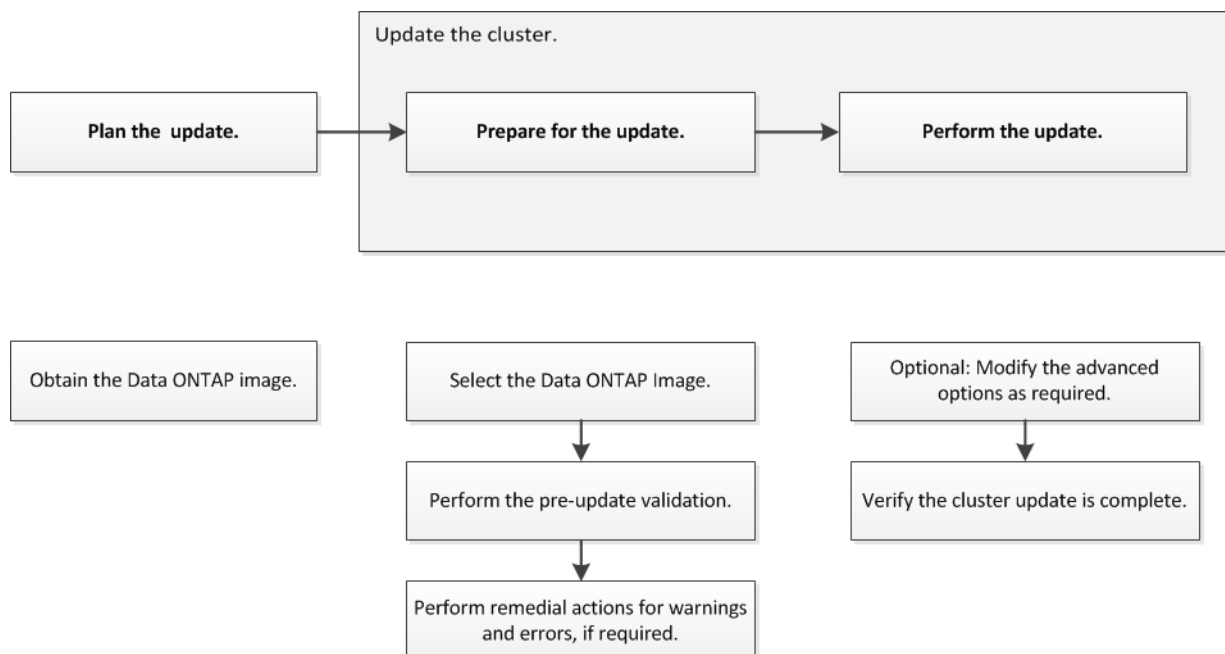
Provides information about creating authenticated cluster peer relationships among clusters to enable the clusters to communicate with each other so that you can replicate data between volumes in different clusters.

- **Creating an SVM and configuring protocols for data access**  
Provides information about creating an SVM and to configure CIFS, NFS, iSCSI, and FC/FCoE protocols on the SVM to serve data to the clients.
- **Disaster recovery using SnapMirror relationships**  
Provides information about configuring and monitoring SnapMirror relationships between peered clusters for volume disaster recovery.
- **Volume backup using SnapVault relationships**  
Provides information about configuring SnapVault relationships between volumes that are located in different clusters to back up and restore data when the data is corrupted or lost.

*[NetApp Documentation: Clustered Data ONTAP Express Guides](#)*

## Performing a nondisruptive cluster update

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



### Steps

1. [Obtaining Data ONTAP software images](#) on page 7  
You must copy a software image from the NetApp Support Site to an HTTP or FTP server on your network so that nodes can access the images.
2. [Updating the cluster nondisruptively](#) on page 8  
You can use System Manager to update a cluster or individual nodes in HA pairs running Data ONTAP 8.3.1 to a particular Data ONTAP version without disrupting access to client data.

## Obtaining Data ONTAP software images

You must copy a software image from the NetApp Support Site to an HTTP or FTP server on your network so that nodes can access the images.

### About this task

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

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

**Steps**

1. Locate the target Data ONTAP software in the **Software Downloads** area of the NetApp Support Site.
2. Copy the software image (for example, `831_q_image.tgz`) from the NetApp Support Site to the directory on the HTTP or FTP server from which the image will be served.

## Updating the cluster nondisruptively

You can use System Manager to update a cluster or individual nodes in HA pairs running Data ONTAP 8.3.1 to a particular Data ONTAP version without disrupting access to client data.

**Before you begin**

- All the nodes must be in HA pairs.  
You cannot update a single-node cluster.
- All the nodes must be healthy.
- The clusters must be running Data ONTAP 8.3.1.  
You can only update to versions later than Data ONTAP 8.3.1 by using System Manager.
- You must have copied the software image from the NetApp Support Site to an HTTP or FTP server on your network so that the nodes can access the image.  
[Obtaining Data ONTAP software images](#) on page 7

**About this task**

- If you try to perform other tasks from System Manager while updating the node hosting the cluster-management LIF, an error message might be displayed.  
You must wait for the update to finish before performing any operations.
- If the cluster consists of less than eight nodes, a rolling update is performed; if there are eight or more nodes in the cluster, a batch update is performed.  
In a rolling update, each node in the cluster is updated at a time. In a batch update, multiple nodes are updated in parallel.

**Steps**

1. Select the software image of the Data ONTAP version to which you want to update your system:
  - a. Expand the **Cluster** hierarchy in the left navigation pane.
  - b. In the navigation pane, click **Cluster Update**.
  - c. In **Cluster Update** tab, perform one of the following operations:

If you want to...	Then...
Add a new software image	<ol style="list-style-type: none"> <li>i. Click <b>Add</b>.</li> <li>ii. In the Add a New Software Image dialog box, enter the URL of the HTTP or FTP server on which you have saved the image downloaded from the NetApp Support Site.</li> <li>iii. Click <b>Add</b>.</li> </ol>



If you want to...	Then...
Select the available image	Choose one of the listed images.

2. Click **Next**.
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 provides the status of each check along with any remedial action you must perform before performing the software update, if required.

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

4. Click **Next**.
5. Optional: Click **Advanced Options** and perform the following steps:

- a. In the **Advanced Options** area, perform one of the following operations:

If you want to...	Then...
Update the entire cluster	Ensure that the <b>Update the entire cluster</b> check box is selected. By default, this check box is selected.
Update particular HA pairs	Clear the <b>Update the entire cluster</b> check box, and then select the HA pair that you want to update.

- b. Specify a different stabilization time if your environment requires more or less time for client stabilization.

Stabilization time specifies the time period for which the update process should wait after completing a task to enable client applications to recover. It should be in the range of 1 through 60 minutes; it is set to 8 minutes by default.

- c. Select the **Pause after every step (not recommended)** check box to automatically pause the update after every major step.

Pausing an update after every major step enables you to review the status of the update and then manually resume the update. This option is disabled by default and the update is not paused unless an error is encountered or you manually pause the update.

- d. Select the **Force Rolling Update** check box to perform a rolling update.

This check box is displayed only if your cluster consists of eight or more nodes.

You can enable this option if the entire cluster is selected or if there are four or more HA pairs for update.

6. Click **Update**.

Validation is performed again.

- When the validation is complete, a table is displayed, which shows the status of each check and any required remedial action that you have to perform before proceeding.
- If the validation is completed with warnings, you can select the **Continue update with warnings** check box, and then click **Continue**.

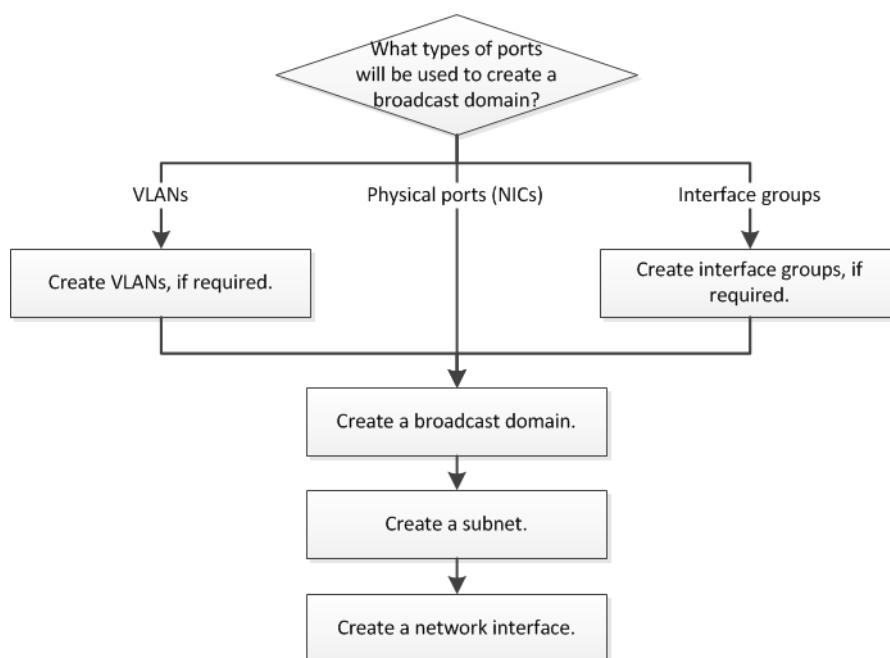
When the validation is complete and the update is in progress, the update might be paused due to errors. You can click the error to view the details and then perform the remedial actions before resuming the update.

After the update is completed successfully, you are redirected to the login page of System Manager.

7. Verify that the cluster is successfully updated to the selected version by clicking **Cluster > Cluster Update > Update History** and viewing the details.

## Setting up the network

After the IPspace is created, you can use System Manager to set up the network for accessing data from Storage Virtual Machines (SVMs) and managing SVMs. You must create a broadcast domain by using any type of port (VLANs, physical ports, or interface groups) and then create a subnet and a network interface.



