



NVIDIA SN2100

Cluster and storage switches

NetApp
April 25, 2024

This PDF was generated from <https://docs.netapp.com/us-en/ontap-systems-switches/switch-nvidia-sn2100/configure-overview-sn2100-cluster.html> on April 25, 2024. Always check docs.netapp.com for the latest.

Table of Contents

- NVIDIA SN2100 1
 - Overview 1
 - Install hardware 3
 - Configure software 13
 - Migrate switches 69
 - Replace switches 128

NVIDIA SN2100

Overview

Overview of installation and configuration for NVIDIA SN2100 switches

The NVIDIA SN2100 is a cluster switch that allows you to build ONTAP clusters with more than two nodes.

Initial configuration overview

To configure a NVIDIA SN2100 switch on systems running ONTAP, follow these steps:

1. [Install the hardware for the NVIDIA SN2100 switch.](#)

Instructions are available in the *NVIDIA Switch Installation Guide*.

2. [Configure the switch.](#)

Instructions are available in NVIDIA's documentation.

3. [Review cabling and configuration considerations.](#)

Review requirements for optical connections, the QSA adapter, and the switchport speed.

4. [Cable the NS224 shelves as switch-attached storage.](#)

Follow the cabling procedures if you have a system in which the NS224 drive shelves need to be cabled as switch-attached storage (not direct-attached storage).

5. [Install Cumulus Linux in Cumulus mode](#) or [install Cumulus Linux in ONIE mode](#).

You can install Cumulus Linux (CL) OS when the switch is running either Cumulus Linux or ONIE.

6. [Install the Reference Configuration File \(RCF\) script.](#)

There are two RCF scripts available for Clustering and Storage applications. The procedure for each is the same.

7. [Configure SNMPv3 for switch log collection.](#)

This release includes support for SNMPv3 for switch log collection and for Switch Health Monitoring (SHM).

The procedures use Network Command Line Utility (NCLU), which is a command line interface that ensures Cumulus Linux is fully accessible to all. The net command is the wrapper utility you use to execute actions from a terminal.

Additional information

Before you begin installation or maintenance, be sure to review the following:

- [Configuration requirements](#)

- [Components and part numbers](#)
- [Required documentation](#)
- [Hardware Universe](#) for all supported ONTAP versions.

Configuration requirements for NVIDIA SN2100 switches

For NVIDIA SN2100 switch installation and maintenance, be sure to review all configuration requirements.

Installation requirements

If you want to build ONTAP clusters with more than two nodes, you need two supported cluster network switches. You can use additional management switches, which are optional.

You install the NVIDIA SN2100 switch (X190006) in the NVIDIA dual/single switch cabinet with the standard brackets that are included with the switch.

For cabling guidelines, see [Review cabling and configuration considerations](#).

ONTAP and Linux support

The NVIDIA SN2100 switch is a 10/25/40/100GbE switch running Cumulus Linux. The switch supports the following:

- ONTAP 9.10.1P3.

The SN2100 switch serves Cluster and Storage applications in ONTAP 9.10.1P3 over different switch-pairs.

- Cumulus Linux (CL) OS version.

In order to download the SN2100 Cumulus software from NVIDIA, you must have login credentials to access NVIDIA's Enterprise Support Portal. See the Knowledge Base article [How to register with NVIDIA for Enterprise Support Portal Access](#). For current compatibility information, see the [NVIDIA Ethernet Switches](#) information page.

- You can install Cumulus Linux when the switch is running Cumulus Linux or ONIE.

Components and part numbers for NVIDIA SN2100 switches

For NVIDIA SN2100 switch installation and maintenance, be sure to review the list of components and part numbers for the cabinet and rail kit.

Cabinet details

You install the NVIDIA SN2100 switch (X190006) in the NVIDIA dual/single switch cabinet with the standard brackets that are included with the switch.

Rail kit details

The following table lists the part number and description for the SN2100 switches and rail kits:

Part number	Description
X190006-PE	Cluster Switch, NVIDIA SN2100, 16PT 100GbE, PTSX
X190006-PI	Cluster Switch, NVIDIA SN2100, 16PT 100GbE, PSIN
X-MTEF-KIT-D	Rail Kit, NVIDIA Dual switch side by side
X-MTEF-KIT-E	Rail Kit, NVIDIA Single switch short depth



See NVIDIA documentation for details on [installing your SN2100 switch and rail kit](#).

Documentation requirements for NVIDIA SN2100 switches

For NVIDIA SN2100 switch installation and maintenance, be sure to review all the recommended documentation.

