



NetApp®

Clustered Data ONTAP® 8.3

Cluster Management Workflows for OnCommand® System Manager



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
Web: www.netapp.com
Feedback: doccomments@netapp.com

Part number: 215-09869_A0
March 2015

Contents

Contents of the Workflow Guide	5
Setting up the network	7
Creating VLAN interfaces	7
Creating interface groups	8
Creating broadcast domains	9
Creating subnets	9
Creating network interfaces	10
Assigning disks to existing aggregates	12
Adding capacity disks (HDDs)	12
Provisioning cache to aggregates by adding SSDs	14
Creating a storage pool	15
Provisioning cache through Flash Pool aggregates	17
Provisioning cache to aggregates by adding SSDs	18
Increasing the cache for Flash Pool aggregates by adding SSDs	19
Creating a storage pool	20
Providing data access to qtrees using export policies	21
Assigning export policies to new qtrees	21
Creating qtrees	21
Assigning export policies to existing qtrees	22
Assigning export policies to qtrees	22
Creating an export policy	23
Adding rules to an export policy	23
Configuring a Service Processor	25
Assigning IP addresses to Service Processors	25
Editing Service Processor settings	26
Setting up the BranchCache configuration	27
Setting up CIFS	27
Setting up BranchCache	28
Editing share settings	29
Copyright information	31
Trademark information	32

How to send comments about documentation and receive update notification	33
Index	34

Contents of the Workflow Guide

This guide provides workflows for some of the new functionality introduced in OnCommand System Manager 8.3. These workflows help you in setting up your network, provisioning space, setting up BranchCache, assigning export policies to qtrees, and configuring Service Processors.

This guide provides workflows for the following functionalities that were introduced or enhanced in System Manager 8.3:

- *Setting up the network* on page 7
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 12
Provides information about adding disks (capacity or cache) to increase the size of aggregates to provide storage for one or more volumes.
- *Provisioning cache through Flash Pool aggregates* on page 17
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 21
Provides information about exporting a specific qtree on a volume by assigning export policies to it.
- *Configuring a Service Processor* on page 25
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 27
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 also includes the following functionalities that are covered in the workflows in the 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.

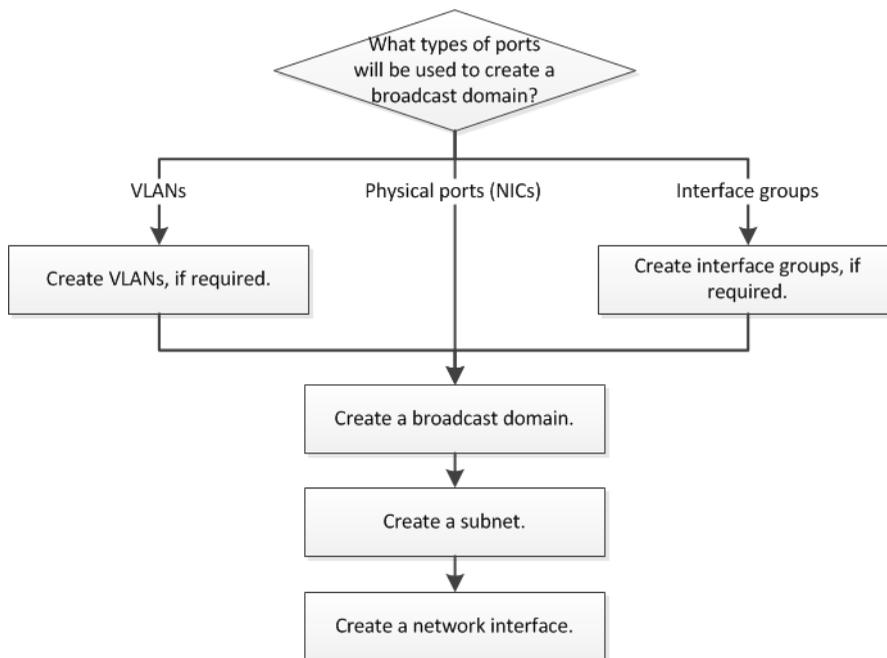
6 | Cluster Management Workflows for OnCommand System Manager

- 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](#)

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.



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 can use System Manager to create an interface group to assemble multiple network interfaces into one logical interface to provide high throughput and fault tolerance, and to avoid single point of failure.

Before you begin

There must be free ports available, which 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 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, or sequential
 - 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.

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 **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 to access data from Storage Virtual Machines (SVMs), manage SVMs, and to provide an interface for intercluster connectivity.