### Steps

1. [Creating VLAN interfaces](#) on page 12  
You can create a VLAN for maintaining separate broadcast domains within the same network domain by using System Manager.
2. [Creating interface groups](#) on page 12  
You create an interface group—single-mode, static multimode, or dynamic multimode (LACP)—to present a single interface to clients by combining the capabilities of the aggregated network ports.
3. [Creating broadcast domains](#) on page 13  
You can create a broadcast domain by using System Manager to provide a logical division of a computer network. In a broadcast domain, all associated nodes can be reached through broadcast at the datalink layer.
4. [Creating subnets](#) on page 13  
You can create a subnet by using System Manager to provide a logical subdivision of an IP network to pre-allocate the IP addresses. A subnet enables you to create interfaces more easily by specifying a subnet instead of an IP address and network mask values for each new interface.
5. [Creating network interfaces](#) on page 14  
You can use System Manager to create a network interface or LIF to access data from Storage Virtual Machines (SVMs), manage SVMs, and to provide an interface for intercluster connectivity.

## Creating VLAN interfaces

You can create a VLAN for maintaining separate broadcast domains within the same network domain by using System Manager.

### Steps

1. Expand the **Cluster** hierarchy in the left navigation pane.
2. Select the cluster and click **Configuration > Network**.
3. In the **Network** window, select the **Ethernet Ports** tab, and click **Create VLAN**.
4. In the **Create VLAN** dialog box, select the node, the physical interface, and the broadcast domain (if required).

The physical interface list includes only Ethernet ports and interface groups. The list does not display interfaces that are in another interface group or an existing VLAN.

5. Type a VLAN tag, and click **Add**.  
You must add unique VLAN tags.
6. Click **Create**.

## Creating interface groups

You create an interface group—single-mode, static multimode, or dynamic multimode (LACP)—to present a single interface to clients by combining the capabilities of the aggregated network ports.

### Before you begin

Free ports must be available that do not belong to any broadcast domain or an interface group, or that host a VLAN.

### Steps

1. Expand the **Cluster** hierarchy in the left navigation pane.
2. Select the cluster, and then click **Configuration > Network**.
3. In the **Network** window, select the **Ethernet Ports** tab.
4. Click **Create Interface Group**.
5. In the **Create Interface Group** dialog box, specify the following settings:
  - Name of the interface group
  - Node
  - Ports that you want to include in the interface group
  - Usage mode of the ports: single, multiple, or LACP
  - Network load distribution: IP based, MAC address based, sequential, or port
  - Broadcast domain for the interface group, if required
6. Click **Create**.

## Creating broadcast domains

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

### Steps

1. Expand the **Cluster** hierarchy in the left navigation pane.
2. Select the cluster and click **Configuration > Network**.
3. In the **Network** window, select the **Broadcast Domain** tab.
4. Click **Create**.
5. In the **Create Broadcast Domain** dialog box, specify the name, MTU size, IPspace, and ports for the broadcast domain that you want to create.
6. Click **Create**.

## Creating subnets

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

### Before you begin

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

### About this task

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

### Steps

1. Expand the **Cluster** hierarchy in the left navigation pane.
2. Select the cluster, and then click **Configuration > Network**.
3. In the **Network** window, select the **Subnet** tab.
4. Click **Create**.
5. In the **Create Subnet** dialog box, specify subnet details, such as the name, subnet IP address or subnet mask, range of IP addresses, gateway address, and broadcast domain.

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

6. Click **Create**.

## Creating network interfaces

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

### Before you begin

The broadcast domain associated with the subnet must have allocated ports.

### About this task

- Dynamic DNS (DDNS) is enabled by default when a LIF is created. However, if you configure the LIF for intercluster communication, for iSCSI and FC/FCoE protocols, with IPv6 addresses, or for management access only, DDNS is disabled.
- You can specify an IP address with or without using a subnet.

### Steps

1. Expand the **Cluster** hierarchy in the left navigation pane.
2. Select the cluster, and then click **Configuration > Network**.
3. In the **Network** window, select the **Network Interfaces** tab.
4. Click **Create**.
5. In the **Create Network Interface** dialog box, specify an interface name.
6. Specify an interface role:

If you want to...	Then...
Associate the network interface with a data LIF	<ol style="list-style-type: none"> <li>a. Select <b>Serves Data</b>.</li> <li>b. Select the SVM for the network interface.</li> </ol>
Associate the network interface with an intercluster LIF	<ol style="list-style-type: none"> <li>a. Select <b>Intercluster Connectivity</b>.</li> <li>b. Select the IPspace for the network interface.</li> </ol>

7. Select the appropriate protocols.  
The interface uses the selected protocols to access data from the SVM.
8. If you want to enable management access on the data LIF, select the **Enable Management Access** check box.  
You cannot enable management access for intercluster LIFs or LIFs with FC/FCoE or iSCSI protocols.
9. Assign the IP address by choosing one of the following options:

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

**10.** Select the required ports from the ports details area.

- For data LIFs, the details area displays all the ports from the broadcast domain associated with the IPspace of the SVM.
- For intercluster LIFs, the details area displays all the ports from the broadcast domain associated with the required IPspace.

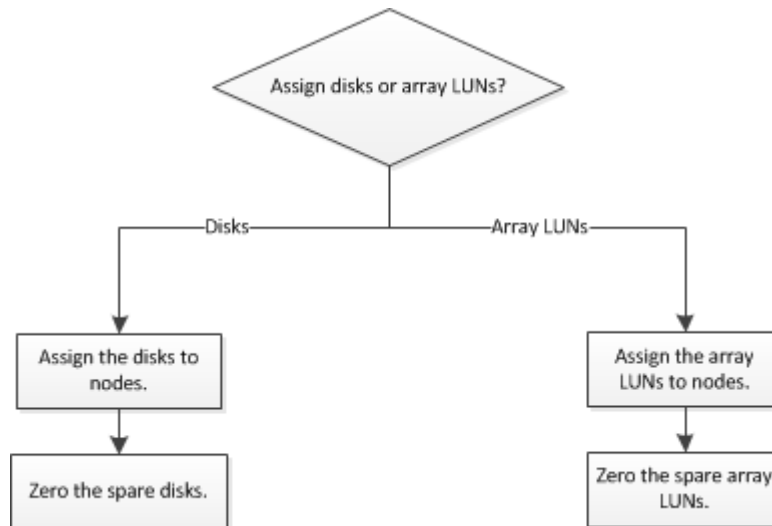
**11.** Optional: Select the **Dynamic DNS (DDNS)** check box to enable DDNS.

**12.** Click **Create**.

## Assigning disks or array LUNs to nodes

---

You can use System Manager to assign the ownership of unassigned disks and unassigned array LUNs to nodes to increase the size of the aggregates.



### Choices

- Assign disks to nodes:
  1. [Assigning disks to nodes](#) on page 16
  2. [Zeroing spare disks](#) on page 17
- Assign array LUNs to nodes:
  1. [Assigning array LUNs](#) on page 17
  2. [Zeroing spare array LUNs](#) on page 18

## Assigning disks to nodes

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

### About this task

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



**Steps**

1. Expand the **Cluster** hierarchy in the left navigation pane.
2. In the navigation pane, click **Storage > Disks**.
3. In the **Disks** window, select the **Inventory** tab.
4. Select the disks that you want to assign, and then click **Assign**.
5. In the **Assign Disks** dialog box, select the node to which you want to assign the disks.
6. Click **Assign**.

## Zeroing spare disks

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

**About this task**

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

**Steps**

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

## Assigning array LUNs

You can use System Manager to assign unassigned array LUNs to an existing aggregate to increase the size of the aggregate.

**About this task**

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

### Steps

1. Expand the **Cluster** hierarchy in the left navigation pane.
2. Click **Storage > Array LUNs**.
3. Select the array LUNs, and then click **Assign**.
4. In the **Assign Array LUNs** dialog box, select the node to which you want to assign the array LUNs.
5. Click **Assign**.