Title	Description
NVIDIA Switch Installation Guide	Describes how to install your NVIDIA SN2100 switches.
NS224 NVMe Drive Shelf Cabling Guide	Overview and illustrations showing how to configure cabling for drive shelves.
NetApp Hardware Universe	Allows you to confirm supported hardware, such as storage switches and cables, for your platform model.

Install hardware

Install the hardware for the NVIDIA SN2100 switch

To install the SN2100 hardware, refer to NVIDIA's documentation.

Steps

1. Review the [configuration requirements](#).
2. Follow the instructions in [NVIDIA Switch Installation Guide](#).

What's next?

[Configure the switch](#).

Configure the NVIDIA SN2100 switch

To configure the SN2100 switch, refer to NVIDIA's documentation.

Steps

1. Review the [configuration requirements](#).

2. Follow the instructions in [NVIDIA System Bring-Up..](#)

What's next?

[Review cabling and configuration considerations.](#)

Review cabling and configuration considerations

Before configuring your NVIDIA SN2100 switch, review the following considerations.

NVIDIA port details

Switch ports	Ports usage
swp1s0-3	4x10GbE breakout cluster port nodes
swp2s0-3	4x25GbE breakout cluster port nodes
swp3-14	40/100GbE cluster port nodes
swp15-16	40/100GbE Inter-Switch Link (ISL) ports

See the [Hardware Universe](#) for more information on switch ports.

Link-up delays with optical connections

If you are experiencing link-up delays of more than five seconds, Cumulus Linux 5.4 and later includes support for fast link-up. You can configure the links by using the `nv set` command as follows:

```
nv set interface <interface-id> link fast-linkup on
nv config apply
reload the switchd
```

Show example

```
cumulus@cumulus-cs13:mgmt:~$ nv set interface swp5 link fast-linkup on
cumulus@cumulus-cs13:mgmt:~$ nv config apply
switchd need to reload on this config change

Are you sure? [y/N] y
applied [rev_id: 22]

Only switchd reload required
```

Support for copper connections

The following configuration changes are required to fix this issue.

Cumulus Linux 4.4.3

1. Identify the name for each interface using 40GbE/100GbE copper cables:

```
cumulus@cumulus:mgmt:~$ net show interface pluggables
```

Interface	Identifier	Vendor Name	Vendor PN	Vendor SN
Vendor Rev				
-----	-----	-----	-----	-----

swp3	0x11 (QSFP28)	Molex	112-00576	93A2229911111
B0				
swp4	0x11 (QSFP28)	Molex	112-00576	93A2229922222
B0				

2. Add the following two lines to the `/etc/cumulus/switchd.conf` file for every port (swp<n>) that is using 40GbE/100GbE copper cables:

- `interface.swp<n>.enable_media_depended_linkup_flow=TRUE`
- `interface.swp<n>.enable_short_tuning=TRUE`

For example:

```
cumulus@cumulus:mgmt:~$ sudo nano /etc/cumulus/switchd.conf
.
.
interface.swp3.enable_media_depended_linkup_flow=TRUE
interface.swp3.enable_short_tuning=TRUE
interface.swp4.enable_media_depended_linkup_flow=TRUE
interface.swp4.enable_short_tuning=TRUE
```

3. Restart the `switchd` service:

```
cumulus@cumulus:mgmt:~$ sudo systemctl restart switchd.service
```

4. Confirm that the ports are up:


```
cumulus@cumulus:mgmt:~$ net show interface all
```

State	Name	Spd	MTU	Mode	LLDP	Summary
UP	swp3	100G	9216	Trunk/L2		Master: bridge(UP)
UP	swp4	100G	9216	Trunk/L2		Master: bridge(UP)

Cumulus Linux 5.x

1. Identify the name for each interface using 40GbE/100GbE copper cables:

```
cumulus@cumulus:mgmt:~$ nv show interface pluggables
```

Interface	Identifier	Vendor Name	Vendor PN	Vendor SN
Vendor Rev				
swp3	0x11 (QSFP28)	Molex	112-00576	93A2229911111
B0				
swp4	0x11 (QSFP28)	Molex	112-00576	93A2229922222
B0				

2. Configure the links using the `nv set` command as follows:

- `nv set interface <interface-id> link fast-linkup on`
- `nv config apply`
- Reload the `switchd` service

For example:

```
cumulus@cumulus:mgmt:~$ nv set interface swp5 link fast-linkup on
cumulus@cumulus:mgmt:~$ nv config apply
switchd need to reload on this config change

Are you sure? [y/N] y
applied [rev_id: 22]

Only switchd reload required
```

3. Confirm that the ports are up:

```
cumulus@cumulus:mgmt:~$ net show interface all
```

State	Name	Spd	MTU	Mode	LLDP	Summary
UP	swp3	100G	9216	Trunk/L2		Master: bridge(UP)
UP	swp4	100G	9216	Trunk/L2		Master: bridge(UP)

See [this KB](#) for further details.

