



ONTAP® 9

Cluster Expansion Express Guide

May 2021 | 215-11193_2021-05_en-us
doccomments@netapp.com

Updated for ONTAP 9.9.1

 **NetApp**®

Contents

Deciding whether to use the Cluster Expansion Express Guide.....3

Cluster expansion workflow.....4

- Verifying the cluster is ready for expansion..... 5
 - Verifying the planned configuration..... 5
 - Gathering the required network information..... 6
 - Adding or replacing switches..... 7
- Preparing the cluster for expansion..... 7
 - Adding node-locked licenses..... 8
 - Verifying the health of the system..... 8
 - Backing up the cluster configuration..... 9
 - Generating an AutoSupport message about starting expansion..... 9
- Installing the new nodes..... 9
 - Installing the controllers..... 10
 - Configuring node-management LIFs..... 11
 - Upgrading or downgrading the nodes..... 11
 - Ensuring hardware-level HA is enabled..... 12
 - Adding nodes to a cluster using System Manager..... 13
 - Joining nodes to the cluster using the CLI..... 13
- Completing the expansion..... 15
 - Configuring the node details in System Manager..... 15
 - Configuring AutoSupport on the new nodes..... 16
 - Configuring the Service Processor network..... 16
 - Validating the configuration of the expanded cluster..... 17
 - Generating an AutoSupport message about completing expansion..... 18
 - Updating LUN paths for the new nodes..... 19

Where to find additional information..... 20

Copyright, trademark, and machine translation..... 21

- Copyright..... 21
- Trademark..... 21
- Machine translation..... 21

Deciding whether to use the Cluster Expansion Express Guide

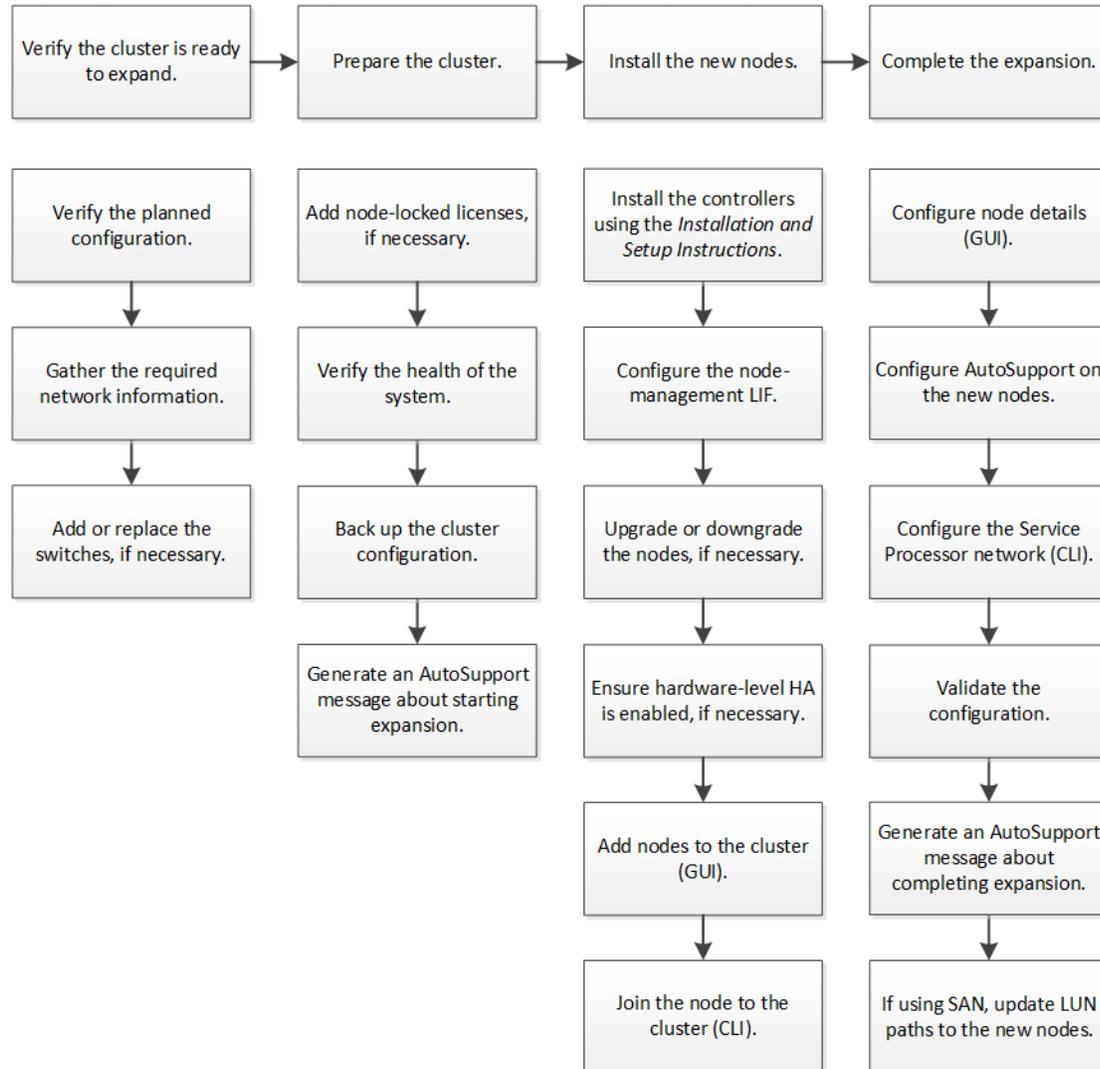
This guide describes how to quickly and nondisruptively expand an existing cluster by adding an HA pair. A larger cluster increases performance and storage capacity available in the cluster.

