



Cluster administration

System Manager Classic

NetApp
January 13, 2022

Table of Contents

- Cluster administration 1
 - Cluster administration 1
 - Volume move management 23
 - SNMP configuration 30

Cluster administration

Cluster administration

Cluster expansion administration

Cluster expansion overview

This content 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 content 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 content use a two-node cluster, it 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 *Classic* 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.



Verify 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.

Verify 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      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 ONTAP software for reference later in this workflow.

Gather 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          Avail/
              Subnet          Domain            Gateway            Total    Ranges
-----
sub1          10.53.33.1/18   Default           10.53.0.1           2/4
10.53.33.3-10.53.33.6
...
```

Add or replace 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.

Procedure

- 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](#)

Prepare 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.

Add 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.

Verify 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.



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. 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.
```

Back 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. 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.

Generate 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.

Steps

1. 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.
```

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
      Is      Is      Install
Node  Image   Default Current Version   Date
-----
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.

Steps

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="border-left: 1px solid #ccc; padding-left: 10px; margin: 10px 0;">  <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> <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.

If you are...	Do this...
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 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>b. Modify the name of the nodes.</p> <p>c. Specify the node licenses.</p> <p>d. Click Submit and Proceed.</p>

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
e1a       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 [#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.
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.

```

Complete 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.

Configure 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.
 - b. Enter the IP address, netmask, and gateway details.
3. Click **Submit and Proceed** to complete the network configuration of the nodes.
4. Verify the details of the nodes in the **Summary** page.

What to do next

- If your cluster is protected, you should create the required number of intercluster LIFs in the newly added nodes to avoid partial peering and unhealthy protection.
- If SAN data protocols are enabled in your cluster, you should create the required number of SAN Data LIFs for serving data.

Configure 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.

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

```

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.

```
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
```

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

Configure the Service Processor network

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.

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.

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
```

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.

Validate 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.

Generate 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.

Steps

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.

