




**NetApp®**

## Clustered Data ONTAP® 8.3

### Cluster Expansion Express Guide



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](http://www.netapp.com)  
Feedback: [doccomments@netapp.com](mailto:doccomments@netapp.com)

Part number: 215-09656\_A0  
January 2015



# Contents

<b>Deciding whether to use this guide .....</b>	<b>4</b>
<b>Cluster expansion workflow .....</b>	<b>5</b>
Verifying the cluster is ready for expansion .....	6
Verifying the planned configuration .....	6
Gathering the required network information .....	8
Adding or replacing switches .....	9
Preparing the cluster for expansion .....	10
Adding node-locked licenses .....	10
Verifying the health of the system .....	11
Backing up the cluster configuration .....	13
Generating an AutoSupport message about starting expansion .....	13
Installing the new nodes .....	14
Installing the controllers .....	15
Configuring node-management LIFs .....	16
Upgrading or downgrading the nodes .....	17
Ensuring hardware-level HA is enabled .....	19
Joining the nodes to the cluster .....	20
Completing the expansion .....	22
Configuring AutoSupport on the new nodes .....	23
Configuring the Service Processor network .....	24
Validating the configuration of the expanded cluster .....	25
Generating an AutoSupport message about completing expansion .....	27
Updating LUN paths for the new nodes .....	27
<b>Where to find additional information .....</b>	<b>30</b>
<b>Copyright information .....</b>	<b>31</b>
<b>Trademark information .....</b>	<b>32</b>
<b>How to send comments about documentation and receive update notification .....</b>	<b>33</b>
<b>Index .....</b>	<b>34</b>

## Deciding whether to use this guide

---

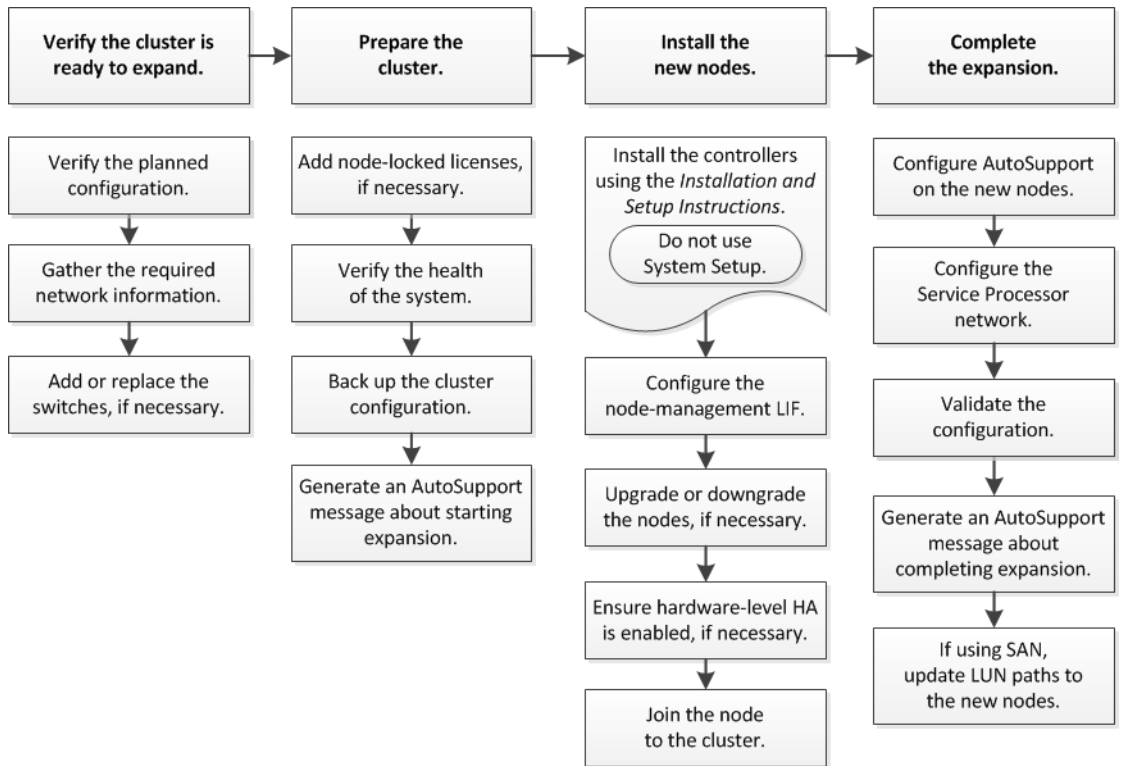
This guide describes how to quickly and nondisruptively expand an existing cluster by adding an HA pair. A larger cluster increases performance and storage capacity available in the cluster.