You should use this guide only if the following is true:

- The existing cluster meets the following requirements:
 - It is running ONTAP 9.
 - It contains at least two nodes.
Although examples in this guide use a two-node cluster, this guide also applies to clusters with more than two nodes.
If you want to add a node to a single-node cluster, you must follow a different procedure.
[Adding a second controller to create an HA pair](#)
 - It does not use IPv6 addressing or Storage Encryption.
 - It is not a MetroCluster configuration.
- The controller modules that you plan to add meet the following requirements:
 - If they are not new, they have been wiped clean, are no longer part of a cluster, and are ready to be added to the new cluster.
 - They support ONTAP 9.
 - They are running a version of the ONTAP 9 release family.
- When completing ONTAP configuration tasks with System Manager, you want to use the *legacy* OnCommand System Manager UI for ONTAP 9.7 and earlier releases, not the ONTAP System Manager UI for ONTAP 9.7 and later.
[ONTAP System Manager documentation](#)
- You want to use best practices, not explore every available option.
- You do not want to read a lot of conceptual background.

Cluster expansion workflow

Adding two nodes to an existing cluster involves verifying that the cluster is ready for expansion, preparing the cluster, installing the new nodes, and completing the expansion.



Steps

1. [Verifying the cluster is ready for expansion](#) on page 5
Before you start expanding a cluster, you must verify the planned configuration, gather the required network information, and add or replace switches, if necessary.
2. [Preparing the cluster for expansion](#) on page 7
To prepare a cluster for expansion, you must add node-locked licenses, verify the system health, back up the cluster's configuration, and generate an AutoSupport message.
3. [Installing the new nodes](#) on page 9
After the cluster is prepared, you must install the controllers and configure the node-management LIFs. If the controllers are not running the same ONTAP version as the existing cluster, or are repurposed and lack hardware-level HA, you must address those issues in Maintenance mode. Finally, you can join the nodes to the cluster.
4. [Completing the expansion](#) on page 15

After both nodes are joined to the cluster, you must finish configuring the newly added nodes by configuring AutoSupport and completing the SP network. You then validate the expanded cluster and generate an AutoSupport message to complete the expansion. If the cluster uses SAN, you should update LUN paths.

Verifying the cluster is ready for expansion

Before you start expanding a cluster, you must verify the planned configuration, gather the required network information, and add or replace switches, if necessary.

Steps

1. [Verifying the planned configuration](#) on page 5
Before you expand a cluster, you must ensure the following: the planned configuration is supported, the required licenses exist, the site is ready, the cluster switches support the expansion, and the existing nodes are using the same version of ONTAP 9.
2. [Gathering the required network information](#) on page 6
Before you expand a cluster, you must obtain networking information required to later configure the node-management LIFs and the Service Processor IP addresses for both of the nodes.
3. [Adding or replacing switches](#) on page 7
Before you expand the cluster, you must ensure that the cluster switches support the expanded configuration. If the cluster is switchless, you must add switches. If the existing switches do not have enough ports available to support the new configuration, you must replace the switches.

Verifying the planned configuration

Before you expand a cluster, you must ensure the following: the planned configuration is supported, the required licenses exist, the site is ready, the cluster switches support the expansion, and the existing nodes are using the same version of ONTAP 9.

Before you begin

You must have two sets of credentials—the user name and password required to log in to the cluster as an administrator, and the user name and password required to log in to the NetApp Support Site.

Steps

1. Verify the planned configuration:
 - a. Verify that the platform of the new controllers can be mixed with the cluster's existing controllers.
 - b. Verify that the expanded cluster does not exceed the system limits for the platforms.
[NetApp Hardware Universe](#)
 - c. If your cluster is configured for SAN, verify that the expanded cluster does not exceed the configuration limits for FC, FCoE, and iSCSI.
[SAN configuration](#)

If these requirements are not met, you cannot proceed with the expansion.

2. Ensure that licenses cover the new nodes:
 - a. On the existing cluster, view the licenses by using the `system license show` command.

```
cluster1::> system license show

Serial Number: 9-99-999999
Owner: cluster1
Package      Type      Description      Expiration
-----
Base         site     Cluster Base License -
NFS          license  NFS License      -
CIFS         license  CIFS License     -
...
```

- b. Review the output to identify the node-locked licenses (identified by the type `license`) that will be required for the additional nodes.
- c. Ensure that the licenses that are included with the additional nodes are consistent with the cluster's existing node-locked licenses.