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.
```

Update 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.

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

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

- [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.

Volume move management

Volume move overview

Using ONTAP System Manager classic interface with ONTAP 9.7 or earlier, you can nondisruptively move a data volume from one node to another node within the same storage virtual machine (SVM) in an ONTAP 9 cluster.

Requirements for using this content

Before you use perform these procedures, ensure that the following conditions are met:

- The cluster is running ONTAP 9.
- You have cluster administrator privileges.
- You want to use best practices, not explore every available option.
- You want to use the *Classic* System Manager UI for ONTAP 9.7 and earlier releases, not the ONTAP System Manager UI for ONTAP 9.7 and later.

For some tasks, you must use the ONTAP command-line interface (CLI).

- You know which volume you want to move.

For help in deciding which volumes to move, you can use Active IQ Unified Manager (formerly OnCommand Unified Manager).

- The volume that will be moved is a data volume.
- Any new or repurposed hardware is fully installed and already has aggregates.
- If the cluster has LUNs, all nodes have two paths per LUN.

- Flow control is not enabled on cluster network ports.
- For volumes containing namespaces, the cluster is running ONTAP 9.6 or later.

Volume move is not supported for NVMe configurations running ONTAP 9.5.

Other ways to do this in ONTAP

To complete this task using...	Refer to...
The redesigned System Manager (available with ONTAP 9.7 and later)	Manage volumes
The ONTAP command-line interface	Logical storage management

Alternatives to volume move

Before moving volumes, you should evaluate whether the following approaches are better suited to your situation:

- If you want to nondisruptively upgrade a controller in place, you can consider using aggregate relocation (ARL), which does not require physical data movement.

[High Availability](#)


- If you want to move only a LUN—without its containing volume—you can use the LUN move process.

[SAN administration](#)

Volume move workflow

You can move a volume by using the ONTAP System Manager *classic* interface with ONTAP 9.7 or earlier.

Before moving a volume, you should select a method for the volume move operation and plan the timing of the operation. After the move, you might have to update the NDMP backup configuration.



Other ways to do this in ONTAP

To complete this workflow using...	Refer to...
The redesigned System Manager (available with ONTAP 9.7 and later)	Manage volumes
The ONTAP command-line interface	Logical storage management

Plan the method and timing of a volume move

You can use the ONTAP System Manager *classic* interface with ONTAP 9.7 or earlier to move a volume and to decide whether to have a manual cutover. If you need to update LUN reporting nodes, you must follow an advanced procedure in the command-line interface (CLI). Optionally, you can also plan the timing of a volume move.

About this task

The source volume of a SnapMirror or SnapVault relationship can be moved while the volume is being mirrored. SnapMirror services encounter a brief pause during the cutover phase of the volume move job.

The destination volume can also be moved. In the iterative phase, SnapMirror or SnapVault updates and volume move operations run concurrently. When evaluating whether a cutover is possible in the cutover phase, priority between the cutover and SnapMirror or SnapVault updates is determined on a first-come, first-served basis. Until the first operation finishes, other operations are blocked.

Steps

1. Decide whether you require a manual cutover.

Cutover is the moment at which the move operation finishes and ONTAP starts serving data from the volume on the new aggregate. The cutover can occur automatically or you can trigger the cutover manually.

If your company's standard practice requires you to control when changes occur in the storage system, you can manually perform the final cutover of the move operation during a maintenance window.

A cutover does not require an outage, but you can use a maintenance window to control *when* it occurs.



The volume move operation is nondisruptive, regardless of whether you choose automatic or manual cutover.

2. If the volume contains LUNs and the cluster contains four or more nodes, use the CLI to update the LUN reporting nodes if the volume moves to a different HA pair.

If the volume does not contain LUNs or if the cluster contains only two nodes, you can skip this step.

3. Plan a time using the following considerations:

- A volume move operation might take more time than expected because moves are designed to occur nondisruptively in the background in a manner that preserves client access and overall system performance.

For example, ONTAP throttles the resources that are available to the volume move operation.

- If you want the move to occur as quickly as possible, you must select a time with less cluster activity, especially the following activities:
 - I/O operations on the volume
 - Jobs using background resources, for example, when controller CPU usage is less than 50 percent
 - Jobs using the cluster interconnect
- A move cannot be started while the volume is affected by the following operations: volume offline, restrict, or destroy; SnapMirror resync, break, or restore; and Snapshot restore.

You must wait for any of these specific operations to finish before you can start the move.

- While the volume move operation occurs, a MetroCluster switchback cannot occur, although a switchover can occur.
- MetroCluster switchbacks are blocked when volume move operations are in progress for volumes belonging to the switched over site. Switchbacks are not blocked when volume move operations are in progress for volumes local to the surviving site.
- Forced MetroCluster switchovers can occur when volume move operations are in progress.

Related information

[Verifying LUN reporting nodes after moving a volume](#)

Move a volume using System Manager

Using the ONTAP System Manager *classic* interface with ONTAP 9.7 or earlier, you can move a volume by selecting a volume and the destination aggregate, starting the volume move operation, and optionally monitoring the volume move job. When using System Manager, a volume move operation finishes automatically.

Before you begin

You should have reviewed the available space on the source aggregate and destination aggregate before the volume move operation and after the volume move operation.

About this task

A volume move operation is supported only within the same cluster. In addition, you should note that the aggregate you are moving the volume to and the aggregate you are moving the volume from must be in the same storage virtual machine (SVM). A volume move does not disrupt client access.

Steps

1. Navigate to the **Volumes** window.
2. Select the volume that you want to move, and then click **Actions > Move**.
3. Select the destination aggregate, and then start the volume move operation:
 - a. Select a destination aggregate from the list of possible aggregates, which includes only the aggregates that have the required capacity.

You should review the available space, total space, RAID type, and storage type of the aggregates. For example, if the goal is to alter the performance characteristics of the volume, you can focus on aggregates with the desired storage type.

- b. Click **Move**, and then click **Move** again to confirm that you want to proceed with the volume move operation.

When the Move Volume dialog box is displayed, leave the dialog box open if you want to monitor the volume move job.

4. Monitor the volume move job:
 - a. In the **Move Volume** dialog box, click the link to the **Job ID** of the volume move job.
 - b. Locate the volume move job, and then review the information in the **Status** column.

The job can be in any one of several phases, such as transferring the initial baseline of data or starting a cutover attempt.

241	03/05/2015 07:3...	Volume Move	node1-1	running	Move "vol1" in V...	Cutover Started:(1 of 3 attempts) Transferring final da...
-----	--------------------	-------------	---------	---------	---------------------	--

- c. Click **Refresh** in the **Jobs** window to view the updated job status.


241	03/05/2015 07:3...	Volume Move	node1-1	success	Move "vol1" in V...	Complete: Successful [0]
-----	--------------------	-------------	---------	---------	---------------------	--------------------------

The job status changes to `Complete: Successful` when the volume move operation finishes.

5. If the volume move job enters the cutover deferred phase, perform a manual cutover.
 - a. From the **Volumes** window, select the volume for which you initiated the volume move job.
 - b. Initiate cutover for the volume:

If you are running...	Perform these steps...
ONTAP 9.3 or later	<ol style="list-style-type: none"> 1. Expand the volume and click the Show More Details link to view more information about the volume. 2. In the Overview tab, click Cutover.
ONTAP 9.2 or earlier	In the Volume Move Details tab, click Cutover .

- c. In the **Cutover** dialog box, click **Advanced Options** .
 - d. Specify the cutover action and the cutover duration.



- e. Click **OK**.
6. Repeat [Step 4](#).

Verify LUN reporting nodes after moving a volume

You can add LUN reporting nodes to help maintain optimized LUN paths.

If the volume that you move contains LUNs, and the destination aggregate is on another high-availability (HA) pair, ONTAP automatically adds an HA pair to the Selective LUN Map reporting-nodes list

Before you begin

Two LIFs must be configured: one LIF on the destination node and the other LIF on the HA partner of the destination node.

About this task

This procedure is required only if you move a volume from one HA pair to a different HA pair. If you move a volume to a different node of the same HA pair—for example, if you have a two-node cluster or a MetroCluster configuration—you can skip this procedure.

Steps

1. Verify that the destination node and its partner node are in the reporting-nodes list of the volume. If the nodes are not in the reporting-nodes list, add the destination node and its partner node to the reporting-nodes list of the volume:

```
lun mapping add-reporting-nodes
```

2. Rescan from the host to discover the newly added paths.
3. Add the new paths to your MPIO configuration.
4. Remove the previous LUN owner and its partner node from the reporting-nodes list:

```
lun mapping remove-reporting-nodes-remote-nodes -vserver vserver_name -path  
lun_path -igroup igroup_name
```

5. Rescan the host to verify removal of old paths.

See your host documentation for specific steps to rescan your hosts.

Update LUN reporting nodes after moving a volume

If the volume that you moved contained LUNs and the volume is now on a different HA pair, you should remove all remote nodes from the Selective LUN Map (SLM) reporting-nodes list. The LUN map then contains only the owner node and its HA partner, which ensures that only optimized LUN paths are used.

About this task

This procedure is necessary only if you moved the volume from its HA pair to a different HA pair. If the volume is on a different node of the same HA pair, you can skip this procedure.

Steps

1. Remove all of the remote nodes from the reporting-nodes list by using the `lun mapping remove-`

reporting-nodes command with the `-remote-nodes` parameter.

```
cluster1::> lun mapping remove-reporting-nodes -vserver SVM1 -volume
voll -igroup ig1 -remote-nodes true
```

2. Verify that the LUN map contains only the owner node and its partner by using the `lun mapping show` command with the `-fields reporting-nodes` parameter.