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

- The existing cluster meets the following requirements:
  - It is running Data ONTAP 8.3.
  - 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.  
[\*Clustered Data ONTAP: Adding a second controller module 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 Data ONTAP 8.3.
  - They are running a version of the Data ONTAP 8.3 release family.
- 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 6  
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 10  
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 14  
After the cluster is prepared, you must install the controllers and configure the node-management LIFs. If the controllers have the wrong version of Data ONTAP 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 22

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 6

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 Data ONTAP 8.3.

#### 2. [Gathering the required network information](#) on page 8

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 9

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 Data ONTAP 8.3.

### Before you begin

You must have the 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 as follows:

- a. Verify that the platform of the new controllers can be mixed with the cluster's existing controllers.

*Clustered Data ONTAP Storage Platform Mixing Rules*

- 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.

*Clustered Data ONTAP 8.3 SAN Configuration Guide*

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.

### Example

```
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 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 Data ONTAP 8.3—including the same minor release and patch, if applicable—by using the `cluster image show` command.

### Example

```
cluster1::> cluster image show
Node           Current          Installation
              Version          Date
-----
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 Data ONTAP for reference later in this workflow.

*[Clustered Data ONTAP 8.3 man page: cluster image show - Display currently running image information](#)*

## 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.

### Example

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                      -
```

### Example

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.

### Example

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.

## 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 10  
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 11  
Before you expand a cluster, you must ensure that all components of the cluster are healthy by running the Config Advisor tool and running several health commands.
3. [Backing up the cluster configuration](#) on page 13  
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 13  
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.

### Example

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

*Clustered Data ONTAP 8.3 man page: [system license add](#) - Add one or more licenses*

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

## Example

```
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      -
...
```

*Clustered Data ONTAP 8.3 man page: [system license show - Display licenses](#)*

- 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 ensure that all components of the cluster are healthy by running the Config Advisor tool and running several health commands.

### Steps

- Ensure 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 the contents in the folder.
- Verify the cabling and configuration by running Config Advisor:
  - Connect your laptop to the management network for the cluster.
  - Click **Collect Data**.  
Config Advisor displays any problems found.
  - If problems are found, correct them and run the tool again.
- Verify that the shelves and disks do not have any issues by using and reviewing the output of the following health commands, which display information that is not checked by Config Advisor:
  - `system health system-connectivity shelf show`
  - `system health node-connectivity shelf show`

## 12 | Cluster Expansion Express Guide

- `system health node-connectivity disk show`

### Example

```
cluster1::> system health system-connectivity shelf show
Shelf UUID                               Shelf Connected Num Status
----- ID ----- Nodes ----- Paths -----
50:0a:09:80:01:b9:73:19                 0 cluster1-1, 4 ok
                                           cluster1-2

cluster1::> system health node-connectivity shelf show
Node Shelf Name Num Num Status
----- Shelf Name ----- Disks Paths -----
cluster1-1 shelf0 24 2 ok
cluster1-1 shelf1 24 2 ok
cluster1-1 shelf2 24 2 ok
...