NetApp Software License Search

If you do not have the required licenses for the additional nodes, you must purchase additional licenses before you proceed.

3. Verify that the site is ready for all the new equipment.

NetApp Hardware Universe

If the site is not ready, you must prepare the site before continuing with the expansion.

4. Verify that the existing switches support the additional controllers.

NetApp Hardware Universe

If the cluster is switchless or if the existing switches do not support the additional nodes, you must obtain cluster switches, which you can install later in the expansion process.

5. Verify that all nodes in the existing cluster are running the same version of ONTAP 9—including the same minor release and patch, if applicable—by using the `cluster image show` command.

```
cluster1::> cluster image show
Node      Current Version      Installation
-----
cluster1-1  8.3RC1      12/15/2014 17:37:26
cluster1-2  8.3RC1      12/15/2014 17:37:42
2 entries were displayed.
```

You should make note of the version of ONTAP software for reference later in this workflow.

Gathering the required network information

Before you expand a cluster, you must obtain networking information required to later configure the node-management LIFs and the Service Processor IP addresses for both of the nodes.

Steps

1. Obtain the following details to configure two node-management LIFs—one for each of the nodes that you plan to add to the cluster:
 - IP address
 - Network mask
 - Gateway
 - Port
2. If your site typically has DNS entries for node-management LIFs, ensure that DNS entries are created for the new nodes.
3. Determine whether the cluster uses automatic or manual network configuration for the SP by using the `system service-processor network auto-configuration show` command.

If a subnet name is displayed in either the `SP IPv4 Subnet Name` or `SP IPv6 Subnet Name` column, the cluster is using automatic SP networking. If both columns are blank, the cluster is using manual SP networking.

In the following output, the `sub1 subnet` indicates that `cluster1 SP` uses automatic network configuration:

```
cluster1::> system service-processor network auto-configuration show
Cluster Name      SP IPv4 Subnet Name      SP IPv6 Subnet Name
```

```
-----  
cluster1          sub1          -  
-----
```

In the following output, the blank subnet fields indicate that cluster1 SP uses manual network configuration:

```
cluster1::> system service-processor network auto-configuration show  
Cluster Name      SP IPv4 Subnet Name      SP IPv6 Subnet Name  
-----  
cluster1          -                          -  
-----
```

4. Depending on the SP network configuration, perform one of the following actions:

- If the SP uses manual network configuration, obtain two IP addresses that you will use later when configuring SP on the new nodes.
- If the SP uses automatic network configuration, verify that the subnet used by the SP has available IP addresses for the two new nodes by using the `network subnet show` command.

In the following output, the sub1 subnet has 2 addresses available:

```
cluster1::> network subnet show  
IPspace: Default  
Subnet  
Name      Subnet      Broadcast Domain      Gateway      Avail/Total      Ranges  
-----  
sub1      10.53.33.1/18      Default      10.53.0.1      2/4      10.53.33.3-10.53.33.6  
...
```

Adding or replacing switches

Before you expand the cluster, you must ensure that the cluster switches support the expanded configuration. If the cluster is switchless, you must add switches. If the existing switches do not have enough ports available to support the new configuration, you must replace the switches.

Choices

- If the cluster is currently a two-node switchless cluster, migrate the cluster to a two-node switched cluster using the type of switch you prefer.
[Migrating to a two-node switched cluster with Cisco cluster switches](#)
[Migrating to a two-node switched cluster with NetApp CN1610 cluster switches](#)
- If the existing switches do not have enough ports available to support the future configuration, replace the switches by using the appropriate replacement procedure.
[NetApp Documentation: Cluster, Management and Storage Switches](#)

Preparing the cluster for expansion

To prepare a cluster for expansion, you must add node-locked licenses, verify the system health, back up the cluster's configuration, and generate an AutoSupport message.

Steps

1. [Adding node-locked licenses](#) on page 8
If the cluster has features that use node-locked licenses (which entitle only specific nodes to the licensed functionality), you must ensure that node-locked licenses are installed for the new nodes. You should add the licenses before the nodes are joined to the cluster.
2. [Verifying the health of the system](#) on page 8
Before you expand a cluster, you must verify that all components of the cluster are healthy by running the Config Advisor tool and running several ONTAP CLI commands.
3. [Backing up the cluster configuration](#) on page 9
Before you expand a cluster, you should use advanced privilege to create a backup file to save the cluster configuration information and optionally save the node configurations.
4. [Generating an AutoSupport message about starting expansion](#) on page 9

Immediately before you expand a cluster, you should send an AutoSupport message to indicate that you are about to start the expansion process. The message informs internal and external support staff about expansion and acts as a timestamp for any troubleshooting that might be required later.

Adding node-locked licenses

If the cluster has features that use node-locked licenses (which entitle only specific nodes to the licensed functionality), you must ensure that node-locked licenses are installed for the new nodes. You should add the licenses before the nodes are joined to the cluster.