On Cumulus Linux 4.4.2, copper connections are not supported on SN2100 switches with X1151A NIC, X1146A NIC, or onboard 100GbE ports. For example:

- AFF A800 on ports e0a and e0b
- AFF A320 on ports e0g and e0h

QSA adapter

When a QSA adapter is used to connect to the 10GbE/25GbE cluster ports on a platform, the link might not come up.

To resolve this issue, do the following:

- For 10GbE, manually set the swp1s0-3 link speed to 10000 and set auto-negotiation to off.
- For 25GbE, manually set the swp2s0-3 link speed to 25000 and set auto-negotiation to off.



When using 10GbE/25GbE QSA adapters, insert them in non-breakout 40GbE/100GbE ports (swp3-swp14). Do not insert the QSA adapter in a port that is configured for breakout.

Setting interface speed on breakout ports

Depending on the transceiver in the switch port, you might need to set the speed on the switch interface to a fixed speed. If using 10GbE and 25GbE breakout ports, verify that auto-negotiation is off and set the interface speed on the switch.

Cumulus Linux 4.4.3

For example:

```
cumulus@cumulus:mgmt:~$ net add int swp1s3 link autoneg off && net com
--- /etc/network/interfaces      2019-11-17 00:17:13.470687027 +0000
+++ /run/nclu/ifupdown2/interfaces.tmp  2019-11-24 00:09:19.435226258
+0000
@@ -37,21 +37,21 @@
     alias 10G Intra-Cluster Node
     link-autoneg off
     link-speed 10000 <---- port speed set
     mstpctl-bpduguard yes
     mstpctl-portadminedge yes
     mtu 9216

auto swp1s3
iface swp1s3
    alias 10G Intra-Cluster Node
-   link-autoneg off
+   link-autoneg on
    link-speed 10000 <---- port speed set
    mstpctl-bpduguard yes
    mstpctl-portadminedge yes
    mtu 9216

auto swp2s0
iface swp2s0
    alias 25G Intra-Cluster Node
    link-autoneg off
    link-speed 25000 <---- port speed set
```

Check the interface and port status to verify that the settings are applied:

```
cumulus@cumulus:mgmt:~$ net show interface
```

State	Name	Spd	MTU	Mode	LLDP	Summary
-----	-----	-----	-----	-----	-----	

.						
.						
UP	swp1s0	10G	9216	Trunk/L2	cs07 (e4c)	Master:
br_default(UP)						
UP	swp1s1	10G	9216	Trunk/L2	cs07 (e4d)	Master:
br_default(UP)						
UP	swp1s2	10G	9216	Trunk/L2	cs08 (e4c)	Master:
br_default(UP)						
UP	swp1s3	10G	9216	Trunk/L2	cs08 (e4d)	Master:
br_default(UP)						
.						
.						
UP	swp3	40G	9216	Trunk/L2	cs03 (e4e)	Master:
br_default(UP)						
UP	swp4	40G	9216	Trunk/L2	cs04 (e4e)	Master:
br_default(UP)						
DN	swp5	N/A	9216	Trunk/L2		Master:
br_default(UP)						
DN	swp6	N/A	9216	Trunk/L2		Master:
br_default(UP)						
DN	swp7	N/A	9216	Trunk/L2		Master:
br_default(UP)						
.						
.						
UP	swp15	100G	9216	BondMember	cs01 (swp15)	Master:
cluster_isl(UP)						
UP	swp16	100G	9216	BondMember	cs01 (swp16)	Master:
cluster_isl(UP)						
.						
.						