## Zeroing spare array LUNs

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

### About this task

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

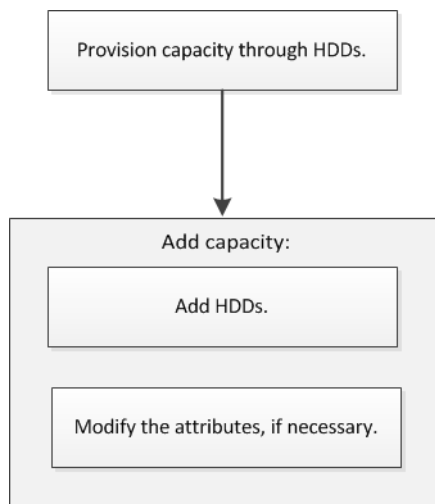
### Steps

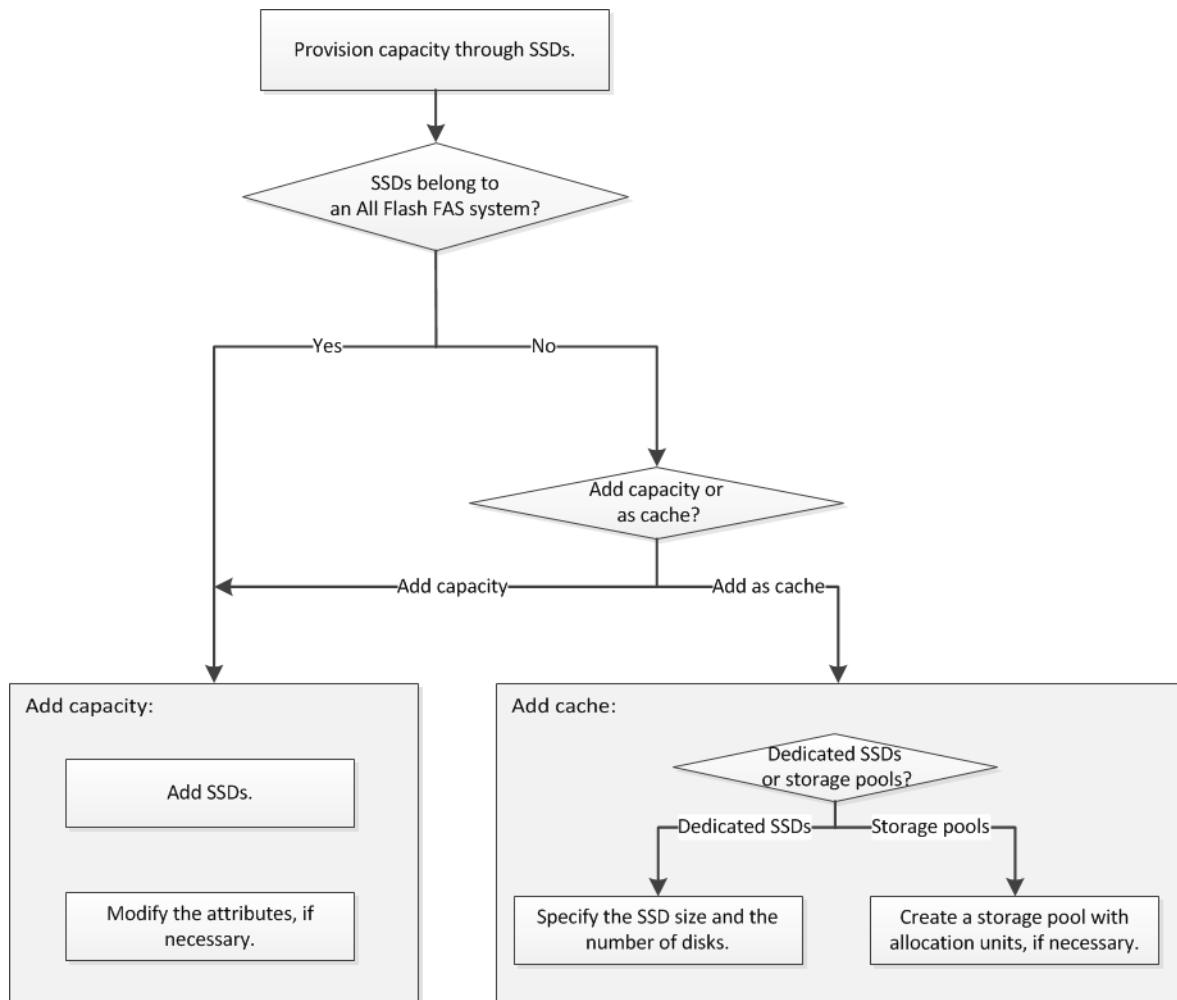
1. Expand the **Cluster** hierarchy in the left navigation pane.
2. Click **Storage > Array LUNs**.
3. Click **Zero Spares**.
4. In the **Zero Spares** dialog box, select a node or “All nodes” from which you want to zero the array LUNs.
5. Select the **Zero all non-zeroed spares** check box to confirm the zeroing operation.
6. Click **Zero Spares**.

## Assigning disks to existing aggregates

---

You can increase the size of existing aggregates by adding disks (HDDs or SSDs) or increase the cache for existing Flash Pool aggregates by adding storage pools or dedicated SSDs (cache disks).





### Steps

1. [Adding capacity disks](#) on page 21  
You can increase the size of an existing non-root aggregate or a root aggregate containing disks by adding capacity disks. You can use System Manager to add HDDs or SSDs of the selected Data ONTAP disk type and modify the RAID group options.
2. [Provisioning cache to aggregates by adding SSDs](#) on page 22  
You can use System Manager to add storage pools or dedicated SSDs to provision cache by converting an existing non-root HDD aggregate or a root aggregate that does not contain partitioned disks to a Flash Pool aggregate.
3. [Creating a storage pool](#) on page 23  
A storage pool is a collection of SSDs (cache disks). You can use System Manager to combine SSDs to create a storage pool, which enables you to share the SSDs and SSD spares between an HA pair for allocation to two or more Flash Pool aggregates at the same time.

## Adding capacity disks

You can increase the size of an existing non-root aggregate or a root aggregate containing disks by adding capacity disks. You can use System Manager to add HDDs or SSDs of the selected Data ONTAP disk type and modify the RAID group options.

### Before you begin

- The aggregate must be online.
- There must be sufficient compatible spare disks.

### About this task

- It is a best practice to add disks of the same size as the other disks in the aggregate. If you add disks that are smaller in size than the other disks in the aggregate, the aggregate becomes suboptimal in configuration, which in turn can cause performance issues. If you add disks that are larger in size than the other disks in the aggregate, the disks are downsized and their space is reduced to that of the other disks in the aggregate. If you add disks that are not of the same size as the other disks in the aggregate, the selected disks might not be added, instead other disks with a usable size between 90% and 105% of the specified size are automatically added. For a 744 GB disk, all disks in the range of 669 GB through 781 GB are eligible for selection. For all the spare disks in this range, Data ONTAP first selects only partitioned disks, then selects only unpartitioned disks, and finally selects both partitioned and unpartitioned disks.
- You cannot use System Manager to add HDDs to the following:
  - Aggregates containing only SSDs
  - Root aggregates containing partitioned disks
 You must use the command-line interface instead.
- For shared disks, RAID-DP is the only supported RAID type.

### Steps

1. Expand either the **Cluster** or the **Nodes** hierarchy in the left navigation pane.
2. In the navigation pane, click **Storage > Aggregates**.
3. In the **Aggregates** window, select the aggregate, and then click **Add Capacity**.
4. In the **Add Capacity** dialog box, perform the following steps:
  - a. Specify the disk type for the capacity disks by using **Disk Type to Add**.
  - b. Specify the number of capacity disks by using **Number of Disks**.
  - c. Specify the RAID group to which the capacity disks are added by using **Add Disks To**.  
By default, System Manager adds the capacity disks to all RAID groups.
5. Modify the RAID group by performing the following steps:
  - a. Click **Change**.
  - b. In the **RAID Group Selection** dialog box, specify the RAID group as **New RAID group** or **Specific RAID group** by using the **Add Disks To** option.

Shared disks can be added to the **New RAID group** option only.

- c. Click **Save**.
6. Click **Add**.  
For mirrored aggregates, an Add Capacity dialog box is displayed with the information that twice the number of selected disks will be added.
7. In the **Add Capacity** dialog box, click **Yes** to confirm.

### Result

The capacity disks are added to the selected aggregate and the aggregate size is increased. The updated information is displayed in the **Details** tab of the Aggregates window.

## Provisioning cache to aggregates by adding SSDs

You can use System Manager to add storage pools or dedicated SSDs to provision cache by converting an existing non-root HDD aggregate or a root aggregate that does not contain partitioned disks to a Flash Pool aggregate.

### Before you begin

- The aggregate must be online.
- There must be sufficient spare SSDs or allocation units in the storage pool that can be assigned as cache disks.
- All nodes in the cluster must be running Data ONTAP 8.3 or later.  
If the cluster is in a mixed-version state, you can use the command-line interface to create a Flash Pool aggregate and provide SSD cache.
- You must have identified a valid 64-bit non-root aggregate composed of HDDs to convert it to a Flash Pool aggregate.
- The aggregate must not contain any array LUNs.
- The aggregate must not provision storage to an Infinite Volume.
- You must be aware of platform-specific and workload-specific best practices for Flash Pool aggregate SSD tier size and configuration.

### Steps

1. Expand either the **Cluster** or the **Nodes** hierarchy in the left navigation pane.
2. In the navigation pane, click **Storage > Aggregates**.
3. In the **Aggregates** window, select the aggregate, and then click **Add Cache**.
4. In the **Add Cache** dialog box, perform the appropriate action:

If you select the cache source as...	Do this...
Storage pools	<ol style="list-style-type: none"> <li>a. Select the storage pool from which cache can be obtained.</li> <li>b. Specify the cache size.</li> <li>c. Modify the RAID type, if required.</li> </ol>

If you select the cache source as...	Do this...
Dedicated SSDs	Select the SSD size and the number of SSDs to include, and optionally modify the RAID configuration: <ol style="list-style-type: none"> <li data-bbox="667 359 846 384">a. Click <b>Change</b>.</li> <li data-bbox="667 405 1333 464">b. In the Change RAID Configuration dialog box, specify the RAID type and RAID group size, and then click <b>Save</b>.</li> </ol>

5. Click **Add**.

For mirrored aggregates, an Add Cache dialog box is displayed with the information that twice the number of selected disks will be added.

6. In the **Add Cache** dialog box, click **Yes** to confirm.

#### Result

The cache disks are added to the selected aggregate. The updated information is displayed in the **Details** tab of the Aggregates window.

## Creating a storage pool

A storage pool is a collection of SSDs (cache disks). You can use System Manager to combine SSDs to create a storage pool, which enables you to share the SSDs and SSD spares between an HA pair for allocation to two or more Flash Pool aggregates at the same time.

#### Before you begin

- Both nodes of the HA pair must be up and running in order to allocate SSDs and SSD spares through a storage pool.
- Storage pools must have a minimum of 3 SSDs.
- All SSDs in a storage pool must be owned by the same HA pair.

#### About this task

System Manager enforces the hot spare rule for SSD RAID groups when you use SSDs for adding disks to a storage pool. For example, if there are 10 SSDs in the SSD RAID group and the option `raid.min_spare_count` is set to 1 at the node level, System Manager leaves 1 SSD as the hot spare and uses the other 9 SSDs for SSD-related operations.