Steps

1. Add each license key by using the `system license add` command.

```
cluster1::> system license add -license-code AAAAAAAAAAAAAA
```

2. View the existing licenses by using the `system license show` command.

```
cluster1::> system license show

Serial Number: 9-99-999999
Owner: cluster1
Package      Type      Description      Expiration
-----
Base         site     Cluster Base License -
NFS          license  NFS License      -
CIFS         license  CIFS License     -
...
```

3. Review the output to ensure that a node-locked license is displayed for all serial numbers, including serial numbers for existing and new nodes.

Verifying the health of the system

Before you expand a cluster, you must verify that all components of the cluster are healthy by running the Config Advisor tool and running several ONTAP CLI commands.

Steps

1. Verify that you have the latest version of Config Advisor:

- If you do not have Config Advisor on your laptop, download it.
[NetApp Downloads: Config Advisor](#)
- If you have Config Advisor, start it, click **Help > Check for Updates**, and follow the prompts to upgrade it if necessary.

Note: Do not uninstall the previous version of the tool or delete the data folder during the upgrade. The tool uninstalls the previous version and replaces it with the latest version. It renames the data folder as the latest folder and retains all of the contents in the folder.

2. Verify the cabling and configuration by running Config Advisor:

- a. Connect your laptop to the management network for the cluster.
- b. Click **Collect Data**.
Config Advisor displays any problems found.
- c. If problems are found, correct them and run the tool again.

3. Optional: Check the health of the system with the following commands:

- a. Verify that the cluster is in a healthy state by using the `system health status show` command and verifying that the Status is `ok`.

```
cluster1::> system health status show
Status
-----
ok
```

- b. Verify that all nodes in the cluster are in a healthy state by using the `cluster show` command and verifying that the Health of each node is `true`.

```
cluster1::> cluster show
Node          Health Eligibility
-----
cluster1-1    true   true
cluster1-2    true   true
2 entries were displayed.
```

Backing up the cluster configuration

Before you expand a cluster, you should use advanced privilege to create a backup file to save the cluster configuration information and optionally save the node configurations.

Steps

1. Set the privilege level to advanced by using the `set -privilege advanced` command.
2. Create a backup file of the cluster configuration by using the `system configuration backup create` command with the `-backup-type cluster` parameter.

```
cluster1::*> system configuration backup create -node cluster1-1 -backup-name clusterbeforeexpansion.7z -backup-type cluster
[Job 5573] Job is queued: Cluster Backup OnDemand Job.
```

3. Optional: Create a backup file of each node's configuration by using the `system configuration backup create` command with the `-backup-type node` parameter for each node.
4. Return the privilege level to admin by using the `set -privilege admin` command.

Generating an AutoSupport message about starting expansion

Immediately before you expand a cluster, you should send an AutoSupport message to indicate that you are about to start the expansion process. The message informs internal and external support staff about expansion and acts as a timestamp for any troubleshooting that might be required later.

Before you begin

AutoSupport must be set up.

Step

For each node in the cluster, send an AutoSupport message by using the `system node autosupport invoke` command.

```
cluster1::> system node autosupport invoke -node * -message "cluster expansion started" -type all
The AutoSupport was successfully invoked on node "cluster1-1". To view the status of the AutoSupport, use the "system node autosupport history show" command.
Note: It may take several minutes for the AutoSupport to appear in the history list.
The AutoSupport was successfully invoked on node "cluster1-2". To view the status of the AutoSupport, use the "system node autosupport history show" command.
Note: It may take several minutes for the AutoSupport to appear in the history list.
2 entries were acted on.
```

Installing the new nodes

After the cluster is prepared, you must install the controllers and configure the node-management LIFs. If the controllers are not running the same ONTAP version as the existing cluster, or are

repurposed and lack hardware-level HA, you must address those issues in Maintenance mode. Finally, you can join the nodes to the cluster.

Steps

1. [Installing the controllers](#) on page 10

When you install controllers that will be added to an existing cluster, you must follow the first three steps of the appropriate *Installation and Setup Instructions*.

2. [Configuring node-management LIFs](#) on page 11

After the controller modules are physically installed, you can power on each one and configure its node-management LIF.

3. [Upgrading or downgrading the nodes](#) on page 11

Before joining the newly installed nodes to the cluster, you must ensure that they are running the same version of ONTAP that the cluster is running. If the nodes are running a different version, you must upgrade or downgrade the nodes to match the cluster.

4. [Ensuring hardware-level HA is enabled](#) on page 12

If the newly installed controller modules are reused—not new—you must enter Maintenance mode and ensure that their HA state is set to HA.

5. [Adding nodes to a cluster using System Manager](#) on page 13

You can use System Manager to increase the size and capabilities of your storage system by adding nodes to an existing cluster. This feature is automatically enabled in System Manager when the effective cluster version is ONTAP 9.2.

6. [Joining nodes to the cluster using the CLI](#) on page 13

When the newly installed controller modules are ready, you can add each one to the cluster by using the `cluster setup` command.