Cumulus Linux 5.x

For example:

```
cumulus@cumulus:mgmt:~$ nv set interface swp1s3 link auto-negotiate off
cumulus@cumulus:mgmt:~$ nv set interface swp1s3 link speed 10G
cumulus@cumulus:mgmt:~$ nv show interface swp1s3
```

```
link
```

auto-negotiate	off	off
duplex	full	full
speed	10G	10G
fec	auto	auto
mtu	9216	9216
[breakout]		
state	up	up

Check the interface and port status to verify that the settings are applied:

```
cumulus@cumulus:mgmt:~$ nv show interface
```

State	Name	Spd	MTU	Mode	LLDP	Summary
-----	-----	-----	-----	-----	-----	

.						
.						
UP	swp1s0	10G	9216	Trunk/L2	cs07 (e4c)	Master:
br_default(UP)						
UP	swp1s1	10G	9216	Trunk/L2	cs07 (e4d)	Master:
br_default(UP)						
UP	swp1s2	10G	9216	Trunk/L2	cs08 (e4c)	Master:
br_default(UP)						
UP	swp1s3	10G	9216	Trunk/L2	cs08 (e4d)	Master:
br_default(UP)						
.						
.						
UP	swp3	40G	9216	Trunk/L2	cs03 (e4e)	Master:
br_default(UP)						
UP	swp4	40G	9216	Trunk/L2	cs04 (e4e)	Master:
br_default(UP)						
DN	swp5	N/A	9216	Trunk/L2		Master:
br_default(UP)						
DN	swp6	N/A	9216	Trunk/L2		Master:
br_default(UP)						
DN	swp7	N/A	9216	Trunk/L2		Master:
br_default(UP)						
.						
.						
UP	swp15	100G	9216	BondMember	cs01 (swp15)	Master:
cluster_isl(UP)						
UP	swp16	100G	9216	BondMember	cs01 (swp16)	Master:
cluster_isl(UP)						
.						
.						

What's next?

[Cable NS224 shelves as switch-attached storage.](#)

Cable the NS224 shelves as switch-attached storage

If you have a system in which the NS224 drive shelves need to be cabled as switch-attached storage (not direct-attached storage), use the information provided here.

- Cable NS224 drive shelves through storage switches:

[Cabling switch-attached NS224 drive shelves](#)

- Confirm supported hardware, such as storage switches and cables, for your platform model:

[NetApp Hardware Universe](#)

What's next?

[Install Cumulus Linux in Cumulus mode](#) or [Install Cumulus Linux in ONIE mode](#).

Configure software

Software install workflow for NVIDIA SN2100 switches

To install and configure software for a NVIDIA SN2100 switch, follow these steps:

1. [Install Cumulus Linux in Cumulus mode](#) or [install Cumulus Linux in ONIE mode](#).

You can install Cumulus Linux (CL) OS when the switch is running either Cumulus Linux or ONIE.

2. [Install the Reference Configuration File \(RCF\) script](#).

There are two RCF scripts available for Clustering and Storage applications. The procedure for each is the same.

3. [Configure SNMPv3 for switch log collection](#).

This release includes support for SNMPv3 for switch log collection and for Switch Health Monitoring (SHM).

The procedures use Network Command Line Utility (NCLU), which is a command line interface that ensures Cumulus Linux is fully accessible to all. The net command is the wrapper utility you use to execute actions from a terminal.

Install Cumulus Linux in Cumulus mode

Follow this procedure to install Cumulus Linux (CL) OS when the switch is running in Cumulus mode.



Cumulus Linux (CL) OS can be installed either when the switch is running Cumulus Linux or ONIE (see [Install in ONIE mode](#)).

What you'll need

- Intermediate-level Linux knowledge.
- Familiarity with basic text editing, UNIX file permissions, and process monitoring. A variety of text editors are pre-installed, including `vi` and `nano`.
- Access to a Linux or UNIX shell. If you are running Windows, use a Linux environment as your command line tool for interacting with Cumulus Linux.
- The baud rate requirement is set to 115200 on the serial console switch for NVIDIA SN2100 switch console access, as follows:

- 115200 baud
- 8 data bits
- 1 stop bit
- parity: none
- flow control: none

About this task

Be aware of the following:



Each time Cumulus Linux is installed, the entire file system structure is erased and rebuilt.



The default password for the cumulus user account is **cumulus**. The first time you log into Cumulus Linux, you must change this default password. Be sure to update any automation scripts before installing a new image. Cumulus Linux provides command line options to change the default password automatically during the installation process.

Example 1. Steps

Cumulus Linux 4.4.3

1. Log in to the switch.

First time log in to the switch requires username/password of **cumulus/cumulus** with **sudo** privileges.

```
cumulus login: cumulus
Password: cumulus
You are required to change your password immediately (administrator
enforced)
Changing password for cumulus.
Current password: cumulus
New password: <new_password>
Retype new password: <new_password>
```

2. Check the Cumulus Linux version: `net show system`

```
cumulus@cumulus:mgmt:~$ net show system
Hostname..... cumulus
Build..... Cumulus Linux 4.4.3
Uptime..... 0:08:20.860000
Model..... Mlnx X86
CPU..... x86_64 Intel Atom C2558 2.40GHz
Memory..... 8GB
Disk..... 14.7GB
ASIC..... Mellanox Spectrum MT52132
Ports..... 16 x 100G-QSFP28
Part Number..... MSN2100-CB2FC
Serial Number.... MT2105T05177
Platform Name.... x86_64-mlnx_x86-r0
Product Name..... MSN2100
ONIE Version..... 2019.11-5.2.0020-115200
Base MAC Address. 04:3F:72:43:92:80
Manufacturer..... Mellanox
```

3. Configure the hostname, IP address, subnet mask, and default gateway. The new hostname only becomes effective after restarting the console/SSH session.