Before you begin

- You must have created a subnet.
The subnet that you created is used for obtaining an IP address or adding the IP address that you enter manually for the network interface.
- The broadcast domain associated with the subnet must have allocated ports.

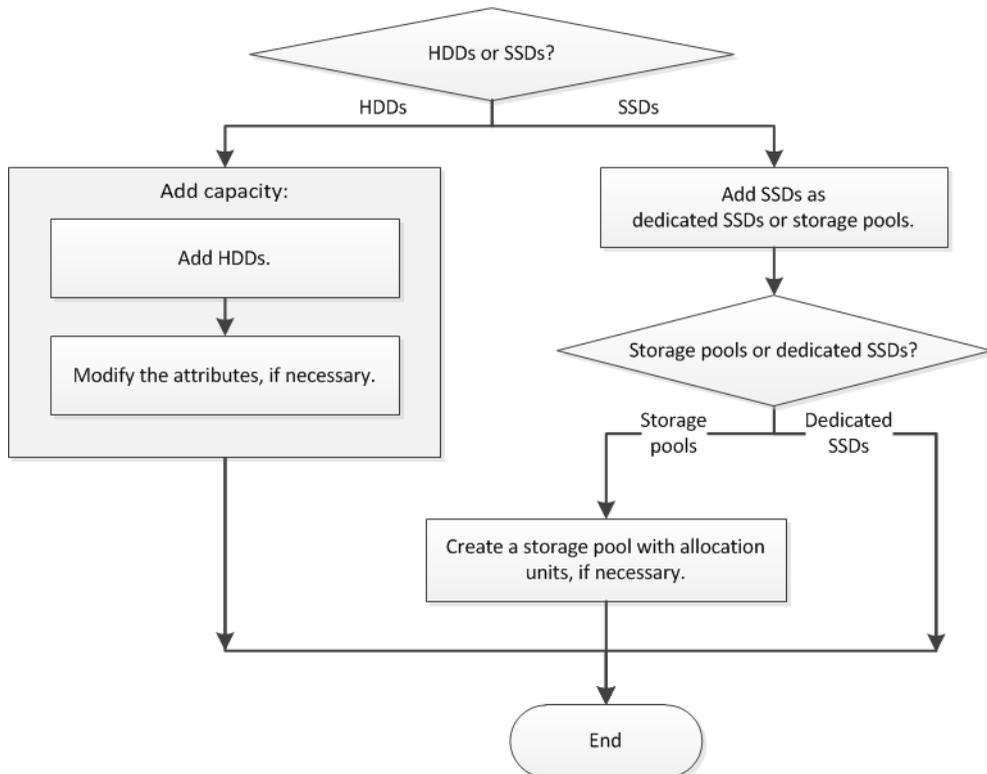
Steps

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, perform the following steps:
 - a. Specify an interface name.
 - b. Specify a role, whether the interface should be used for serving data (data LIF) or for intercluster connectivity (intercluster LIF).
 - c. Select the appropriate protocols, as required.
The interface uses the selected protocols to access data from the SVM.
 - d. Optional: If you want to enable management access on a 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.
 - e. Specify the IP address by choosing one of the following:
 - Select the subnet from which the IP address will be assigned automatically to the interface.
 - Select **Use this IP Address** and type the IP address.
The IP address you specify will be added to the subnet if it is not already present in the subnet range.

6. Select the required ports from the ports details area, which displays all the ports from the broadcast domain associated with the selected subnet.
7. Click **Create**.

Assigning disks to existing aggregates

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



Adding capacity disks (HDDs)

You can increase the size of an existing non-root aggregate or a root aggregate containing non-partitioned disks by adding capacity disks (HDDs). You can use System Manager to add HDDs 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.
- You cannot use System Manager to add HDDs to the following:
 - SyncMirror aggregates
 - 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 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**.

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 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"> a. Select the storage pool from which cache can be obtained. b. Specify the cache size. c. Modify the RAID type, if required.
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"> a. Click Change. b. In the Change RAID Configuration dialog box, specify the RAID type and RAID group size, and then click Save.

5. Click **Add**.

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.