Installing the controllers

When you install controllers that will be added to an existing cluster, you must follow the first three steps of the appropriate *Installation and Setup Instructions*.

About this task

Note: As of ONTAP 9.0, HA mode is enabled by default on new hardware.

Steps

1. Obtain the *Installation and Setup Instructions* for the FAS model number of the controller module that you plan to add to the cluster.

- For a new controller module, the document is available in the box.
- For a repurposed controller module, you can download the document.

[NetApp Documentation](#)

2. Follow the *Prepare for installation* section with the following exceptions:

- You can skip any instructions about downloading software or a worksheet.
- You must provide a serial console connection even if it is not mentioned in the *Installation and Setup Instructions*.

You require a serial console because you must use the nodeshell CLI to configure node-management LIFs.

If the ONTAP section does not mention the serial console, you can see the 7-mode section.

3. Follow the *Install hardware* section.

4. Follow the *Cable storage* section.

5. Skip most of the *Complete System Setup* section with the following exceptions:

- If instructed to, you must power on all disk shelves and check IDs.
- You must cable the serial console so that you can access the node.
If the ONTAP section does not mention the serial console, you can see the 7-mode section.

6. Skip the *Complete configuration* section.

Configuring node-management LIFs

After the controller modules are physically installed, you can power on each one and configure its node-management LIF.

About this task

You must perform this procedure on both the nodes.

Steps

1. Access the controller module through the serial console.
2. Power on the controller module, and wait while the node boots and the Cluster Setup wizard automatically starts on the console.

```
Welcome to the cluster setup wizard.
```

```
You can enter the following commands at any time:  
"help" or "?" - if you want to have a question clarified,  
"back" - if you want to change previously answered questions, and  
"exit" or "quit" - if you want to quit the cluster setup wizard.  
Any changes you made before quitting will be saved.
```

```
You can return to cluster setup at any time by typing "cluster setup".  
To accept a default or omit a question, do not enter a value.
```

3. Follow the prompts in the web-based Cluster Setup wizard to configure a node management LIF using the networking information you gathered earlier.

4. Type

exit

after node management LIF configuration is complete to exit the setup wizard and complete the administration tasks.

```
Use your web browser to complete cluster setup by accessing  
https://10.63.11.29
```

```
Otherwise, press Enter to complete cluster setup using the command line interface:  
exit
```

5. Log in to the node as the **admin** user, which does not require a password.

```
Tue Mar 4 23:13:33 UTC 2015  
login: admin  
*****  
* This is a serial console session. Output from this *  
* session is mirrored on the SP console session.      *
```

6. Repeat the entire procedure for the second newly installed controller module.

Upgrading or downgrading the nodes

Before joining the newly installed nodes to the cluster, you must ensure that they are running the same version of ONTAP that the cluster is running. If the nodes are running a different version, you must upgrade or downgrade the nodes to match the cluster.

Steps

1. Determine the version of ONTAP installed on the cluster:

```
cluster image show
```

2. View the current version of ONTAP on the nodes:

- a. On the first node, view the software version:

system node image show

```
::*> system node image show
Node      Image  Is Default  Is Current  Version  Install
-----  -----  - - - - -  - - - - -  -----  -----
localhost image1  false  false   9.3      MM/DD/YYYY TIMESTAMP
          image1  true   true    9.3      MM/DD/YYYY TIMESTAMP
2 entries were displayed.
```

- b. Repeat the previous step for the second node.
- 3. Compare the versions of ONTAP on the cluster and the nodes, and perform either of the following actions:
 - If the versions of ONTAP on the cluster and the nodes are the same, no upgrade or downgrade is needed.
 - If the versions of ONTAP on the cluster and the nodes are different, use the procedures in the following guides to upgrade or downgrade the version of ONTAP on the nodes to match the version of ONTAP on the cluster.

Software express upgrade

Upgrade, revert, or downgrade

Ensuring hardware-level HA is enabled

If the newly installed controller modules are reused—not new—you must enter Maintenance mode and ensure that their HA state is set to HA.

About this task

If you are using new controller modules, you can skip this procedure because HA is enabled by default. Otherwise, you must perform this procedure on both the nodes.

Steps

1. On the first node, enter Maintenance mode:
 - a. Exit the nodeshell by entering `halt`.
The `LOADER` prompt is displayed.
 - b. Enter Maintenance mode by entering `boot_ontap maint`.
After some information is displayed, the Maintenance mode prompt is displayed.
2. In Maintenance mode, ensure that the controller module and chassis are in HA state:
 - a. Display the HA state of the controller module and chassis by entering `ha-config show`.
 - b. If the displayed state of the controller is not HA, enter `ha-config modify controller ha`.
 - c. If the displayed state of the chassis is not HA, enter `ha-config modify chassis ha`.
 - d. Verify that HA is enabled on both the controller module and chassis by entering `ha-config show`.
3. Return to ONTAP:
 - a. Enter `halt` to exit Maintenance mode.
 - b. Boot ONTAP by entering `boot_ontap`
 - c. Wait while the node boots and the Cluster Setup wizard automatically starts on the console.
 - d. Press Enter four times to accept the existing settings for the node-management LIF.
 - e. Log in to the node as the `admin` user, which does not require a password.
4. Repeat this procedure on the other node that you are adding to the cluster.