cluster1::> system health node-connectivity disk show
Node Disk Name Bay Num Status
----- Disk Name ----- Bay Paths -----
cluster1-1 0a.00.01 21 2 ok
cluster1-1 0a.00.02 18 2 ok
cluster1-1 0a.00.03 10 2 ok
...
```

*Clustered Data ONTAP 8.3 man page: [system health system-connectivity shelf show](#) - Show shelf resources and connectivity status*

*Clustered Data ONTAP 8.3 man page: [system health node-connectivity shelf show](#) - Show shelf resources and connectivity status*

*Clustered Data ONTAP 8.3 man page: [system health node-connectivity disk show](#) - Show disk resources and connectivity status*

4. Optional: Check the health of the system with the following commands:
  - a. Ensure that the cluster is in a healthy state by using the `system health status show` command and verifying that the Status is `ok`.

### Example

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

- b. Ensure 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`.

**Example**

```
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.

**Example**

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

*[Clustered Data ONTAP 8.3 man page: system configuration backup create - Create a configuration backup](#)*

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**

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

**Example**

```
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.
```

*Clustered Data ONTAP 8.3 man page: [system node autosupport invoke](#) - Generate and send an AutoSupport message*

## Installing the new nodes

After the cluster is prepared, you must install the controllers and configure the node-management LIFs. If the controllers have the wrong version of Data ONTAP 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 15  
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*, but you must *not* use System Setup to create a cluster.
2. [Configuring node-management LIFs](#) on page 16  
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 17  
Before joining the newly installed nodes to the cluster, you must ensure that they are running the same version of Data 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 19  
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. [Joining the nodes to the cluster](#) on page 20  
When the newly installed controller modules are ready, you can add each one to the cluster by using the `cluster setup` command, which first sets the node in HA mode and then joins it to the cluster.

## 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*, but you must *not* use System Setup to create a cluster.

### 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 clustered Data 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 clustered Data ONTAP section does not mention the serial console, you can see the 7-mode section.  
  
**Attention:** You must not run System Setup software, which will create a new cluster with the new nodes rather than joining the nodes to an existing cluster.
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 Node Setup wizard automatically starts on the console.

### Example

```
Welcome to node setup.

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 setup wizard.
Any changes you made before quitting will be saved.

To accept a default or omit a question, do not enter a value.

This system will send event messages and weekly reports to NetApp Technical
Support.
To disable this feature, enter "autosupport modify -support disable" within 24
hours.
Enabling AutoSupport can significantly speed problem determination and
resolution should a problem occur on your system.
For further information on AutoSupport, see:
http://support.netapp.com/autosupport/
```

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

### Example

```
Type yes to confirm and continue {yes}: yes

Enter the node management interface port [e0a]: e0M
Enter the node management interface IP address: 172.17.178.29
Enter the node management interface netmask: 255.255.255.0
Enter the node management interface default gateway: 172.17.178.1
A node management interface on port e0M with IP address 172.17.178.29 has been created.

This node has its management address assigned and is ready for cluster setup.

To complete cluster setup after all nodes are ready, download and run the System
Setup utility from the NetApp Support Site and use it to discover the configured nodes.

For System Setup, this node's management address is: 172.17.178.29.
```



You must ignore the onscreen instructions to run the System Setup utility because they apply only to a new cluster, not to joining an existing cluster.

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

### Example

```
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.     *
```

5. 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 Data 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.

### Before you begin

You must know the version of Data ONTAP installed on the cluster, which you obtained earlier in the workflow. If you do not know the version, you can use the `cluster image show` command on the cluster.

### Steps

1. View the current version of Data ONTAP on the nodes:
  - a. On the first node, view the software version by using the `system node image show` command.

### Example

```
::> system node image show
      Is      Is
Node  Image  Default Current Version          Install
-----
localhost
      image1 false  false  8.3RC1          12/19/2014 03:01:02
      image1 true   true   8.3RC1          01/05/2015 17:36:06
2 entries were displayed.
```

[Clustered Data ONTAP 8.3 man page: system node image show - Display software image information](#)

- b. Repeat the previous step for the second new node.
2. Compare the versions of Data ONTAP on the cluster and the nodes, and perform one of the following actions:

## 18 | Cluster Expansion Express Guide

- If the versions are the same, skip the remaining steps in this procedure.
- If the versions are different, install a different version of Data ONTAP on the nodes by following the remaining steps this procedure.