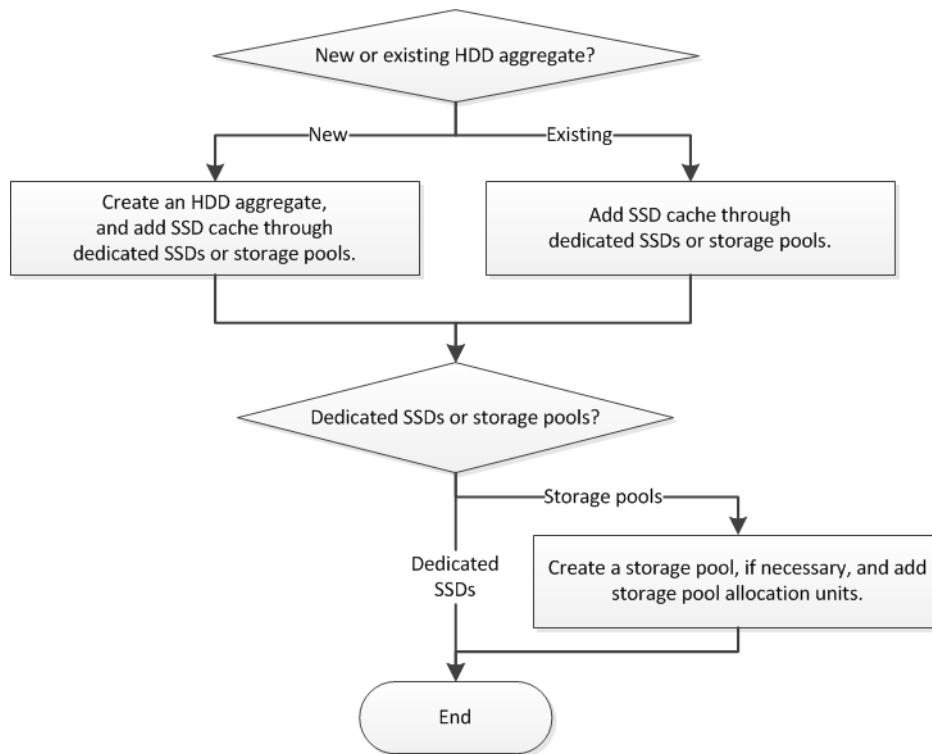
You cannot use partitioned SSDs when creating a storage pool by using System Manager.

#### Steps

1. Expand the **Cluster** hierarchy in the left navigation pane.
2. In the navigation pane, click **Storage > Storage Pools**.
3. In the **Storage Pool** window, click **Create**.
4. In the **Create Storage Pool** dialog box, specify the name for the storage pool, disk size, and the number of disks.
5. Click **Create**.

## Provisioning cache through Flash Pool aggregates

You can provision cache by creating a Flash Pool aggregate or convert an existing aggregate to a Flash Pool aggregate by adding SSDs. You can create a new HDD aggregate (if one does not exist already) and add cache through dedicated SSDs or storage pools.



### Steps

1. [Provisioning cache to aggregates by adding SSDs](#) on page 25  
You can use System Manager to add storage pools or dedicated SSDs to provision cache by converting an existing non-root HDD aggregate or a root aggregate that does not contain partitioned disks to a Flash Pool aggregate.
2. [Increasing the cache for Flash Pool aggregates by adding SSDs](#) on page 26  
You can add SSDs as either storage pools or dedicated SSDs to increase the size of a Flash Pool aggregate by using System Manager.
3. [Creating a storage pool](#) on page 26  
A storage pool is a collection of SSDs (cache disks). You can use System Manager to combine SSDs to create a storage pool, which enables you to share the SSDs and SSD spares between an HA pair for allocation to two or more Flash Pool aggregates at the same time.



## Provisioning cache to aggregates by adding SSDs

You can use System Manager to add storage pools or dedicated SSDs to provision cache by converting an existing non-root HDD aggregate or a root aggregate that does not contain partitioned disks to a Flash Pool aggregate.

### Before you begin

- The aggregate must be online.
- There must be sufficient spare SSDs or allocation units in the storage pool that can be assigned as cache disks.
- All nodes in the cluster must be running Data ONTAP 8.3 or later.  
If the cluster is in a mixed-version state, you can use the command-line interface to create a Flash Pool aggregate and provide SSD cache.
- You must have identified a valid 64-bit non-root aggregate composed of HDDs to convert it to a Flash Pool aggregate.
- The aggregate must not contain any array LUNs.
- The aggregate must not provision storage to an Infinite Volume.
- You must be aware of platform-specific and workload-specific best practices for Flash Pool aggregate SSD tier size and configuration.

### Steps

1. Expand either the **Cluster** or the **Nodes** hierarchy in the left navigation pane.
2. In the navigation pane, click **Storage > Aggregates**.
3. In the **Aggregates** window, select the aggregate, and then click **Add Cache**.
4. In the **Add Cache** dialog box, perform the appropriate action:

If you select the cache source as...	Do this...
Storage pools	<ol style="list-style-type: none"> <li>a. Select the storage pool from which cache can be obtained.</li> <li>b. Specify the cache size.</li> <li>c. Modify the RAID type, if required.</li> </ol>
Dedicated SSDs	<p>Select the SSD size and the number of SSDs to include, and optionally modify the RAID configuration:</p> <ol style="list-style-type: none"> <li>a. Click <b>Change</b>.</li> <li>b. In the Change RAID Configuration dialog box, specify the RAID type and RAID group size, and then click <b>Save</b>.</li> </ol>

5. Click **Add**.

For mirrored aggregates, an Add Cache dialog box is displayed with the information that twice the number of selected disks will be added.

6. In the **Add Cache** dialog box, click **Yes** to confirm.

**Result**

The cache disks are added to the selected aggregate. The updated information is displayed in the **Details** tab of the Aggregates window.

## Increasing the cache for Flash Pool aggregates by adding SSDs

You can add SSDs as either storage pools or dedicated SSDs to increase the size of a Flash Pool aggregate by using System Manager.

**Before you begin**

- The Flash Pool aggregate must be online.
- There must be sufficient spare SSDs or allocation units in the storage pool that can be assigned as cache disks.

**Steps**

1. Expand either the **Cluster** or the **Nodes** hierarchy in the left navigation pane.
2. In the navigation pane, click **Storage > Aggregates**.
3. In the **Aggregates** window, select the Flash Pool aggregate, and then click **Add Cache**.
4. In the **Add Cache** dialog box, perform the appropriate action:

<b>If you selected the cache source as...</b>	<b>Do this...</b>
Storage pools	Select the storage pool from which cache can be obtained and specify the cache size.
Dedicated SSDs	Select the SSD size and the number of SSDs to include.

5. Click **Add**.

For mirrored aggregates, an Add Cache dialog box is displayed with the information that twice the number of selected disks will be added.

6. In the **Add Cache** dialog box, click **Yes** to confirm.

**Result**

The cache disks are added to the selected Flash Pool aggregate. The updated information is displayed in the **Details** tab of the Aggregates window.

## Creating a storage pool

A storage pool is a collection of SSDs (cache disks). You can use System Manager to combine SSDs to create a storage pool, which enables you to share the SSDs and SSD spares between an HA pair for allocation to two or more Flash Pool aggregates at the same time.

**Before you begin**

- Both nodes of the HA pair must be up and running in order to allocate SSDs and SSD spares through a storage pool.
- Storage pools must have a minimum of 3 SSDs.

- All SSDs in a storage pool must be owned by the same HA pair.

**About this task**

System Manager enforces the hot spare rule for SSD RAID groups when you use SSDs for adding disks to a storage pool. For example, if there are 10 SSDs in the SSD RAID group and the option `raid.min_spare_count` is set to 1 at the node level, System Manager leaves 1 SSD as the hot spare and uses the other 9 SSDs for SSD-related operations.

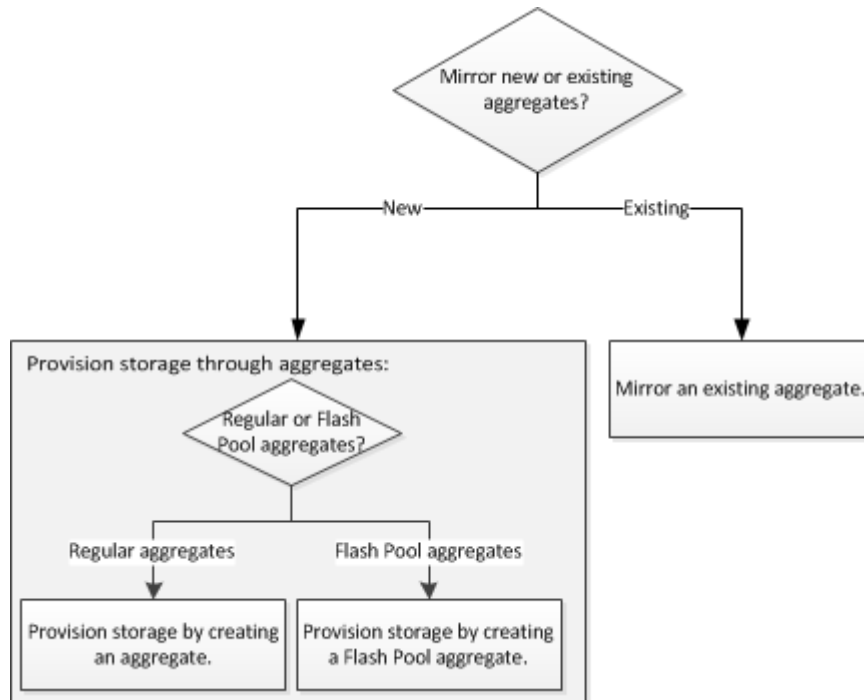
You cannot use partitioned SSDs when creating a storage pool by using System Manager.