Adding nodes to a cluster using System Manager

You can use System Manager to increase the size and capabilities of your storage system by adding nodes to an existing cluster. This feature is automatically enabled in System Manager when the effective cluster version is ONTAP 9.2.

Before you begin

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

Step

Add the new compatible nodes to the cluster:

If you are...	Do this...
Not logged in to System Manager	<p>a. Log in to System Manager.</p> <p>Note: The new compatible nodes are automatically detected by System Manager at login. System Manager prompts you to add the new compatible nodes to the cluster.</p> <p>b. Click Add Nodes to Cluster.</p> <p>c. Modify the name of the nodes.</p> <p>d. Specify the node licenses.</p> <p>e. Click Submit and Proceed.</p>
Logged in to System Manager	<p>a. Depending on the System Manager version that you are running, perform one of the following steps:</p> <ul style="list-style-type: none">• ONTAP 9.4 or earlier: Click Configuration > Cluster Expansion.• Starting with ONTAP 9.5: Click Configuration > Cluster > Expansion <p>System Manager searches for newly added nodes. If any warnings are displayed, you must fix them before proceeding. If new compatible nodes are discovered, proceed to the next step.</p> <p>b. Modify the name of the nodes.</p> <p>c. Specify the node licenses.</p> <p>d. Click Submit and Proceed.</p>

Joining nodes to the cluster using the CLI

When the newly installed controller modules are ready, you can add each one to the cluster by using the `cluster setup` command.

About this task

- You must perform this procedure on both nodes.
- You must join each node one at a time, not concurrently.

Steps

1. Start the Cluster Setup wizard by using the `cluster setup` command at the CLI prompt.

```
::> cluster setup
Welcome to the cluster setup wizard...

Use your web browser to complete cluster setup by accessing
https://10.63.11.29
```

Otherwise, press Enter to complete cluster setup using the command line interface:

Note: For instructions using the GUI-based cluster setup wizard, see the topic *Adding nodes to the cluster using System Manager* in this guide.

2. Press Enter to use the CLI to complete this task. When prompted to create a new cluster or join an existing one, enter **join**.

```
Do you want to create a new cluster or join an existing cluster? {create, join}:  
join
```

3. When prompted with the existing cluster interface configuration, press

Enter
to accept it.

Existing cluster interface configuration found:

Port	MTU	IP	Netmask
ela	9000	169.254.87.75	255.255.0.0

```
Do you want to use this configuration? {yes, no} [yes]:
```

4. Follow the prompts to join the existing cluster.

```
Step 1 of 3: Join an Existing Cluster  
You can type "back", "exit", or "help" at any question.  
  
Enter the name of the cluster you would like to join [cluster1]: cluster1  
  
Joining cluster cluster1  
  
Starting cluster support services ..  
  
This node has joined the cluster cluster1.  
  
Step 2 of 3: Configure Storage Failover (SFO)  
You can type "back", "exit", or "help" at any question.  
  
SFO will be enabled when the partner joins the cluster.  
  
Step 3 of 3: Set Up the Node  
  
Cluster setup is now complete.
```

The node is automatically renamed to match the name of the cluster.

5. On the cluster, verify that the node is part of the cluster by using the `cluster show` command.

```
cluster1::> cluster show  
Node           Health  Eligibility  
-----  
cluster1-1     true    true  
cluster1-2     true    true  
cluster1-3     true    true  
3 entries were displayed.
```

6. Repeat steps 1 on page 13 through 5 on page 14 for the second newly installed controller module.

The Cluster Setup wizard differs on the second node in the following ways:

- It defaults to joining the existing cluster because its partner is already part of a cluster.
- It automatically enables storage failover on both nodes.

7. Verify that storage failover is enabled and possible by using the `storage failover show` command.

The following output shows that storage failover is enabled and possible on all nodes of the cluster, including the newly added nodes:

```
cluster1::> storage failover show  
Node           Partner      Takeover  
Possible State  
-----  
cluster1-1     cluster1-2   true    Connected to cluster1-2
```

```
cluster1-2    cluster1-1    true    Connected to cluster1-1
cluster1-3    cluster1-4    true    Connected to cluster1-3
cluster1-4    cluster1-3    true    Connected to cluster1-4
4 entries were displayed.
```

Completing the expansion

After both nodes are joined to the cluster, you must finish configuring the newly added nodes by configuring AutoSupport and completing the SP network. You then validate the expanded cluster and generate an AutoSupport message to complete the expansion. If the cluster uses SAN, you should update LUN paths.