### 3. Obtain the appropriate software image:

- a. Identify an HTTP or FTP server from which the software image will be served.
- b. Locate the target Data ONTAP software in the **Software Downloads** area of the NetApp Support Site, ensuring that you select the software image for the platform module.

#### *NetApp Downloads*

- c. Copy the software image (for example, `830_q_image.tgz`) to the HTTP or FTP server.
- d. Identify the IP address of the HTTP or FTP server, because the server's host name is not resolvable by the node.

### 4. Install the software on the first node:

- a. On the first node, install the software image by using the `system node image update` command.

#### **Example**

```
::> system node image update -package http://10.56.94.15/ONTAP_update/830_q_image.tgz -  
replace-package true -background true
```

*Clustered Data ONTAP 8.3 man page: `system node image update` - Perform software image upgrade/downgrade*

- b. Monitor the progress of the installation by using the `system node image show-update-progress` command, and wait for the installation to finish.

*Clustered Data ONTAP 8.3 man page: `system node image show-update-progress` - Show progress information for a currently running update*

- c. When the installation is finished, restart the node by using the `system node reboot` command.

The `-inhibit-takeover` parameter is not required because the nodes are not yet part of an HA pair.

*Clustered Data ONTAP 8.3 man page: `system node reboot` - Reboot a node*

- d. Wait while the node boots and the Node Setup wizard automatically starts on the console.
- e. Press Enter four times to accept the existing settings for the node-management LIF.
- f. Log in to the node as the `admin` user, which does not require a password.

- g. Verify that the software version on the node is correct by using the `system node image show` command.
5. Repeat the previous step to install the software on the second node.

## 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 Data ONTAP:
  - a. Enter `halt` to exit Maintenance mode.
  - b. Boot Data ONTAP by entering `boot_ontap`
  - c. Wait while the node boots and the Node 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.

## Joining the nodes to the cluster

When the newly installed controller modules are ready, you can add each one to the cluster by using the `cluster setup` command, which first sets the node in HA mode and then joins it to the cluster.

### 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.

#### Example

```

::> cluster setup

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.

```

2. When prompted to create a new cluster or join an existing one, enter **join**.

#### Example

```

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

```

3. Set storage failover to HA mode by rebooting when prompted:
  - a. When prompted to reboot to set storage failover to HA mode, press Enter.

#### Example

```

Non-HA mode, Reboot node to activate HA

Warning: Ensure that the HA partner has started disk initialization before
rebooting this node to enable HA.
Do you want to reboot now to set storage failover (SFO) to HA mode?
{yes, no} [yes]:

Rebooting now

```

- b. Wait while the node boots and the Node Setup wizard automatically starts on the console.

- c. Press Enter four times to accept the existing settings for the node-management LIF.
  - d. Log in to the node as the **admin** user, which does not require a password.
4. Complete the cluster join process:
- a. Start the Cluster Setup wizard by using the `cluster setup` command.
  - b. When prompted to create a new cluster or join an existing one, enter **join**.
  - c. When prompted with the existing cluster interface configuration, press Enter to accept it.

### Example

```
Existing cluster interface configuration found:

Port      MTU      IP           Netmask
e1a       9000     169.254.87.75 255.255.0.0

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

5. Follow the prompts to join the existing cluster.

### Example

```
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.

To begin storing and serving data on this cluster, log in to the command-line
interface (for example, ssh admin@172.17.178.28) and complete the following
additional tasks if they have not already been completed:

- Join additional nodes to the cluster by running "cluster setup" on
  those nodes.
- For HA configurations, verify that storage failover is enabled by
  running the "storage failover show" command.
  Create a Vserver by running the "vserver create" command.