```
cluster1::> lun mapping show -vserver SVM1 -volume voll -fields
reporting-nodes
vserver path          igroup  reporting-nodes
-----
SVM1     /vol/voll  ig1     cluster1-3,cluster1-4
```

3. Remove stale device entries for the host operating system.
4. Rescan from the host to refresh the host's available paths.

See your host documentation for specific steps to rescan your hosts.

Update NDMP backup after moving a volume

If the volume that you moved was previously backed up to tape using NDMP in a specific configuration, after moving the volume, you can perform one of the following actions to ensure the volume continues to be backed up successfully: create a baseline or migrate the backup LIF to the node containing the moved volume.

About this task

- This procedure is necessary only if the backup application does not support the cluster-aware backup (CAB) extension and the backup process uses node-scoped NDMP.

If the backup application supports CAB and it is configured to use the SVM-scoped NDMP mode, you can skip this procedure.

- You must perform only one of these actions, not both.

Procedure

- From the backup application, create a new baseline.
- Identify the LIF that is configured for the backup process, and then migrate the LIF to the node where the volume now resides.

SNMP configuration

SNMP configuration overview

Using the ONTAP System Manager *classic* interface with ONTAP 9.7 or earlier, you can configure SNMP at the cluster management level; add communities, security users, and traphosts; and test the SNMP communication.

You should use these procedures if you want to configure SNMP access to a cluster in the following way:

- You are working with clusters running ONTAP 9.
- You want to use best practices, not explore every available option.



There are a few steps in these procedures for which you must use the command-line interface.


Other ways to do this in ONTAP

You can configure SNMP access to a cluster using for all ONTAP 9 versions. You should use the appropriate procedure for your version of ONTAP.

To perform these tasks with...	Refer to...
The redesigned System Manager (available with ONTAP 9.7 and later)	Manage SNMP on the cluster (cluster administrators only) > Overview
The ONTAP command-line interface (CLI)	Commands for managing SNMP

SNMP configuration workflow

Configuring SNMP involves enabling SNMP, optionally configuring an SNMPv1 or SNMPv2c community, optionally adding an SNMPv3 user, adding SNMP traphosts, and testing SNMP polling and traps.



Verify that SNMP is enabled

You can use the ONTAP System Manager *classic* interface with ONTAP 9.7 or earlier to verify whether SNMP is enabled on the cluster.


About this task

In all versions of ONTAP, SNMPv3 is enabled by default at the cluster level and SNMPv1 and SNMPv2c are disabled by default. SNMPv1 and SNMPv2c are enabled when you create an SNMP community.

SNMP is disabled by default on data LIFs. For information about enabling SNMP on data LIFs, see [Network management](#).

Steps

1. Click the groove icon.
2. In the **Setup** pane, navigate to the **SNMP** window.



You can view the current SNMP status for the cluster.

If SNMP is not enabled, click **Enable**.

Add an SNMP community

You can use the ONTAP System Manager *classic* interface with ONTAP 9.7 or earlier to add a community to the administrative storage virtual machine (SVM) for a cluster that is running SNMPv1 or SNMPv2c. System Manager uses SNMP protocols SNMPv1 and SNMPv2c, and an SNMP community to discover storage systems.

About this task

This procedure is for adding an SNMP community to the administrative SVM for the cluster. The procedure for adding an SNMP community to a data SVM is described in [Network management](#).

In new installations of ONTAP, SNMPv1 and SNMPv2c are disabled by default. SNMPv1 and SNMPv2c are enabled when you create an SNMP community.

Steps

1. In the SNMP window, click **Edit** to open the **Edit SNMP Settings** dialog box.
2. In the **General** tab, specify the contact personnel and location for the ONTAP system.
3. Click **Add**, enter a community name, and then click **OK** in the **Community Names** pane.

You can add multiple community names. A community name can be a maximum of 32 characters and must not contain the following special characters: , / : " ' |

4. When you finish adding community names, click **OK** in the **Edit SNMP Settings** dialog box.

Add an SNMPv3 security user

You can use the ONTAP System Manager *classic* interface with ONTAP 9.7 or earlier to add an SNMPv3 user at the cluster level.

The SNMPv3 user can run SNMP utilities from the traphost (SNMP manager) using the authentication and privacy settings that you specify. SNMPv3 offers advanced security by using passphrases and encryption.

About this task

When you add an SNMPv3 user at the cluster level, that user can access the cluster through all the LIFs that

have the “mgmt” firewall policy applied.

Steps

1. In the SNMP window, click **Edit** to open the **Edit SNMP Settings** dialog box.
2. In the **SNMPv3** tab, click **Add** to open the **Add SNMPv3 User** dialog box.
3. Enter the following values:

- a. Enter an SNMPv3 user name.

A security user name must not exceed 31 characters and must not contain the following special characters:

, / : " ' |

- b. For Engine ID, select the default value `Local Engine ID`.

The Engine ID is used to generate authentication and encryption keys for SNMPv3 messages.

- c. Select an authentication protocol and enter an authentication password.

A password must contain a minimum of eight characters.

- d. Optional: Select a privacy protocol and enter a password for it.

4. Click **OK** in the **Add SNMPv3 User** dialog box.

You can add multiple security user names, clicking **OK** after each addition. For example, if you use SNMP to monitor different applications that require different privileges, you might need to add an SNMPv3 user for each monitoring or management function.

5. When you finish adding user names, click **OK** in the **Edit SNMP Settings** dialog box.

Add an SNMP traphost

You can use the ONTAP System Manager *classic* interface with ONTAP 9.7 or earlier to add a traphost (SNMP manager) to receive SNMP notifications (SNMP trap protocol data units) when traps are generated in the cluster.

Before you begin

IPv6 must be enabled on the cluster if you configure SNMP traphosts that have IPv6 addresses.

About this task

SNMP and SNMP traps are enabled by default. The NetApp Technical Report TR-4220 on SNMP support contains lists of all default events that are supported by SNMP traps.

[NetApp Technical Report 4220: SNMP Support in Data ONTAP](#)

Steps

1. In the SNMP window, click **EDIT** to open the **Edit SNMP Settings** dialog box.
2. In the **Trap Hosts** tab, verify that the **Enable traps** check box is selected and click **Add**.
3. Enter the traphost IP address, and then click **OK** in the **Trap Hosts** pane.

The IP address of an SNMP trap host can be IPv4 or IPv6.

4. To add another trap host, repeat [Step 2](#) and [Step 3](#).
5. When you finish adding trap hosts, click **OK** in the **Edit SNMP Settings** dialog box.

Test SNMP traps

You can use the ONTAP System Manager *classic* interface with ONTAP 9.7 or earlier to test SNMP traps. Because communication with a trap host is not automatically validated when you add it, you should verify that the SNMP trap host can correctly receive traps.

Steps

1. Navigate to the **SNMP** screen.
2. Click **Test Trap Host** to generate a trap from the cluster in which you added a trap host.
3. From the trap host location, verify that the trap was received.

Use whatever software you ordinarily use to manage the SNMP trap host.

Test SNMP polling

After you configure SNMP, you should verify that you can poll the cluster.

About this task

To poll a cluster, you need to use a third-party command such as `snmpwalk`.

Steps

1. Send an SNMP command to poll the cluster from a different cluster.

For systems running SNMPv1, use the CLI command `snmpwalk -v version -c community_string ip_address_or_host_name system` to discover the contents of the MIB (Management Information Base).

In this example, the IP address of the cluster management LIF that you are polling is 10.11.12.123. The command displays the requested information from the MIB:

```
C:\Windows\System32>snmpwalk -v 1 -c public 10.11.12.123 system

SNMPv1-MIB::sysDescr.0 = STRING: NetApp Release 8.3.0
                        Cluster-Mode: Tue Apr 22 16:24:48 EDT 2014
SNMPv1-MIB::sysObjectID.0 = OID: SNMPv1-SMI::enterprises.789.2.5
DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks: (162644448) 18 days,
19:47:24.48
SNMPv1-MIB::sysContact.0 = STRING:
SNMPv1-MIB::sysName.0 = STRING: systemname.testlabs.com
SNMPv1-MIB::sysLocation.0 = STRING: Floor 2 Row B Cab 2
SNMPv1-MIB::sysServices.0 = INTEGER: 72
```

For systems running SNMPv2c, use the CLI command `snmpwalk -v version -c community_string ip_address_or_host_name system` to discover the contents of the MIB (Management Information Base).

In this example, the IP address of the cluster management LIF that you are polling is 10.11.12.123. The command displays the requested information from the MIB:

```
C:\Windows\System32>snmpwalk -v 2c -c public 10.11.12.123 system

SNMPv2-MIB::sysDescr.0 = STRING: NetApp Release 8.3.0
    Cluster-Mode: Tue Apr 22 16:24:48 EDT 2014
SNMPv2-MIB::sysObjectID.0 = OID: SNMPv2-SMI::enterprises.789.2.5
DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks: (162635772) 18 days,
19:45:57.72
SNMPv2-MIB::sysContact.0 = STRING:
SNMPv2-MIB::sysName.0 = STRING: systemname.testlabs.com
SNMPv2-MIB::sysLocation.0 = STRING: Floor 2 Row B Cab 2
SNMPv2-MIB::sysServices.0 = INTEGER: 72
```

For systems running SNMPv3, use the CLI command `snmpwalk -v 3 -a MD5 or SHA -l authnopriv -u username -A password ip_address_or_host_name system` to discover the contents of the MIB (Management Information Base).

In this example, the IP address of the cluster management LIF that you are polling is 10.11.12.123. The command displays the requested information from the MIB:

```
C:\Windows\System32>snmpwalk -v 3 -a MD5 -l authnopriv -u snmpv3
-a password123 10.11.12.123 system

SNMPv3-MIB::sysDescr.0 = STRING: NetApp Release 8.3.0
    Cluster-Mode: Tue Apr 22 16:24:48 EDT 2014
SNMPv3-MIB::sysObjectID.0 = OID: SNMPv3-SMI::enterprises.789.2.5
DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks: (162666569) 18 days,
19:51:05.69
SNMPv3-MIB::sysContact.0 = STRING:
SNMPv3-MIB::sysName.0 = STRING: systemname.testlabs.com
SNMPv3-MIB::sysLocation.0 = STRING: Floor 2 Row B Cab 2
SNMPv3-MIB::sysServices.0 = INTEGER: 72
```

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.