Steps

1. [Configuring the node details in System Manager](#) on page 15
You can use System Manager to configure the node management LIF and Service Processor settings for the newly added nodes.
2. [Configuring AutoSupport on the new nodes](#) on page 16
After you add nodes to a cluster, you must configure AutoSupport on the nodes.
3. [Configuring the Service Processor network](#) on page 16
After you expand a cluster, you must configure the Service Processor (SP) network on the new nodes. If the SP uses manual network configuration, you must configure the IP addresses for the SP on the new nodes. If the SP uses automatic network configuration, you must identify the IP addresses that were selected.
4. [Validating the configuration of the expanded cluster](#) on page 17
After you expand the cluster, you must validate the configuration by running Config Advisor and using some commands that verify cluster health and cluster replication rings.
5. [Generating an AutoSupport message about completing expansion](#) on page 18
After you expand a cluster, you should send an AutoSupport message to indicate that the expansion process is complete. This message communicates to internal and external support staff that the expansion is complete and acts as a timestamp for any troubleshooting that might be required later.
6. [Updating LUN paths for the new nodes](#) on page 19
If your cluster is configured for SAN, you must create SAN LIFs on the newly added nodes and then update paths.

Configuring the node details in System Manager

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

Before you begin

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

Steps

1. Configure node management:
 - a. Enter the IP address in the **IP Address** field.
 - b. Select the port for node management in the **Port** field.
 - c. Enter the netmask and gateway details.
2. Configure Service Processor settings:
 - a. Select the **Override defaults** check box to override the default values.

2. Verify that the SP network is configured correctly on both the new nodes by using the `system service-processor network show` command for each node.

The status should be `succeeded`. Verification is required in all situations. Even if the SP network was automatically configured, you should verify that it was configured successfully, and you must determine which IP addresses were assigned.

The following output indicates that both the `cluster1-3` and the `cluster1-4` nodes have successful SP network setup:

```
cluster1::> system service-processor network show -node cluster1-3
Address
Node      Status      Family      Link State  IP Address
-----
cluster1-3  online      IPv4        up          192.168.123.98

          DHCP: none
          MAC Address: 00:a0:98:43:a1:1e
          Network Gateway: 10.60.172.1
          Network Mask (IPv4 only): 255.255.255.0
          Prefix Length (IPv6 only): -
          IPv6 RA Enabled: -
          Subnet Name: -
          SP Network Setup Status: succeeded
          ...

cluster1::> system service-processor network show -node cluster1-4
Address
Node      Status      Family      Link State  IP Address
-----
cluster1-4  online      IPv4        up          192.168.123.99

          DHCP: none
          MAC Address: 00:a0:98:43:a1:1e
          Network Gateway: 10.60.172.1
          Network Mask (IPv4 only): 255.255.255.0
          Prefix Length (IPv6 only): -
          IPv6 RA Enabled: -
          Subnet Name: -
          SP Network Setup Status: succeeded
          ...
```

3. If your site typically has DNS entries for the SP network, verify that the DNS entries are created for the new nodes.

Validating the configuration of the expanded cluster

After you expand the cluster, you must validate the configuration by running Config Advisor and using some commands that verify cluster health and cluster replication rings.

Steps

1. Check the health of the configuration by running Config Advisor:
 - a. Start Config Advisor, and then click **Collect Data**.
Config Advisor displays any problems found.
 - b. If problems are found, correct them and run the tool again.
2. Ensure that all nodes in the cluster are in a healthy state by using the `cluster show` command.

```
cluster-1::> cluster show
Node      Health  Eligibility
-----
cluster1-1  true   true
cluster1-2  true   true
cluster1-3  true   true
cluster1-4  true   true
4 entries were displayed.
```

3. Ensure that the cluster replication rings have the same epoch, database epoch, and database transaction numbers on all nodes in the cluster:

The easiest way to compare transaction numbers is to view them for one unit name at a time.

- a. Set the privilege level to advanced by using the `set -privilege advanced` command.

- b. View cluster ring information about the first unit name by using the `cluster ring show` command with the `-unitname mgmt` parameter, and verify that all nodes have the same number in the Epoch, DB Epoch, and DB Trnxs columns.

```
cluster-1::*> cluster ring show -unitname mgmt
Node      UnitName Epoch  DB Epoch DB Trnxs Master  Online
-----
cluster1-1
      mgmt    2      2      959    cluster1-1
                        master
cluster1-2
      mgmt    2      2      959    cluster1-2
                        secondary
cluster1-3
      mgmt    2      2      959    cluster1-3
                        master
cluster1-4
      mgmt    2      2      959    cluster1-3
                        secondary
4 entries were displayed.
```

- c. Repeat the command with the `-unitname vldb` parameter.
- d. Repeat the command with the `-unitname vifmgr` parameter.
- e. Repeat the command with the `-unitname bcomd` parameter.
- f. Repeat the command with the `-unitname crs` parameter.
- g. Return the privilege level to admin by using the `set -privilege admin` command.

Generating an AutoSupport message about completing expansion

After you expand a cluster, you should send an AutoSupport message to indicate that the expansion process is complete. This message communicates to internal and external support staff that the expansion is complete and acts as a timestamp for any troubleshooting that might be required later.

Before you begin

AutoSupport must be set up.

Step

For each node in the cluster, send an AutoSupport message by using the `system node autosupport invoke` command.