In addition to using the CLI to perform cluster management tasks, you can manage
your cluster using OnCommand System Manager, which features a graphical user
interface that simplifies many cluster management tasks. This software is
available from the NetApp Support Site.

Exiting the cluster setup wizard.
```

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

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

### Example

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

7. Repeat steps 1 through 6 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.

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

### Example

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.
```

*Clustered Data ONTAP 8.3 man page: [storage failover show](#) - Display storage failover status*

## 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 AutoSupport on the new nodes](#) on page 23

After you add nodes to a cluster, you must configure AutoSupport on the nodes.

2. [Configuring the Service Processor network](#) on page 24  
After you expand a cluster, you must ensure that the Service Processor (SP) network is configured on the new nodes. If the SP uses manual network configuration, you must first manually configure the IP addresses for SP on the new nodes. If the SP uses automatic network configuration, you must identify the IP addresses that were selected.
3. [Validating the configuration of the expanded cluster](#) on page 25  
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.
4. [Generating an AutoSupport message about completing expansion](#) on page 27  
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.
5. [Updating LUN paths for the new nodes](#) on page 27  
If your cluster is configured for SAN, you must create SAN LIFs on the newly added nodes and then update paths.

## Configuring AutoSupport on the new nodes

After you add nodes to a cluster, you must configure AutoSupport on the nodes.

### Before you begin

AutoSupport must be set up on the cluster's existing nodes.

### About this task

You must perform this procedure on both the nodes.

### Steps

1. View the AutoSupport configuration using the `system node autosupport show` command with the `-node` parameter set to one of the nodes in the original cluster.

### Example

```
cluster1::> system node autosupport show -node cluster1-1
Node: cluster1-1
State: enable
SMTP Mail Hosts: smtp.example.com
...
```

*[Clustered Data ONTAP 8.3 man page: system node autosupport show - Display AutoSupport configuration](#)*

2. On one of the newly added nodes, configure AutoSupport in the same way that it is configured on the existing nodes by using the `system node autosupport modify` command.

### Example

```
cluster1::> system node autosupport modify -node cluster1-3 -state enable -mail-hosts
smtp.example.com -from alerts@node3.example.com -to support@example.com -support enable -
transport https -noteto pda@example.com -retry-interval 23m
```

*Clustered Data ONTAP 8.3 man page: [system node autosupport modify](#) - [Modify AutoSupport configuration](#)*

3. Repeat the previous step for the other newly added node.

## Configuring the Service Processor network

After you expand a cluster, you must ensure that the Service Processor (SP) network is configured on the new nodes. If the SP uses manual network configuration, you must first manually configure the IP addresses for SP on the new nodes. If the SP uses automatic network configuration, you must identify the IP addresses that were selected.

### Steps

1. If the cluster SP uses manual network configuration, configure IP addresses on both nodes for the SP network by using the `system service-processor network modify` command.

### Example

The following commands configure the SP network in cluster1-3 and cluster1-4 nodes.

```
cluster1::> system service-processor network modify -node cluster1-3 -address-family IPv4
-enable true -ip-address 192.168.123.98-netmask 255.255.255.0 -gateway 192.168.123.1
cluster1::> system service-processor network modify -node cluster1-4 -address-family IPv4
-enable true -ip-address 192.168.123.99 -netmask 255.255.255.0 -gateway 192.168.123.1
```

*Clustered Data ONTAP 8.3 man page: [system service-processor network modify](#) - [Modify the network configuration](#)*

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

Verification is required in all situations. Even if the SP network was automatically configured, you should ensure that it was configured successfully, and you have to learn which IP addresses were assigned.

### Example

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
                        ...

```

*Clustered Data ONTAP 8.3 man page: [system service-processor network show](#) - Display the network configuration*

3. If your site typically has DNS entries for the SP network, ensure that 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.

**Example**

```
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.

**Example**

```
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.
```

*Clustered Data ONTAP 8.3 man page: [cluster ring show](#) - Display cluster node member's replication rings*

- 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

1. 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.