A Cumulus Linux switch provides at least one dedicated Ethernet management port called `eth0`. This interface is specifically for out-of-band management use. By default, the management interface uses DHCPv4 for addressing.



Do not use an underscore (_), apostrophe ('), or non-ASCII characters in the hostname.

```
cumulus@cumulus:mgmt:~$ net add hostname sw1
cumulus@cumulus:mgmt:~$ net add interface eth0 ip address
10.233.204.71
cumulus@cumulus:mgmt:~$ net add interface eth0 ip gateway
10.233.204.1
cumulus@cumulus:mgmt:~$ net pending
cumulus@cumulus:mgmt:~$ net commit
```

This command modifies both the /etc/hostname and /etc/hosts files.

4. Confirm that the hostname, IP address, subnet mask, and default gateway have been updated.

```
cumulus@sw1:mgmt:~$ hostname sw1
cumulus@sw1:mgmt:~$ ifconfig eth0
eth0: flags=4163<UP,BROADCAST,RUNNING,MULTICAST> mtu 1500
inet 10.233.204.71 netmask 255.255.254.0 broadcast 10.233.205.255
inet6 fe80::bace:f6ff:fe19:1df6 prefixlen 64 scopeid 0x20<link>
ether b8:ce:f6:19:1d:f6 txqueuelen 1000 (Ethernet)
RX packets 75364 bytes 23013528 (21.9 MiB)
RX errors 0 dropped 7 overruns 0 frame 0
TX packets 4053 bytes 827280 (807.8 KiB)
TX errors 0 dropped 0 overruns 0 carrier 0 collisions 0 device
memory 0xdfc00000-dfc1ffff

cumulus@sw1::mgmt:~$ ip route show vrf mgmt
default via 10.233.204.1 dev eth0
unreachable default metric 4278198272
10.233.204.0/23 dev eth0 proto kernel scope link src 10.233.204.71
127.0.0.0/8 dev mgmt proto kernel scope link src 127.0.0.1
```

5. Configure the time zone using NTP interactive mode.

- a. On a terminal, run the following command:

```
cumulus@sw1:~$ sudo dpkg-reconfigure tzdata
```

- b. Follow the on-screen menu options to select the geographic area and region.
 - c. To set the time zone for all services and daemons, reboot the switch.
 - d. Verify that the date and time on the switch are correct and update if necessary.
6. Install Cumulus Linux 4.4.3:

```
cumulus@sw1:mgmt:~$ sudo onie-install -a -i http://<web-server>/<path>/cumulus-linux-4.4.3-mlx-amd64.bin
```

The installer starts the download. Type **y** when prompted.

7. Reboot the NVIDIA SN2100 switch:

```
cumulus@sw1:mgmt:~$ sudo reboot
```

8. The installation starts automatically, and the following GRUB screen choices appear. Do **not** make any selections.

- Cumulus-Linux GNU/Linux
- ONIE: Install OS
- CUMULUS-INSTALL
- Cumulus-Linux GNU/Linux

9. Repeat steps 1 to 4 to log in.

10. Verify that the Cumulus Linux version is 4.4.3: `net show version`

```
cumulus@sw1:mgmt:~$ net show version  
NCLU_VERSION=1.0-cl4.4.3u0  
DISTRIB_ID="Cumulus Linux"  
DISTRIB_RELEASE=4.4.3  
DISTRIB_DESCRIPTION="Cumulus Linux 4.4.3"
```