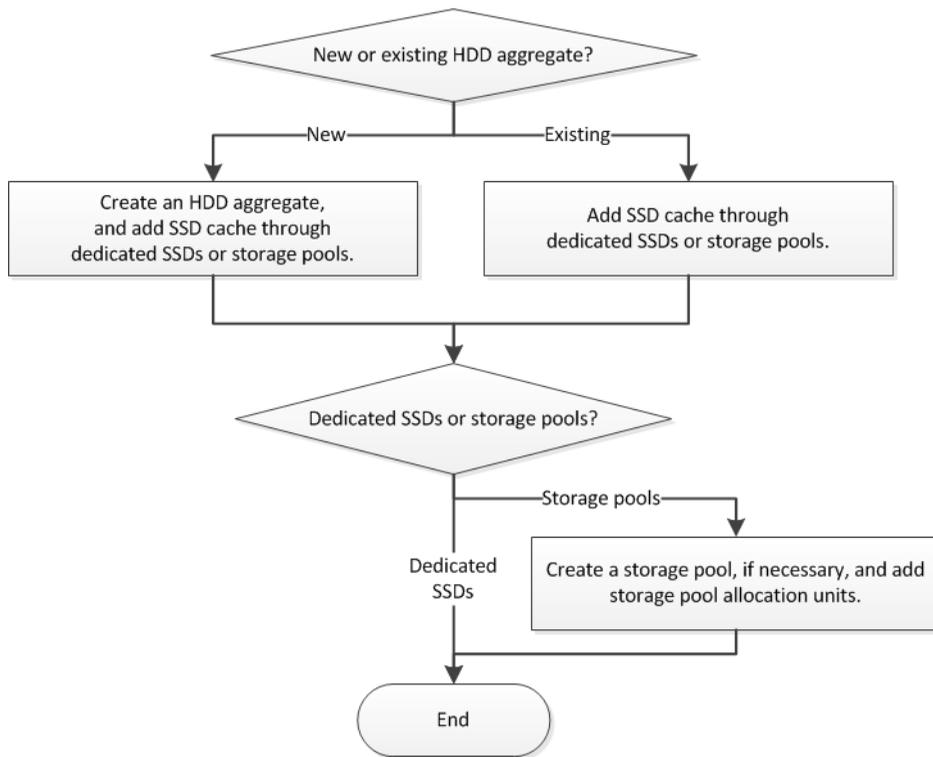
16 | Cluster Management Workflows for OnCommand System Manager

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.



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 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"> a. Select the storage pool from which cache can be obtained. b. Specify the cache size. c. Modify the RAID type, if required.

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"> a. Click Change. b. In the Change RAID Configuration dialog box, specify the RAID type and RAID group size, and then click Save.

5. Click **Add**.

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 click **Add Cache**.
4. In the **Add Cache** dialog box, perform the appropriate action:

If you selected the cache source as...	Do this...
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**.

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.

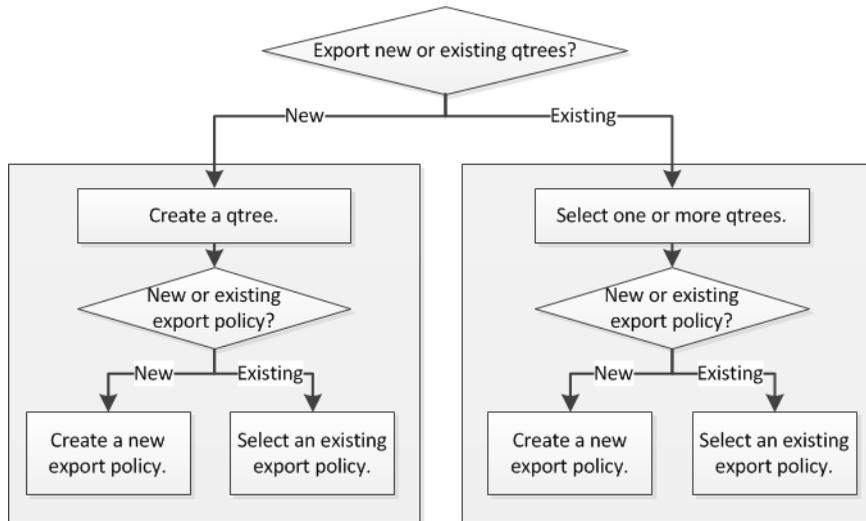
You cannot use partitioned SSDs when creating a storage pool by using System Manager.

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 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.



Assigning export policies to new qtrees

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 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 23

9. If you want to restrict disk space usage, click the **Quotas** tab.
 - a. If you want to apply quotas on the qtree, click **Qtree quota** and specify the disk space limit.
 - b. If you want to apply quotas for all the users on the qtree, click **User quota** and 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 existing qtrees

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 23

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 Storage Virtual Machine (SVM) and 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 select the SVM and the 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 23
7. Click **Create**.
8. Verify that the export policy you created is displayed in the **Export Policies** window.