### Example

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.
```

*[Clustered Data ONTAP 8.3 man page: system node autosupport invoke - Generate and send an AutoSupport message](#)*

## 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.

### Example

```
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.

### Example

```
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
```

*Clustered Data ONTAP 8.3 man page: [network interface create](#) - Create a logical interface*

- 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.

*Clustered Data ONTAP 8.3 man page: [network interface show](#) - Display logical interfaces*

2. Update port sets:
  - a. Determine whether port sets exist by using the `lun portset show` command.
 

*Clustered Data ONTAP 8.3 man page: [lun portset show](#) - Displays a list of portsets*
  - 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.
 

*Clustered Data ONTAP 8.3 man page: [lun portset add](#) - Add iSCSI/FCP LIFs to a portset*
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.

*[Clustered Data ONTAP 8.3 man page: lun mapping add-reporting-nodes - Add Reporting Nodes](#)*

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**

*[Clustered Data ONTAP 8.3 SAN Configuration Guide](#)*

*[Clustered Data ONTAP 8.3 SAN Administration Guide](#)*

## 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.

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

### Comprehensive guides

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

- *[Clustered Data ONTAP 8.3 SAN Administration Guide](#)*  
Describes how to configure and manage the iSCSI, FCoE, and FC protocols for clustered SAN environments, including configuration of LUNs, igroups, and targets.
- *[Clustered Data ONTAP 8.3 Logical Storage Management Guide](#)*  
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.

## 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](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**
- about this guide
    - deciding whether to use [4](#)
  - adding
    - switches before cluster expansion [9](#)
  - ASUP
    - See AutoSupport
  - audience
    - for the guide [4](#)
  - AutoSupport
    - configuring on new nodes [23](#)
    - generating message after expansion [27](#)
    - generating message before expansion [13](#)
- B**
- backup files
    - creating before expansion [13](#)
  - block access
    - updating LUN paths [27](#)
- C**
- cabling
    - new nodes [15](#)
  - cluster expansion
    - where to find additional information after [30](#)
  - cluster rings
    - verifying after expansion [25](#)
  - clusters
    - adding nodes to [20](#)
    - backing up configuration information before expansion [13](#)
    - cabling and installing nodes [15](#)
    - preparing for expansion [10](#)
    - requirements for using this guide to expand [4](#)
    - upgrading or downgrading nodes for expansion [17](#)
    - verifying cluster is ready for expansion [6](#)
    - verifying expansion limits for [6](#)
    - verifying health after expansion [25](#)
    - verifying health before expansion [11](#)
    - workflow for adding two nodes to an existing cluster [5](#)
  - comments
    - how to send feedback about documentation [33](#)
  - Config Advisor
    - running after expansion [25](#)
    - updating and verifying configuration health before expansion [11](#)
  - configurations
    - saving before expansion [13](#)
  - controller modules
    - reused, ensuring HA state [19](#)
- D**
- Data ONTAP
    - upgrading or downgrading nodes to match cluster version [17](#)
  - documentation
    - additional information [30](#)
    - how to receive automatic notification of changes to [33](#)
    - how to send feedback about [33](#)
  - downgrading
    - nodes for cluster expansion [17](#)
- E**
- expansion
    - adding licenses [10](#)
    - adding or replacing switches for cluster [9](#)
    - cabling and installing nodes [15](#)
    - cluster, requirements for using this guide [4](#)
    - completing [22](#)
    - configuring AutoSupport after [23](#)
    - creating backups before [13](#)
    - gathering networking and login information [8](#)
    - generating AutoSupport message after [27](#)
    - generating AutoSupport message before [13](#)
    - preparing cluster for [10](#)
    - verifying limits for cluster [6](#)
    - verifying system health after [25](#)
    - verifying system health before [11](#)
    - verifying the cluster is ready for expansion [6](#)
    - workflow for cluster [5](#)
  - express guides
    - additional documentation [30](#)
    - cluster expansion workflow [5](#)
    - requirements for using this guide [4](#)

