



Install the new nodes

System Manager Classic

NetApp
July 01, 2022

This PDF was generated from https://docs.netapp.com/us-en/ontap-sm-classic/expansion/task_installing_controllers.html on July 01, 2022. Always check docs.netapp.com for the latest.

Table of Contents

- Install the new nodes 1
 - Install the controllers 1
 - Configure node-management LIFs 1
 - Upgrade or downgrade the nodes 2
 - Ensure hardware-level HA is enabled 3
 - Add nodes to a cluster using System Manager 4
 - Join nodes to the cluster using the CLI 5

Install 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.

Install 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

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.

Configure 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.

Upgrade or downgrade 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, you can [Upgrade ONTAP](#) on nodes with earlier versions or you can [Revert ONTAP](#) for nodes with later versions.

Ensure 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.

Add 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

1. Add the new compatible nodes to the cluster:

If you are...	Do this...
Not logged in to System Manager	<ol style="list-style-type: none"> a. Log in to System Manager. <div style="display: flex; align-items: center; margin: 10px 0;">  <div style="border-left: 1px solid #ccc; padding-left: 10px;"> <p>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> </div> </div> <ol style="list-style-type: none"> b. Click Add Nodes to Cluster. c. Modify the name of the nodes. d. Specify the node licenses. e. Click Submit and Proceed.

Logged in to System Manager

- a. Depending on the System Manager version that you are running, perform one of the following steps:
 - ONTAP 9.4 or earlier: Click **Configuration > Cluster Expansion**.
 - Beginning with ONTAP 9.5: Click **Configuration > Cluster > Expansion** System Manager searches for newly added nodes. If any warnings are displayed, you must fix them before proceeding. If new compatible nodes are discovered, proceed to the next step.
- b. Modify the name of the nodes.
- c. Specify the node licenses.
- d. Click **Submit and Proceed**.

Join 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:
```



For instructions using the GUI-based cluster setup wizard, see [Adding nodes to the cluster using System Manager](#).

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
eth	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.
```


- Repeat steps [#STEP_3D8223C5AC7145EE8C9A9397270D0610](#) through [#STEP_F6678CB6B1A94AF08F86F83BA8BA8E35](#) 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.
- 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.
```

Copyright Information

Copyright © 2022 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, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.