**Steps**

1. Expand the **Cluster** hierarchy in the left navigation pane.
2. In the navigation pane, click **Storage > Storage Pools**.
3. In the **Storage Pool** window, click **Create**.
4. In the **Create Storage Pool** dialog box, specify the name for the storage pool, disk size, and the number of disks.
5. Click **Create**.

## Providing data resiliency by mirroring aggregates

You can use System Manager to protect data and provide increased resiliency by mirroring data in real-time, within a single aggregate. Mirroring aggregates removes single points of failure when connecting to disks and array LUNs.



## Provisioning storage through aggregates

You can create an aggregate or a Flash Pool aggregate to provide storage for one or more volumes by using System Manager.

### Before you begin

You must have enough spare disks to create an aggregate.

### About this task

You cannot perform the following actions by using System Manager:

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

## Provisioning storage by creating an aggregate

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

### Before you begin

All disks must be of the same size.

### Steps

1. Expand either the **Cluster** or the **Nodes** hierarchy in the left navigation pane.
2. In the navigation pane, click **Storage > Aggregates**.
3. Click **Create**.
4. In the **Create Aggregate** dialog box, perform the following steps:
  - a. Specify the name of the aggregate, the disk type, and the number of disks to include in the aggregate.  
The minimum hot spare rule is applied to the disk group that has the largest disk size.
  - b. Optional: Modify the RAID configuration of the aggregate:
    - i. Click **Change**.
    - ii. In the Change RAID Configuration dialog box, specify the RAID type and the RAID group size.  
Shared disks support the RAID type RAID-DP only.
    - iii. Click **Save**.
  - c. If you want to mirror the aggregate, select the **Mirror this aggregate** check box.  
For MetroCluster configurations, creating unmirrored aggregates is restricted and therefore, this option is enabled by default.
5. Click **Create**.

### Result

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

## Provisioning storage by creating a Flash Pool aggregate

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

### Before you begin

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

**About this task**

- You cannot use partitioned SSDs when creating the Flash Pool aggregate.
- You cannot mirror the aggregates if the cache source is storage pools.

**Steps**

1. Expand either the **Cluster** or the **Nodes** hierarchy in the left navigation pane.
2. In the navigation pane, click **Storage > Aggregates**.
3. Click **Create**.
4. In the **Create Aggregate** dialog box, specify the name of the aggregate, the disk type, and the number of HDD disks to include.
5. If you want to mirror the aggregate, select the **Mirror this aggregate** check box.

For MetroCluster configurations, creating unmirrored aggregates is restricted and therefore, this option is enabled by default.

6. Click **Use Flash Pool Cache with this aggregate**.
7. Specify the cache source by choosing one of the following actions:

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

8. Click **Create**.

**Result**

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

## Mirroring aggregates

You can use System Manager to protect data and provide increased resiliency by mirroring data in real-time, within a single aggregate. Mirroring aggregates removes single points of failure in connecting to disks and array LUNs.

### Before you begin

There must be sufficient free disks in the other pool to mirror the aggregate.

### About this task

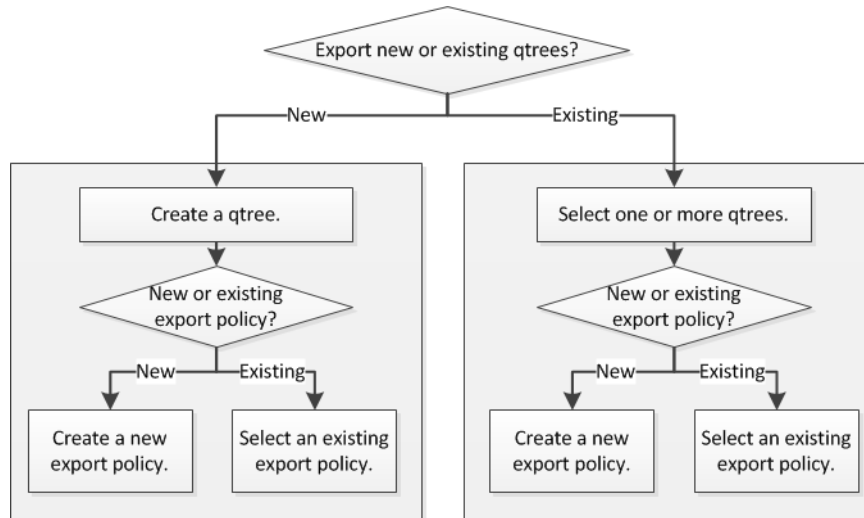
You cannot mirror Flash Pool aggregate when the cache source is storage pools.

### Steps

1. Expand either the **Cluster** or the **Nodes** hierarchy in the left navigation pane.
2. In the navigation pane, click **Storage > Aggregates**.
3. Select the aggregate that you want to mirror, and then click **Mirror**.
4. In the **Mirror this aggregate** dialog box, click **Mirror** to initiate the mirroring.

## Providing data access to qtrees using export policies

You can export a qtree by assigning an export policy to it. This enables you to export a specific qtree on a volume and make it directly accessible to clients instead of exporting the entire volume.



### Choices

- Assign export policies to new qtrees.

[Creating qtrees](#) on page 32

When creating a qtree, you can create a new export policy and add rules to it or select an existing export policy.

- Assign export policies to existing qtrees.

[Assigning export policies to qtrees](#) on page 33

When assigning an export policy to existing qtrees, you can create a new export policy and add rules to it or select an existing export policy.

## Creating qtrees

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

### Steps

1. Expand the **Storage Virtual Machines** hierarchy in the left navigation pane.
2. In the navigation pane, select the Storage Virtual Machine (SVM), and then click **Storage > Qtrees**.
3. Click **Create**.
4. In the **Details** tab of the **Create Qtree** dialog box, type a name for the qtree.



5. Select the volume to which you want to add this qtree.  
The Volume browse list includes only volumes that are online.
6. If you want to disable oplocks for the qtree, clear the **Enable Oplocks for files and directories in this Qtree** check box.  
By default, oplocks are enabled for each qtree.
7. If you want to change the default inherited security style, select a new one.  
The default security style of the qtree is the security style of the volume that contains the qtree.
8. If you want to change the default inherited export policy, select an existing export policy or create a new export policy.  
The default export policy of the qtree is the export policy assigned to the volume that contains the qtree.  
*Creating an export policy* on page 34
9. If you want to restrict the disk space usage, click the **Quotas** tab.
  - a. If you want to apply quotas on the qtree, click **Qtree quota**, and then specify the disk space limit.
  - b. If you want to apply quotas for all the users on the qtree, click **User quota**, and then specify the disk space limit.
10. Click **Create**.
11. Verify that the new qtree you created is included in the list of qtrees in the **Qtrees** window.

## Assigning export policies to qtrees

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

### Steps

1. Expand the **Storage Virtual Machines** hierarchy in the left navigation pane.
2. In the navigation pane, select the SVM, and click **Storage > Qtrees**.
3. Select one or more qtrees for which you want to assign an export policy and click **Change Export Policy**.
4. In the **Export Policy** dialog box, create a new export policy or select an existing export policy.  
*Creating an export policy* on page 34
5. Click **Save**.
6. Verify that the export policy and its related export rules that you assigned to the qtrees is displayed in the **Details** tab of the appropriate qtree.

## Creating an export policy

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

### Steps

1. Expand the **Storage Virtual Machines** hierarchy in the left navigation pane.
2. In the navigation pane, select the SVM, and then click **Policies > Export Policies**.
3. Click **Create**.
4. In the **Create Export Policy** dialog box, specify a name for the export policy.
5. If you want to create a new export policy by copying the rules from an existing export policy, select the **Copy Rules from** check box, and then select the SVM and the export policy.

You should not select the destination SVM for disaster recovery from the drop-down menu to create an export policy.

6. In the **Export Rules** area, click **Add** to add rules to the export policy.  
*[Adding rules to an export policy](#) on page 34*
7. Click **Create**.
8. Verify that the export policy you created is displayed in the **Export Policies** window.

### Related tasks

*[Creating qtrees](#) on page 32*

*[Assigning export policies to qtrees](#) on page 33*

## Adding rules to an export policy

You can use System Manager to add rules to an export policy, which enables you to define client access to data.

### Before you begin

You must have created the export policy to which you want to add the export rules.

### Steps

1. Expand the **Storage Virtual Machines** hierarchy in the left navigation pane.
2. In the navigation pane, select the Storage Virtual Machine (SVM) and click **Policies > Export Policies**.
3. Select the export policy and from the **Export Rules** tab, click **Add**.
4. In the **Create Export Rule** dialog box, perform the following steps:
  - a. Specify the client that requires access to the data.

You can specify the client in any of the following formats:

- As a host name; for instance, host1
- As an IPv4 address; for instance, 10.1.12.24

- As an IPv4 address with a network mask; for instance, 10.1.16.0/255.255.255.0
- As an IPv6 address; for instance, FE80::0202:B3FF:FE1E:8329
- As an IPv6 address with a network mask; for instance, 2001:db8::/32
- As a netgroup, with the netgroup name preceded by an at sign (@); for instance, @netgroup
- As a domain name preceded by a period (.); for instance, .example.com

**Note:** You must not enter an IP address range, such as 10.1.12.10 through 10.1.12.70. Entries in this format are interpreted as a text string and treated as a host name.

You can enter the IPv4 address `0.0.0.0/0` to provide access to all the hosts.

- b. If you want to modify the rule index number, select the appropriate rule index number.
- c. Select one or more access protocols.

If you do not select any access protocol, the default value “Any” is assigned to the export rule.

- d. Select one or more security types and access rules.
5. Click **OK**.
  6. Verify that the export rule you added is displayed in the **Export Rules** tab for the selected export policy.