11. Create a new user and add this user to the `sudo` group. This user only becomes effective after the console/SSH session is restarted.

```
sudo adduser --ingroup netedit admin
```

```

cumulus@sw1:mgmt:~$ sudo adduser --ingroup netedit admin
[sudo] password for cumulus:
Adding user 'admin' ...
Adding new user 'admin' (1001) with group `netedit' ...
Creating home directory '/home/admin' ...
Copying files from '/etc/skel' ...
New password:
Retype new password:
passwd: password updated successfully
Changing the user information for admin
Enter the new value, or press ENTER for the default
Full Name []:
Room Number []:
Work Phone []:
Home Phone []:
Other []:
Is the information correct? [Y/n] y

cumulus@sw1:mgmt:~$ sudo adduser admin sudo
[sudo] password for cumulus:
Adding user `admin' to group `sudo' ...
Adding user admin to group sudo
Done.
cumulus@sw1:mgmt:~$ exit
logout
Connection to 10.233.204.71 closed.

[admin@cycrh6svl01 ~]$ ssh admin@10.233.204.71
admin@10.233.204.71's password:
Linux sw1 4.19.0-cl-1-amd64 #1 SMP Cumulus 4.19.206-1+cl4.4.1u1
(2021-09-09) x86_64
Welcome to NVIDIA Cumulus (R) Linux (R)

For support and online technical documentation, visit
http://www.cumulusnetworks.com/support

The registered trademark Linux (R) is used pursuant to a sublicense
from LMI, the exclusive licensee of Linus Torvalds, owner of the
mark on a world-wide basis.
admin@sw1:mgmt:~$

```

Cumulus Linux 5.x

1. Log in to the switch.

First time log in to the switch requires username/password of **cumulus/cumulus** with **sudo**

privileges.

```
cumulus login: cumulus
Password: cumulus
You are required to change your password immediately (administrator
enforced)
Changing password for cumulus.
Current password: cumulus
New password: <new_password>
Retype new password: <new_password>
```

2. Check the Cumulus Linux version: `nv show system`

```
cumulus@cumulus:mgmt:~$ nv show system
```

operational	applied	description
hostname	cumulus	cumulus
build	Cumulus Linux 5.3.0	system build version
uptime	6 days, 8:37:36	system uptime
timezone	Etc/UTC	system time zone

3. Configure the hostname, IP address, subnet mask, and default gateway. The new hostname only becomes effective after restarting the console/SSH session.



A Cumulus Linux switch provides at least one dedicated Ethernet management port called `eth0`. This interface is specifically for out-of-band management use. By default, the management interface uses DHCPv4 for addressing.



Do not use an underscore (`_`), apostrophe (`'`), or non-ASCII characters in the hostname.

```
cumulus@cumulus:mgmt:~$ nv set system hostname sw1
cumulus@cumulus:mgmt:~$ nv set interface eth0 ip address
10.233.204.71/24
cumulus@cumulus:mgmt:~$ nv set interface eth0 ip gateway
10.233.204.1
cumulus@cumulus:mgmt:~$ nv config apply
cumulus@cumulus:mgmt:~$ nv config save
```

This command modifies both the `/etc/hostname` and `/etc/hosts` files.

4. Confirm that the hostname, IP address, subnet mask, and default gateway have been updated.

```

cumulus@sw1:mgmt:~$ hostname sw1
cumulus@sw1:mgmt:~$ ifconfig eth0
eth0: flags=4163<UP,BROADCAST,RUNNING,MULTICAST> mtu 1500
inet 10.233.204.71 netmask 255.255.254.0 broadcast 10.233.205.255
inet6 fe80::bace:f6ff:fe19:1df6 prefixlen 64 scopeid 0x20<link>
ether b8:ce:f6:19:1d:f6 txqueuelen 1000 (Ethernet)
RX packets 75364 bytes 23013528 (21.9 MiB)
RX errors 0 dropped 7 overruns 0 frame 0
TX packets 4053 bytes 827280 (807.8 KiB)
TX errors 0 dropped 0 overruns 0 carrier 0 collisions 0 device
memory 0xdfc00000-dfc1ffff

cumulus@sw1::mgmt:~$ ip route show vrf mgmt
default via 10.233.204.1 dev eth0
unreachable default metric 4278198272
10.233.204.0/23 dev eth0 proto kernel scope link src 10.233.204.71
127.0.0.0/8 dev mgmt proto kernel scope link src 127.0.0.1

```

5. Configure the time zone using NTP interactive mode.

- a. On a terminal, run the following command:

```

cumulus@sw1:~$ sudo dpkg-reconfigure tzdata

```

- b. Follow the on-screen menu options to select the geographic area and region.
c. To set the time zone for all services and daemons, reboot the switch.
d. Verify that the date and time on the switch are correct and update if necessary.

6. Install Cumulus Linux 5.4:

```

cumulus@sw1:mgmt:~$ sudo onie-install -a -i http://<web-
server>/<path>/cumulus-linux-5.4-mlx-amd64.bin

```

The installer starts the download. Type **y** when prompted.