**F**

- FC
  - updating paths [27](#)
- feedback
  - how to send comments about documentation [33](#)
- flowcharts
  - cluster expansion workflow [5](#)

**H**

- health
  - verifying after expansion [25](#)
  - verifying before expansion [11](#)
- high availability
  - enabling at hardware-level [19](#)

**I**

- information
  - how to send feedback about improving documentation [33](#)
- installing
  - nodes before cluster expansion [15](#)
- IP addresses
  - for SPs [24](#)
  - gathering for node management LIFs [8](#)
  - gathering for SPs [8](#)
- iSCSI
  - updating paths [27](#)

**J**

- joining
  - nodes to the cluster [20](#)

**L**

- licenses
  - adding node-locked [10](#)
  - verifying before a cluster expansion [6](#)
- LIFs
  - configuring node management [16](#)
  - creating for new nodes [27](#)
- limits
  - verifying before a cluster expansion [6](#)
- LUNs
  - updating paths [27](#)

**M**

- messages
  - generating AutoSupport [13](#)
  - generating AutoSupport message after expansion [27](#)
- MPIO
  - updating [27](#)

**N**

- netmask
  - See* network mask
- network mask
  - gathering before cluster expansion [8](#)
- networking information
  - gathering before cluster expansion [8](#)
- networks
  - for SPs [24](#)
  - SPs [8](#)
- node-locked licenses
  - verifying before a cluster expansion [6](#)
- nodes
  - adding [15](#)
  - adding licenses [10](#)
  - adding to the cluster [20](#)
  - backing up configuration information before expansion [13](#)
  - configuring AutoSupport [23](#)
  - configuring node management LIFs [16](#)
  - installing new [14](#)
  - requirements for using this guide to add to an existing cluster [4](#)
  - upgrading or downgrading to match the cluster [17](#)
  - workflow for adding two nodes to an existing cluster [5](#)

**P**

- paths
  - updating to LUNs [27](#)
- platforms
  - verifying they can be mixed [6](#)
  - verifying they support Data ONTAP 8.3 [6](#)
- ports
  - gathering before cluster expansion [8](#)
- post-requisites
  - configuring AutoSupport after expansion [23](#)
  - generating AutoSupport message after expansion [27](#)
  - verifying system health after expansion [25](#)
- powering on

## 36 | Cluster Expansion Express Guide

- nodes [16](#)
- prerequisites
  - adding licenses [10](#)
  - backing up configuration information [13](#)
  - generating AutoSupport message before expansion [13](#)
  - verifying system health before expansion [11](#)
- protocols
  - where to get information about setting up [30](#)

## R

- racking
  - nodes [15](#)
- replacing
  - switches before cluster expansion [9](#)
- repurposed
  - controller modules, ensuring HA state [19](#)

## S

- SAN
  - updating paths [27](#)
  - verifying iSCSI, FC/FCoE configuration limits before cluster expansion [6](#)
- Service Processors
  - See* SPs
- SPs
  - gathering before cluster expansion [8](#)
  - network configuration [24](#)
- subnets
  - for SPs [8](#)
- suggestions
  - how to send feedback about documentation [33](#)
- SVMs
  - where to get information about creating [30](#)
- switches

- adding or replacing before cluster expansion [9](#)
- verifying they are ready for cluster expansion [6](#)
- switchless
  - adding switches before cluster expansion [9](#)
- system health
  - verifying after expansion [25](#)
  - verifying before expansion [11](#)
- system limits
  - verifying before a cluster expansion [6](#)

## T

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

## U

- upgrading
  - nodes for cluster expansion [17](#)

## V

- verifying system health
  - after expansion [25](#)
  - before expansion [11](#)
- version
  - verifying consistency before a cluster expansion [6](#)
- volumes
  - where to get information about creating [30](#)

## W

- workflows
  - cluster expansion [5](#)