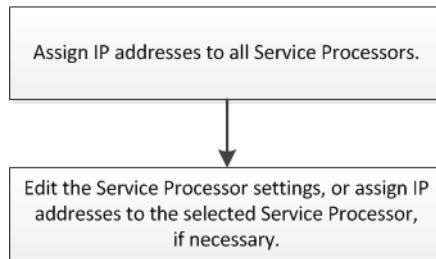
#### **Related tasks**

[Creating an export policy](#) on page 34

## Configuring a Service Processor

---

You can configure a Service Processor to monitor and manage various system parameters of your storage systems. Configuring a Service Processor involves assigning IP addresses to all the Service Processors and modifying the settings. If required, you can assign IP addresses to the selected Service Processors.



### Steps

1. [Assigning IP addresses to Service Processors](#) on page 36  
You can use System Manager to assign IP addresses to all your Service Processors at the same time and use these Service Processors to monitor and manage various system parameters of your storage systems.
2. [Editing Service Processor settings](#) on page 37  
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.

## Assigning IP addresses to Service Processors

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

### Steps

1. Expand the **Cluster** hierarchy in the left navigation pane.
2. In the navigation pane, select the cluster, and click **Configuration > Service Processor**.
3. In the **Service Processor** window, click **Global Settings**.
4. In the **Global Settings** dialog box, choose the source from which to assign the IP addresses:

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

5. Click **Save**.

## Editing Service Processor settings

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

### About this task

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


### Steps

1. Expand the **Cluster** hierarchy in the left navigation pane.
2. Select the cluster, and click **Configuration > Service Processor**.
3. In the **Service Processor** window, select the Service Processor, and click **Edit**.
4. In the **Edit Service Processor** dialog box, make the necessary changes, and click **Save and Close**.

## Setting up a BranchCache configuration

---

You can configure BranchCache on a CIFS-enabled Storage Virtual Machine (SVM) and enable BranchCache on the required shares to enable caching of content on computers local to requesting clients.



### Steps

1. [Setting up CIFS](#) on page 38  
You can use System Manager to enable and configure CIFS servers to allow CIFS clients to access files on the cluster.
2. [Setting up BranchCache](#) on page 39  
You can use System Manager to configure BranchCache on a CIFS-enabled Storage Virtual Machine (SVM) to enable caching of content on computers local to requesting clients.
3. [Editing share settings](#) on page 40  
You can use System Manager to modify the settings of a share, such as the symbolic link settings, share access permissions of users or groups, and the type of access to the share. You can also enable or disable continuous availability of a share over Hyper-V, and enable or disable access-based enumeration (ABE).

## Setting up CIFS

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

### Before you begin

- The CIFS license must be installed on your storage system.
- While configuring CIFS in the Active Directory domain, the following requirements must be met:
  - DNS must be enabled and configured correctly.
  - The storage system must be able to communicate with the domain controller using the fully qualified domain name (FQDN).
  - The time difference (clock skew) between the cluster and the domain controller must not be more than five minutes.

- If CIFS is the only protocol configured on the Storage Virtual Machine (SVM), the following requirements must be met:
  - The root volume security style must be NTFS.  
By default, System Manager sets the security style as UNIX.
  - Superuser access must be set to **Any** for CIFS protocol.

### Steps

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

## Setting up BranchCache

You can use System Manager to configure BranchCache on a CIFS-enabled Storage Virtual Machine (SVM) to enable caching of content on computers local to requesting clients.

### Before you begin

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

### About this task

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

### Steps

1. Expand the **Storage Virtual Machines** hierarchy in the left navigation pane.
2. In the navigation pane, select the SVM and click **Configuration > Protocols > CIFS**.
3. In the **CIFS** window, click **BranchCache > Set Up**.

4. In the **BranchCache Setup** dialog box, enter the following information:
  - a. Specify the path to the hash store.  
The path can be to an existing directory where you want the hash data to be stored. The destination path must be read-writable. Read-only paths, such as Snapshot directories, are not allowed.
  - b. Specify the maximum size (in KB, MB, GB, TB, or PB) for a hash data store.  
If the hash data exceeds this value, older hashes are deleted to provide space for newer hashes. The default size for the hash store is 1 GB.
  - c. Specify the operating mode for the BranchCache configuration.  
The default operating mode is set to all shares.
  - d. Specify a server key to prevent clients from impersonating the BranchCache server.  
You can set the server key to a specific value so that if multiple servers are providing BranchCache data for the same files, clients can use hashes from any server using that same server key. If the server key contains any spaces, you must enclose the server key in quotation marks.
  - e. Select the required BranchCache version.  
By default, all the versions supported by the client are selected.
5. Click **Set Up**.

## Editing share settings

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

### Steps

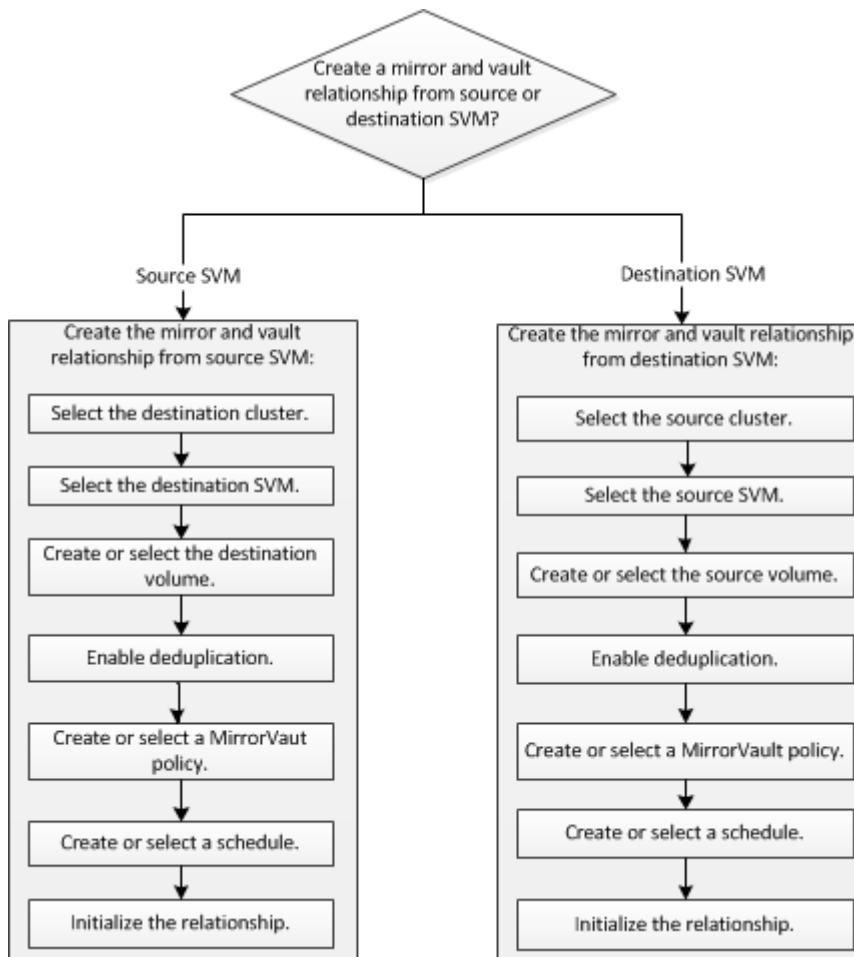
1. Expand the **Storage Virtual Machines** hierarchy in the left navigation pane.
2. In the navigation pane, select the Storage Virtual Machine (SVM), and click **Storage > Shares**.
3. Select the share that you want to modify from the share list and click **Edit**.
4. In the **Edit Share Settings** dialog box, modify the share settings as required:
  - a. In the **General** tab, enable continuous availability of a share over Hyper-V.  
Enabling continuous availability permits SMB 3.0 and later clients that support SMB 3.0 to open files persistently during nondisruptive operations. Files that are opened persistently are protected from disruptive events, such as failover, giveback, and LIF migration.
  - b. In the **Permissions** tab, add users or groups and assign permissions to specify the type of access.
  - c. In the **Options** tab, perform the following actions on the share:
    - Select the settings for the symbolic links.
    - Enable opportunistic locks (oplocks).
    - Enable clients to browse through the share.



- View Snapshot copies.
  - Notify changes.
  - Enable ABE.
  - Enable BranchCache.
  - Enable data to be encrypted using SMB 3.0 while accessing this share.
5. Click **Save and Close**.
  6. Verify the changes you made to the selected share in the **Shares** window.

## Protecting data using mirror and vault relationship

You can create a mirror and vault relationship using the MirrorVault policy to protect your data by periodically transferring data from the source volume to the destination volume. The relationship also enables you to retain data for long periods by creating backups of the source volume.



### Creating a mirror and vault relationship from a source volume

You can use System Manager to create a mirror and vault relationship between a source volume and a destination volume. Creating this relationship enables you to better protect your data by periodically transferring data from the source volume to the destination volume. It also enables you to retain data for long periods by creating backups of the source volume.

#### Before you begin

- The source cluster must be running Data ONTAP 8.3.2 or later.
- The SnapMirror and SnapVault licenses must be enabled on both the source and the destination clusters that contain the source and destination volumes.

- The clusters must not have the same name.
- The source and destination clusters must be in a healthy peer relationship.
- The destination aggregate must have available space.
- The source aggregate and the destination aggregate must be 64-bit aggregates.
- If the destination volume exists, the volume must not be the destination for any other protection relationship.
- The capacity of the destination volume must be greater than or equal to the capacity of the source volume.
- If autogrow is disabled, the free space on the destination volume must be at least five percent more than the used space on the source volume.

#### About this task

- System Manager does not support a cascade relationship.  
For example, a destination volume in a relationship cannot be the source volume in another relationship.
- The destination volume that is created for a mirror and vault relationship is not thin-provisioned.