7. Reboot the NVIDIA SN2100 switch:

```

cumulus@sw1:mgmt:~$ sudo reboot

```

8. The installation starts automatically, and the following GRUB screen choices appear. Do **not** make any selections.

- Cumulus-Linux GNU/Linux
- ONIE: Install OS

- CUMULUS-INSTALL
- Cumulus-Linux GNU/Linux

9. Repeat steps 1 to 4 to log in.

10. Verify that the Cumulus Linux version is 5.4: `nv show system`

```
cumulus@cumulus:mgmt:~$ nv show system
```

operational	applied	description
hostname	cumulus	cumulus
build	Cumulus Linux 5.4.0	system build version
uptime	6 days, 13:37:36	system uptime
timezone	Etc/UTC	system time zone

11. Verify that the nodes each have a connection to each switch:

```
cumulus@sw1:mgmt:~$ net show lldp
```

LocalPort	Speed	Mode	RemoteHost
RemotePort			
eth0	100M	Mgmt	mgmt-sw1
Eth110/1/29			
swp2s1	25G	Trunk/L2	node1
e0a			
swp15	100G	BondMember	sw2
swp15			
swp16	100G	BondMember	sw2
swp16			

12. Create a new user and add this user to the `sudo` group. This user only becomes effective after the console/SSH session is restarted.

```
sudo adduser --ingroup netedit admin
```

```

cumulus@sw1:mgmt:~$ sudo adduser --ingroup netedit admin
[sudo] password for cumulus:
Adding user 'admin' ...
Adding new user 'admin' (1001) with group `netedit' ...
Creating home directory '/home/admin' ...
Copying files from '/etc/skel' ...
New password:
Retype new password:
passwd: password updated successfully
Changing the user information for admin
Enter the new value, or press ENTER for the default
Full Name []:
Room Number []:
Work Phone []:
Home Phone []:
Other []:
Is the information correct? [Y/n] y

cumulus@sw1:mgmt:~$ sudo adduser admin sudo
[sudo] password for cumulus:
Adding user `admin' to group `sudo' ...
Adding user admin to group sudo
Done.
cumulus@sw1:mgmt:~$ exit
logout
Connection to 10.233.204.71 closed.

[admin@cycrh6svl01 ~]$ ssh admin@10.233.204.71
admin@10.233.204.71's password:
Linux sw1 4.19.0-cl-1-amd64 #1 SMP Cumulus 4.19.206-1+cl4.4.1u1
(2021-09-09) x86_64
Welcome to NVIDIA Cumulus (R) Linux (R)

For support and online technical documentation, visit
http://www.cumulusnetworks.com/support

The registered trademark Linux (R) is used pursuant to a sublicense
from LMI, the exclusive licensee of Linus Torvalds, owner of the
mark on a world-wide basis.
admin@sw1:mgmt:~$

```

13. Add additional user groups for the admin user to access `nv` commands:


```
cumulus@sw1:mgmt:~$ sudo adduser admin nvshow
[sudo] password for cumulus:
Adding user 'admin' to group 'nvshow' ...
Adding user admin to group nvshow
Done.
```

See [NVIDIA User Accounts](#) for more information.

What's next?

[Install the Reference Configuration File \(RCF\) script.](#)

Install Cumulus Linux in ONIE mode

Follow this procedure to install Cumulus Linux (CL) OS when the switch is running in ONIE mode.



Cumulus Linux (CL) OS can be installed either when the switch is running ONIE or Cumulus Linux (see [Install in Cumulus mode](#)).

About this task

You can install Cumulus Linux using Open Network Install Environment (ONIE) that allows for automatic discovery of a network installer image. This facilitates the system model of securing switches with an operating system choice, such as Cumulus Linux. The easiest way to install Cumulus Linux with ONIE is with local HTTP discovery.



If your host is IPv6-enabled, make sure it is running a web server. If your host is IPv4-enabled, make sure it is running DHCP in addition to a web server.

This procedure demonstrates how to upgrade Cumulus Linux after the admin has booted in ONIE.

Example 2. Steps

Cumulus Linux 4.4.3

1. Download the Cumulus Linux installation file to the root directory of the web server. Rename this file to: `onie-installer`.
2. Connect your host to the management Ethernet port of the switch using an Ethernet cable.
3. Power on the switch.

The switch downloads the ONIE image installer and boots. After the installation completes, the Cumulus Linux login prompt appears in the terminal window.



Each time Cumulus Linux is installed, the entire file system structure is erased and rebuilt.