Related tasks

[Creating qtrees](#) on page 21

[Assigning export policies to qtrees](#) on page 22

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.

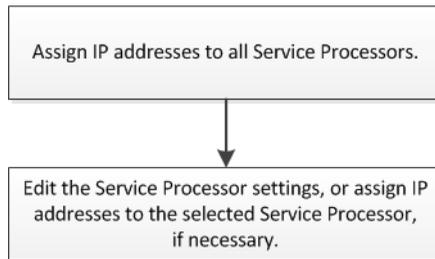
- b. If you want to modify the rule index number, select the appropriate rule index number.
 - c. Select one or more access protocols.
If you do not select any access protocol, the default value “Any” is assigned to the export rule.
 - d. Select one or more security types and access rules.
5. Click **OK**.
 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 23

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.



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

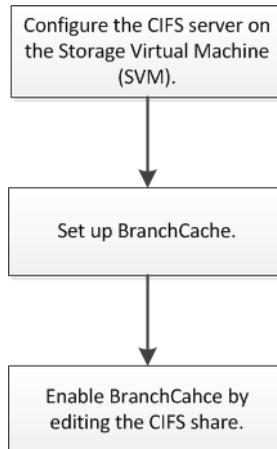
- You can edit the settings of a Service Processor that was assigned IP addresses manually.
- You cannot edit the settings of a Service Processor that was assigned IP addresses through a DHCP server or through a subnet.

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 the 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.



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 click **Configuration > Protocols > CIFS**.
3. In the **Configuration** tab, click **Setup**.
4. In the **CIFS Server Setup** dialog box, specify the NetBIOS name and the Active Directory domain details.
5. 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.

30 | Cluster Management Workflows for OnCommand System Manager

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.
5. Click **Save and Close**.
 6. Verify the changes you made to the selected share in the **Shares** window.

Copyright information

Copyright © 1994–2015 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, 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, SANtricity, SecureShare, Simplicity, Simulate ONTAP, Snap Creator, SnapCopy, SnapDrive, SnapIntegrator, SnapLock, SnapManager, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapValidator, SnapVault, StorageGRID, Tech OnTap, Unbound Cloud, and WAFL are trademarks or registered trademarks of NetApp, Inc., in the United States, and/or other countries. A current list of NetApp trademarks is available on the web at <http://www.netapp.com/us/legal/netapptmlist.aspx>.

Cisco and the Cisco logo are trademarks of Cisco in the U.S. and other countries. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such.

How to send comments about documentation and receive update notification

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. 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 [29](#)
 - access-based enumeration
 - See* ABE
 - adding
 - capacity disks to aggregates [12](#)
 - export rules [23](#)
 - rules to export policies [23](#)
 - SSDs to aggregates [14, 18](#)
 - VLAN interfaces [7](#)
 - aggregates
 - adding capacity disks to [12](#)
 - creating storage pools [15, 20](#)
 - provisioning cache by adding SSDs [14, 18](#)
 - workflow for assigning disks to existing [12](#)
 - all-share cache
 - creating for BranchCache configuration [28](#)
 - assigning
 - export policies to qtrees [22](#)
- B**
- BranchCache
 - enabling [29](#)
 - setting up [28](#)
 - broadcast domains
 - creating [9](#)
- C**
- cache
 - increasing for Flash Pool aggregates by adding SSDs [19](#)
 - provisioning to aggregates by adding SSDs [14, 18](#)
 - caching
 - configuring BranchCache for [28](#)
 - CIFS
 - setting up [27](#)
 - CIFS servers
 - configuring BranchCache on [28](#)
 - client access
 - adding rules to export policies [23](#)
 - creating export policies [23](#)
 - creating export policy for [23](#)
 - setting up CIFS [27](#)
 - comments
 - how to send feedback about documentation [33](#)
 - configuring
 - BranchCache [28](#)
 - continuous availability
 - enabling or disabling for shares [29](#)
 - creating
 - broadcast domains [9](#)
 - export policies [23](#)
 - interface groups [8](#)
 - network interfaces [10](#)
 - qtrees [21](#)
 - storage pools [15, 20](#)
 - subnets [9](#)
- D**
- Data ONTAP
 - workflow for setting up a network [7](#)
 - dedicated SSDs
 - adding to Flash Pool aggregates [19](#)
 - DHCP
 - assigning IP addresses to Service Processors [25](#)
 - disks
 - increasing the size of aggregates [12](#)
 - documentation
 - how to receive automatic notification of changes to [33](#)
 - how to send feedback about [33](#)
 - domains
 - creating broadcast [9](#)
- E**
- editing
 - Service Processor settings [26](#)
 - share settings [29](#)
 - editing shares
 - to enable BranchCache [27](#)
 - export policies
 - adding export rules to [23](#)
 - adding rules [23](#)
 - assigning to qtrees [22](#)
 - creating [23](#)
 - workflow for assigning to qtrees [21](#)