You must issue the message once for each node in the cluster, including the newly added nodes.

If you added two nodes to a two-node cluster, you must send the message four times.

```
cluster1::> system node autosupport invoke -node * -message "cluster expansion complete" -type all
The AutoSupport was successfully invoked on node "cluster1-1". To view the status
of the AutoSupport, use the "system node autosupport history show" command.
Note: It may take several minutes for the AutoSupport to appear in the history list.
The AutoSupport was successfully invoked on node "cluster1-2". To view the status
of the AutoSupport, use the "system node autosupport history show" command.
Note: It may take several minutes for the AutoSupport to appear in the history list.
The AutoSupport was successfully invoked on node "cluster1-3". To view the status
of the AutoSupport, use the "system node autosupport history show" command.
Note: It may take several minutes for the AutoSupport to appear in the history list.
The AutoSupport was successfully invoked on node "cluster1-4". To view the status
of the AutoSupport, use the "system node autosupport history show" command.
Note: It may take several minutes for the AutoSupport to appear in the history list.
4 entries were acted on.
```

Updating LUN paths for the new nodes

If your cluster is configured for SAN, you must create SAN LIFs on the newly added nodes and then update paths.

About this task

This procedure is required only if the cluster contains LUNs. If the cluster contains only files, you can skip this procedure.

Steps

1. For each storage virtual machine (SVM) in the cluster, create new LIFs on the newly added nodes:

- a. Identify the SVMs that use FC or iSCSI protocols by using the `vserver show` command with the `-fields allowed-protocols` parameter and reviewing the output.

```
cluster1::> vserver show -fields allowed-protocols
vserver allowed-protocols
-----
vs1      cifs,ndmp
vs2      fcp
vs3      iscsi
...
```

- b. For each SVM that uses FC or iSCSI, create at least two data LIFs on each of the newly added nodes by using the `network interface create` command with the `-role data` parameter.

```
cluster1::> network interface create -vserver vs1 -lif lif5 -role data
-data-protocol iscsi -home-node cluster1-3 -home-port e0b
-address 192.168.2.72 -netmask 255.255.255.0
```

- c. For each SVM, verify that it has LIFs on all nodes in the cluster by using the `network interface show` command with the `-vserver` parameter.

2. Update port sets:

- a. Determine whether port sets exist by using the `lun portset show` command.
- b. If you want to make the new LIFs visible to existing hosts, add each new LIF to the port sets by using the `lun portset add` command—once for each LIF.

3. If you use FC or FCoE, update zoning:

- a. Verify that zoning is set up correctly to enable the existing initiator ports on the host to connect to the new target ports on the new nodes.
- b. Update switch zoning to connect the new nodes to existing initiators.
Zoning setup varies depending on the switch that you use.
- c. If you plan to move LUNs to the new nodes, expose the new paths to the hosts by using the `lun mapping add-reporting-nodes` command.

4. On all host operating systems, rescan to discover the newly added paths.

5. Depending on the host operating systems, remove any stale paths.

6. Add or remove paths to your MPIO configuration.

Related information

[SAN configuration](#)

[SAN administration](#)

Where to find additional information

After you expand a cluster, you can start storing data on the new nodes either by creating new volumes on the new nodes or by moving existing data to the new nodes. There are express guides and comprehensive documentation to help you achieve these goals.

Express guides

If you want to create storage virtual machines (SVMs) or volumes on the newly added nodes, you can use any of the express guides about setting up CIFS, NFS, iSCSI, and FC protocols.

[ONTAP 9 Documentation Center](#)

Data protection

If you are using SnapMirror or SnapVault relationships in your cluster to protect your data, see the [Cluster and SVM peering express configuration](#) guide to set up appropriate intercluster LIFs for your new nodes.

Comprehensive guides

If you want to move data to the newly added nodes, you can use the following comprehensive guides:

- [SAN administration](#)
Describes how to configure and manage the iSCSI, FCoE, and FC protocols for clustered SAN environments, including configuration of LUNs, igroups, and targets.
- [Logical storage management](#)
Describes how to manage logical storage resources in clusters, including FlexVol volumes, FlexClone volumes, files and LUNs, and FlexCache volumes, using deduplication, compression, qtrees, and quotas.
- [ONTAP concepts](#)
Describes conceptual information about logical storage resources in clusters, including FlexVol volumes, FlexClone volumes, files and LUNs, and FlexCache volumes, using deduplication, compression, qtrees, and quotas.

Copyright, trademark, and machine translation

Copyright

Copyright © 2021 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.

Data contained herein pertains to a commercial item (as defined in FAR 2.101) and is proprietary to NetApp, Inc. The U.S. Government has a non-exclusive, non-transferrable, non-sublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b).

Trademark

NETAPP, the NETAPP logo, and the marks listed on the NetApp Trademarks page are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.

<https://www.netapp.com/company/legal/trademarks/>

Machine translation

See important information about localized content at [netapp.com](https://www.netapp.com).

<https://www.netapp.com/company/legal/machine-translation/>