#### Steps

1. Expand the **Storage Virtual Machines** hierarchy in the left navigation pane.
2. In the navigation pane, select the Storage Virtual Machine (SVM), and then click **Storage > Volumes**.
3. Select the volume for which you want to create a mirror and vault relationship, and then click **Protect**.  
The **Protect** option is available only for a read/write volume.
4. In the **Create Protection Relationship** dialog box, select **Mirror and Vault** from the Relationship Type drop-down list.
5. Specify the cluster, the SVM, and the destination volume.
6. If the selected SVM is not peered, use the **Authenticate** link to enter the credentials of the remote cluster and create the SVM peer relationship.
7. Create a new destination volume or select an existing volume:

If you want to...	Do the following...
Create a new volume	<ol style="list-style-type: none"> <li>a. If you want to change the default name, displayed in the format <i>source_SVM_name_source_volume_name_mirror_vault</i>, specify a new name, and then select the containing aggregate for the destination volume.</li> <li>b. Select <b>Enable dedupe</b> to enable deduplication on the new destination volume. If deduplication is disabled on the source volume, then the check box for the new volume is selected by default.</li> </ol>
Select an existing volume	<p>Select the <b>Select Volume</b> option.</p> <p><b>Note:</b> Only those volumes with the same language attribute as the source volume are listed.</p>

## 8. Select an existing policy or create a new policy:

If you want to...	Do the following...
Select an existing policy	Click <b>Browse</b> , and then select a mirror and vault policy.  You can select a policy that has the maximum number of matching labels with the Snapshot policy that is attached to the source volume.
Create a new policy	<p>a. Click <b>Create Policy</b>.</p> <p>b. Specify the policy name and set the schedule transfer priority. Low indicates that the transfer has the least priority and is usually scheduled after normal priority transfers. By default, the priority is set to normal.</p> <p>c. Select the <b>Enable Network Compression</b> check box to compress the data that is being transferred.</p> <p>d. Click <b>Create</b>.</p> <p>You can also specify the SnapMirror label and destination retention count for the policy. However, you must ensure that a Snapshot copy with the same label is created on the source volume for the new SnapMirror label to be effective.</p>

## 9. Specify a schedule for the relationship:

If...	Do the following...
You want to assign an existing schedule	From the list of schedules, select an existing schedule.
You want to create a new schedule	<p>a. Click <b>Create Schedule</b>.</p> <p>b. Specify a name for the schedule.</p> <p>c. Select either <b>Basic</b> or <b>Advanced</b>.</p> <ul style="list-style-type: none"> <li>Basic: You can select this option to specify only the day of the week, time, and the transfer interval.</li> <li>Advanced: You can select this option to specify a cron style schedule.</li> </ul> <p>d. Click <b>Create</b>.</p>
You do not want to assign a schedule	Select <b>None</b> .

10. Optional: Select **Initialize Relationship** to initialize the relationship.11. Click **Create**.

## Creating a mirror and vault relationship from a destination SVM

You can use System Manager to create a mirror and vault relationship between a source volume and a destination volume from the destination Storage Virtual Machine (SVM). Creating this relationship enables you to better protect your data by periodically transferring data from the source volume to the

destination volume. It also enables you to retain data for long periods by creating backups of the source volume.

### Before you begin

- The destination cluster must be running Data ONTAP 8.3.2 or later.
- The SnapMirror and SnapVault licenses must be enabled on both the source and the destination clusters that contain the source and destination volumes.
- The clusters must not have the same name.
- The source and destination clusters must be in a healthy peer relationship.
- The destination aggregate must have available space.
- The source aggregate and the destination aggregate must be 64-bit aggregates.
- A source volume of type read/write (rw) must already exist.
- If the destination volume exists, the capacity of the destination volume must be greater than or equal to the capacity of the source volume.
- The destination volume must not be the root volume of a storage system.
- If the destination volume exists, the volume must not be the destination for any other mirror relationship.
- If autogrow functionality is disabled, the free space on the destination volume must be at least five percent more than the used space on the source volume.

### About this task

- System Manager does not support a cascade relationship.  
For example, a destination volume in a relationship cannot be the source volume in another relationship.
- You cannot create a mirror and vault relationship between a sync-source SVM and a sync-destination SVM in a MetroCluster configuration.
- You can create a mirror and vault relationship between sync-source SVMs in a MetroCluster configuration.
- You can create a mirror and vault relationship from a volume on a sync-source SVM to a data-serving SVM.
- You can create a mirror and vault relationship from a volume on a data-serving SVM to a DP volume on a sync-source SVM.
- The destination volume that is created for a mirror relationship is not thin provisioned.

### Steps

1. Expand the **Storage Virtual Machines** hierarchy in the left navigation pane.
2. In the navigation pane, select the SVM, and then click **Protection**.
3. In the **Protection** window, click **Create**.
4. In the **Create Protection Relationship** dialog box, select **Mirror and Vault** from the Relationship Type drop-down list.
5. Specify the cluster, the SVM, and the source volume.

6. If the selected SVM is not peered, use the **Authenticate** link to enter the credentials of the remote cluster and create the SVM peer relationship.

7. Create a new destination volume or select an existing volume:

If you want to...	Do the following...
Create a new volume	<p>a. If you want to change the default name, displayed in the format <i>source_SVM_name_source_volume_name_mirror_vault</i>, specify a new name and select the containing aggregate for the destination volume.</p> <p>b. Select <b>Enable dedupe</b> to enable deduplication on the new destination volume. If deduplication is disabled on the source volume, then the check box for the new volume is selected by default.</p>
Select an existing volume	<p>Select the <b>Select Volume</b> option.</p> <p><b>Note:</b> Only those volumes with the same language attribute as the source volume are listed.</p>

8. Select an existing policy or create a new policy:

If you want to...	Do the following...
Select an existing policy	<p>Click <b>Browse</b>, and then select a mirror and vault policy.</p> <p>You can select a policy that has the maximum number of matching labels with the Snapshot policy that is attached to the source volume.</p>
Create a new policy	<p>a. Click <b>Create Policy</b>.</p> <p>b. Specify a policy name and set the schedule transfer priority. Low indicates that the transfer has the least priority and is usually scheduled after normal priority transfers. By default, the priority is set to normal.</p> <p>c. Select the <b>Enable Network Compression</b> check box to compress the data that is being transferred.</p> <p>d. Click <b>Create</b>.</p> <p>You can also specify the SnapMirror label and destination retention count for the policy. However, you must ensure that a Snapshot copy with the same label is created on the source volume for the new SnapMirror label to be effective.</p>

9. Specify a schedule for the relationship:

If...	Do the following...
You want to assign an existing schedule	From the list of schedules, select an existing schedule.

---

<b>If...</b>	<b>Do the following...</b>
You want to create a new schedule	<ol style="list-style-type: none"><li data-bbox="651 254 1372 289"><b>a.</b> Click <b>Create Schedule</b>.</li><li data-bbox="651 304 1372 340"><b>b.</b> Specify a name for the schedule.</li><li data-bbox="651 354 1372 548"><b>c.</b> Select either <b>Basic</b> or <b>Advanced</b>.<ul style="list-style-type: none"><li data-bbox="699 407 1372 464">• <b>Basic:</b> You can select this option to specify only the day of the week, time, and the transfer interval.</li><li data-bbox="699 485 1372 548">• <b>Advanced:</b> You can select this option to specify a cron style schedule.</li></ul></li><li data-bbox="651 562 1372 598"><b>d.</b> Click <b>Create</b>.</li></ol>
You do not want to assign a schedule	Select <b>None</b> .

---

**10.** Optional: Select **Initialize Relationship** to initialize the relationship.

**11.** Click **Create**.

## Copyright information

---

Copyright © 1994–2016 NetApp, Inc. All rights reserved. Printed in the U.S.

No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

**RESTRICTED RIGHTS LEGEND:** Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).



## Trademark information

---

NetApp, the NetApp logo, Go Further, Faster, AltaVault, ASUP, AutoSupport, Campaign Express, Cloud ONTAP, Clustered Data ONTAP, Customer Fitness, Data ONTAP, DataMotion, Fitness, Flash Accel, Flash Cache, Flash Pool, FlashRay, FlexArray, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexVol, FPolicy, GetSuccessful, LockVault, Manage ONTAP, Mars, MetroCluster, MultiStore, NetApp Insight, OnCommand, ONTAP, ONTAPI, RAID DP, RAID-TEC, SANtricity, SecureShare, Simplicity, Simulate ONTAP, Snap Creator, SnapCenter, SnapCopy, SnapDrive, SnapIntegrator, SnapLock, SnapManager, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapValidator, SnapVault, StorageGRID, Tech OnTap, Unbound Cloud, and WAFL and other names are trademarks or registered trademarks of NetApp, Inc., in the United States, and/or other countries. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such. A current list of NetApp trademarks is available on the web at <http://www.netapp.com/us/legal/netapptmlist.aspx>.

## How to send comments about documentation and receive update notifications

---

You can help us to improve the quality of our documentation by sending us your feedback. You can receive automatic notification when production-level (GA/FCS) documentation is initially released or important changes are made to existing production-level documents.

If you have suggestions for improving this document, send us your comments by email to [doccomments@netapp.com](mailto:doccomments@netapp.com). To help us direct your comments to the correct division, include in the subject line the product name, version, and operating system.

If you want to be notified automatically when production-level documentation is released or important changes are made to existing production-level documents, follow Twitter account @NetAppDoc.