- export rules
 - adding to export policies [23](#)
- exporting
 - qtrees [22](#)
- express guides
 - System Manager workflows in [5](#)

F

- feedback
 - how to send comments about documentation [33](#)
- Flash Pool aggregates
 - creating storage pools [15](#), [20](#)
 - increasing the size by adding SSDs [19](#)
 - workflow for provisioning cache [17](#)
- flowcharts
 - Service Processor configuration workflow [25](#)
 - setting up a network [7](#)
 - setting up the BranchCache configuration [27](#)
 - workflow for assigning disks to existing aggregates [12](#)
 - workflow for assigning export policies to qtrees [21](#)
 - workflow for provisioning SSD cache [17](#)

H

- hash stores
 - specifying path and maximum size for BranchCache configuration [28](#)
- HDDs
 - adding to aggregates [12](#)

I

- ifgroups
 - See* interface groups
- information
 - how to send feedback about improving documentation [33](#)
- intercluster connectivity
 - creating network interfaces for [10](#)
- interface groups
 - creating [8](#)
- IP addresses
 - assigning to multiple Service Processors [25](#)
 - editing for a Service Processor [26](#)

M

- modifying
 - Service Processor settings [26](#)
 - the share settings [29](#)

N

- network
 - creating subnets [9](#)
 - creating VLANs [7](#)
- network interfaces
 - creating [10](#)
 - creating VLANs [7](#)
- network mask
 - editing for a Service Processor [26](#)
- non-root aggregates
 - adding capacity disks to [12](#)

O

- oplocks
 - enabling [29](#)

P

- per-share cache
 - creating for BranchCache configuration [28](#)
- policies
 - assigning export, to qtrees [22](#)

Q

- qtrees
 - assigning an export policy to [22](#)
 - creating [21](#)
 - exporting [22](#)
 - workflow for assigning export policies [21](#)

R

- root aggregates
 - adding capacity disks to [12](#)

S

- server keys
 - specifying for BranchCache configuration [28](#)
- Service Processors

36 | Cluster Management Workflows for OnCommand System Manager

- assigning IP addresses globally [25](#)
- configuration workflow [25](#)
- editing the settings [26](#)
- setting up
 - BranchCache [28](#)
 - CIFS [27](#)
- shares
 - editing permissions and options [29](#)
- SMB
 - configuring BranchCache on [28](#)
- SPs
 - See* Service Processors
- SSDs
 - combining to create storage pools [15, 20](#)
 - dedicated SSDs [14, 18](#)
 - increasing the size of Flash Pool aggregates by adding [19](#)
 - provisioning cache by adding [14, 18](#)
 - storage pools [14, 18](#)
 - workflow for provisioning cache [17](#)
- storage pools
 - adding SSDs as [19](#)
 - creating [15, 20](#)
- subnets
 - assigning IP addresses to Service Processors [25](#)
 - creating [9](#)
 - specifying the IP address for the network interface [10](#)
- suggestions
 - how to send feedback about documentation [33](#)
- SVMs
 - creating network interfaces for managing [10](#)
 - setting up BranchCache on [28](#)
- symbolic links

- editing the settings of [29](#)
- System Manager
 - workflows in the 8.3 release [5](#)

T

- twitter
 - how to receive automatic notification of documentation changes [33](#)

V

- VLAN interfaces
 - creating [7](#)
- volumes
 - creating an export policy for client access to [23](#)
 - creating qtrees [21](#)

W

- workflow guides
 - contents of this guides [5](#)
- workflows
 - assigning disks to existing aggregates [12](#)
 - configuring a Service Processor [25](#)
 - in System Manager 8.3 [5](#)
 - providing data access to qtrees by using export policies [21](#)
 - provisioning SSD cache through Flash pool aggregates [17](#)
 - setting up a network [7](#)
 - setting up the BranchCache configuration [27](#)