You can also contact us in the following ways:

- NetApp, Inc., 495 East Java Drive, Sunnyvale, CA 94089 U.S.
- Telephone: +1 (408) 822-6000
- Fax: +1 (408) 822-4501
- Support telephone: +1 (888) 463-8277

# Index

## A

### ABE

enabling [40](#)

### access-based enumeration

*See* ABE

### adding

capacity disks to aggregates [21](#)

export rules [34](#)

rules to export policies [34](#)

SSDs [29](#)

SSDs to aggregates [22, 25](#)

VLAN interfaces [12](#)

### aggregates

adding capacity disks to [21](#)

assigning array LUNs to [17](#)

assigning disks to increase capacity [16](#)

creating storage pools [23, 26](#)

Flash Pool, workflow for provisioning cache [24](#)

mirroring [31](#)

provisioning cache by adding SSDs [22, 25](#)

provisioning storage by creating [28](#)

provisioning storage by creating Flash Pool [29](#)

provisioning storage by creating HDD and SSD [29](#)

workflow for assigning disks to existing [19](#)

workflow for mirroring data [28](#)

zeroing spare array LUNs [18](#)

zeroing spare disks [17](#)

### all-share cache

creating for BranchCache configuration [39](#)

### array LUNs

assigning [17](#)

erasing data from [18](#)

formatting [18](#)

zeroing [18](#)

### assigning

array LUNs [17](#)

disks to nodes [16](#)

export policies to qtrees [33](#)

## B

### BranchCache

enabling [40](#)

setting up [39](#)

workflow for configuring [38](#)

### broadcast domains

creating [13](#)

## C

### cache

increasing for Flash Pool aggregates by adding SSDs [26](#)

provisioning to aggregates by adding SSDs [22, 25](#)

### caching

configuring BranchCache for [39](#)

### CIFS

setting up [38](#)

workflow for configuring BranchCache [38](#)

### CIFS servers

configuring BranchCache on [39](#)

### client access

adding rules to export policies [34](#)

creating export policies [34](#)

creating export policy for [34](#)

setting up CIFS [38](#)

### clusters

nondisruptively updating using System Manager [8](#)

workflow for a nondisruptive update [7](#)

### comments

how to send feedback about documentation [50](#)

### configuring

BranchCache [39](#)

### continuous availability

enabling or disabling for shares [40](#)

### creating

a mirror and vault relationship from a destination

SVM [44](#)

aggregates [28](#)

broadcast domains [13](#)

export policies [34](#)

Flash Pool aggregates [28, 29](#)

HDD and SSD aggregates [29](#)

interface groups [12](#)

mirror and vault relationships from a source volume

[42](#)

network interfaces [14](#)

qtrees [32](#)

storage pools [23, 26](#)

subnets [13](#)

## D

### Data ONTAP

workflow for setting up a network [11](#)

### Data ONTAP software images

obtaining [7](#)

selecting for nondisruptive update [8](#)

### dedicated SSDs

adding to Flash Pool aggregates [26](#)

### DHCP

assigning IP addresses to Service Processors [36](#)

### disks

assigning to nodes [16](#)

erasing data from [17](#)

formatting [17](#)

increasing the size of aggregates [21](#)

mirroring aggregates [31](#)

zeroing [17](#)

### documentation

how to receive automatic notification of changes to [50](#)

how to send feedback about [50](#)

### domains

creating broadcast [13](#)

### downgrades

obtaining Data ONTAP software images [7](#)

## E

editing

- Service Processor settings [37](#)
- share settings [40](#)
- shares, workflow to enable BranchCache [38](#)

encrypt data

- while accessing this share [40](#)

export policies

- adding export rules to [34](#)
- adding rules [34](#)
- assigning to qtrees [33](#)
- creating [34](#)
- workflow for assigning to qtrees [32](#)

export rules

- adding to export policies [34](#)

exporting

- qtrees [33](#)

express guides

- System Manager workflows in [5](#)

## F

feedback

- how to send comments about documentation [50](#)

Flash Pool aggregates

- creating [29](#)
- creating storage pools [23, 26](#)
- increasing the size by adding SSDs [26](#)
- provisioning storage by creating [28](#)
- workflow for provisioning cache [24](#)

flowcharts

- Service Processor configuration workflow [36](#)
- setting up a network [11](#)
- setting up the BranchCache configuration [38](#)
- workflow for assigning array LUNs to nodes [16](#)
- workflow for assigning disks to existing aggregates [19](#)
- workflow for assigning disks to nodes [16](#)
- workflow for assigning export policies to qtrees [32](#)
- workflow for mirroring aggregates [28](#)
- workflow for nondisruptive cluster update [7](#)
- workflow for protecting data using mirror and vault relationship [42](#)
- workflow for provisioning SSD cache [24](#)

## H

hash stores

- specifying path and maximum size for BranchCache configuration [39](#)

HDDs

- converting existing HDD aggregate to a Flash Pool aggregate [29](#)
- creating aggregates [29](#)

HDDs and SSDs

- adding to aggregates [21](#)

## I

ifgroups

- See* interface groups

information

- how to send feedback about improving documentation [50](#)

intercluster connectivity

- creating network interfaces for [14](#)

interface groups

- creating [12](#)

IP addresses

- assigning to multiple Service Processors [36](#)
- editing for a Service Processor [37](#)

## L

LIFs

- creating [14](#)

LUNs, array

- assigning [17](#)

## M

mirror and vault

- creating a relationship from a source volume [42](#)

mirror and vault relationship

- creating from a destination SVM [44](#)

mirror and vault relationships

- workflow for protecting data [42](#)
- workflow for providing data protection [42](#)

mirroring

- aggregates [31](#)

mirroring aggregates

- workflow for providing data resiliency [28](#)

modifying

- Service Processor settings [37](#)
- the share settings [40](#)

## N

NDU

- See* nondisruptive updates

NetApp Support Site

- obtaining Data ONTAP software images [7](#)

network

- creating subnets [13](#)
- creating VLANs [12](#)
- setup workflow [11](#)

network interfaces

- creating [14](#)
- creating VLANs [12](#)

network mask

- editing for a Service Processor [37](#)

nodes

- assigning disks [16](#)
- workflow for assigning array LUNs to [16](#)
- workflow for assigning disks to [16](#)

non-root aggregates

- adding capacity disks to [21](#)

nondisruptive updates

- performing for clusters using System Manager [8](#)

workflow for updating clusters [7](#)

## O

oplocks  
enabling [40](#)

## P

per-share cache  
creating for BranchCache configuration [39](#)  
policies  
assigning export, to qtrees [33](#)  
pre-update checks  
validating the cluster before updating [8](#)  
provisioning storage  
by creating HDD and SSD aggregates [29](#)  
creating aggregates or Flash Pool aggregates [28](#)  
creating Flash Pool aggregates [29](#)

## Q

qtrees  
assigning an export policy to [33](#)  
creating [32](#)  
exporting [33](#)  
workflow for assigning export policies [32](#)

## R

relationships  
mirror and vault, creating from a destination SVM [44](#)  
mirror and vault, creating from a source volume [42](#)  
reversions  
obtaining Data ONTAP software images [7](#)  
root aggregates  
adding capacity disks to [21](#)

## S

server keys  
specifying for BranchCache configuration [39](#)  
Service Processors  
assigning IP addresses globally [36](#)  
configuration workflow [36](#)  
editing the settings [37](#)  
setting up  
BranchCache [39](#)  
CIFS [38](#)  
shares  
editing permissions and options [40](#)  
SMB  
configuring BranchCache on [39](#)  
SMB encryption  
setting up [38](#)  
software images  
obtaining [7](#)  
selecting the Data ONTAP image for update [8](#)  
spare array LUNs  
zeroing [18](#)

spare disks  
zeroing [17](#)  
SPs  
*See* Service Processors  
SSDs  
adding to HDDs to convert as a Flash Pool aggregate [29](#)  
combining to create storage pools [23, 26](#)  
creating aggregates [29](#)  
dedicated SSDs [22, 25](#)  
increasing the size of Flash Pool aggregates by adding [26](#)  
provisioning cache by adding [22, 25](#)  
storage pools [22, 25](#)  
workflow for provisioning cache [24](#)  
storage  
provisioning by creating aggregates [29](#)  
storage pools  
adding SSDs as [26](#)  
assigning disks to increase capacity [16](#)  
creating [23, 26](#)  
subnets  
assigning IP addresses to Service Processors [36](#)  
creating [13](#)  
specifying the IP address for the network interface [14](#)  
suggestions  
how to send feedback about documentation [50](#)  
SVMs  
creating network interfaces for managing [14](#)  
setting up BranchCache on [39](#)  
symbolic links  
editing the settings of [40](#)  
System Manager  
workflows in the 8.3 release [5](#)  
workflows in the 8.3.1 release [5](#)  
workflows in the 8.3.2 release [5](#)

## T

Twitter  
how to receive automatic notification of documentation changes [50](#)

## U

unowned disks  
assigning ownership [16](#)  
updates  
performing a nondisruptive upgrade through System Manager [8](#)  
upgrades  
obtaining Data ONTAP software images [7](#)  
*See also* nondisruptive updates

## V

validating  
checking the cluster before an update [8](#)  
VLAN interfaces  
creating [12](#)  
volumes

creating an export policy for client access to [34](#)  
creating qtrees [32](#)

## W

workflow guides

contents of this guides [5](#)

workflows

assigning array LUNs to existing aggregates [16](#)

assigning disks to existing aggregates [16](#), [19](#)

configuring a Service Processor [36](#)

for nondisruptive cluster update [7](#)

in System Manager 8.3 [5](#)

in System Manager 8.3.1 [5](#)

in System Manager 8.3.2 [5](#)

providing data access to qtrees by using export policies [32](#)

providing data protection by creating mirror and vault relationships [42](#)

providing data resiliency by mirroring aggregates [28](#)

provisioning SSD cache through Flash pool aggregates [24](#)

setting up a network [11](#)

setting up the BranchCache configuration [38](#)

## Z

zeroing

array LUNs [18](#)

disks [17](#)