4. Reboot the SN2100 switch:

```
cumulus@cumulus:mgmt:~$ sudo reboot
```

5. Press the **Esc** key at the GNU GRUB screen to interrupt the normal boot process, select **ONIE**, and press **Enter**.
6. On the next screen, select **ONIE: Install OS**.
7. The ONIE installer discovery process runs searching for the automatic installation. Press **Enter** to temporarily stop the process.
8. When the discovery process has stopped:

```
ONIE:/ # onie-stop
discover: installer mode detected.
Stopping: discover...start-stop-daemon: warning: killing process
427:
No such process done.
```

9. If the DHCP service is running on your network, verify that the IP address, subnet mask, and the default gateway are correctly assigned:

```
ifconfig eth0
```

```

ONIE:/ # ifconfig eth0
eth0    Link encap:Ethernet  HWaddr B8:CE:F6:19:1D:F6
        inet addr:10.233.204.71  Bcast:10.233.205.255
Mask:255.255.254.0
        inet6 addr: fe80::bace:f6ff:fe19:1df6/64 Scope:Link
        UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
        RX packets:21344 errors:0 dropped:2135 overruns:0 frame:0
        TX packets:3500 errors:0 dropped:0 overruns:0 carrier:0
        collisions:0 txqueuelen:1000
        RX bytes:6119398 (5.8 MiB)  TX bytes:472975 (461.8 KiB)
        Memory:dfc00000-dfc1ffff

```

```

ONIE:/ # route
Kernel IP routing table

```

Destination	Gateway	Genmask	Flags	Metric	Ref
Use Iface					
default	10.233.204.1	0.0.0.0	UG	0	0
0 eth0					
10.233.204.0	*	255.255.254.0	U	0	0
0 eth0					

10. If the IP addressing scheme is manually defined, do the following:

```

ONIE:/ # ifconfig eth0 10.233.204.71 netmask 255.255.254.0
ONIE:/ # route add default gw 10.233.204.1

```

11. Repeat step 9 to verify that the static information is correctly entered.

12. Install Cumulus Linux:

```

# onie-nos-install http://<web-server>/<path>/cumulus-linux-4.4.3-
mlx-amd64.bin

```

```

ONIE:/ # route

Kernel IP routing table

ONIE:/ # onie-nos-install http://<web-server>/<path>/cumulus-
linux-4.4.3-mlx-amd64.bin

Stopping: discover... done.
Info: Attempting
http://10.60.132.97/x/eng/testbedN,svl/nic/files/cumulus-linux-
4.4.3-mlx-amd64.bin ...
Connecting to 10.60.132.97 (10.60.132.97:80)
installer          100% |*|    552M  0:00:00 ETA
...
...

```

13. After the installation has completed, log in to the switch.

```

cumulus login: cumulus
Password: cumulus
You are required to change your password immediately (administrator
enforced)
Changing password for cumulus.
Current password: cumulus
New password: <new_password>
Retype new password: <new_password>

```

14. Verify the Cumulus Linux version: `net show version`

```

cumulus@cumulus:mgmt:~$ net show version
NCLU_VERSION=1.0-cl4.4.3u4
DISTRIB_ID="Cumulus Linux"
DISTRIB_RELEASE=4.4.3
DISTRIB_DESCRIPTION="Cumulus Linux 4.4.3"

```

Cumulus Linux 5.x

1. Download the Cumulus Linux installation file to the root directory of the web server. Rename this file to: `onie-installer`.
2. Connect your host to the management Ethernet port of the switch using an Ethernet cable.
3. Power on the switch.

The switch downloads the ONIE image installer and boots. After the installation completes, the Cumulus Linux login prompt appears in the terminal window.



Each time Cumulus Linux is installed, the entire file system structure is erased and rebuilt.

4. Reboot the SN2100 switch:

```
cumulus@cumulus:mgmt:~$ sudo reboot
.
.
GNU GRUB version 2.06-3
+-----+
-----+
| Cumulus-Linux GNU/Linux
|
| Advanced options for Cumulus-Linux GNU/Linux
|
| ONIE
|
|
|
|
|
|
|
|
|
|
|
|
|
|
|
|
|
|
|
|
+-----+
-----+
```

5. Press the Esc key at the GNU GRUB screen to interrupt the normal boot process, select ONIE, and press Enter.

```

.
.
Loading ONIE ...

GNU GRUB version 2.02
+-----+
-----+
| ONIE: Install OS
|
| ONIE: Rescue
|
| ONIE: Uninstall OS
|
| ONIE: Update ONIE
|
| ONIE: Embed ONIE
|
|
|
|
|
|
|
|
|
|
+-----+
-----+

```

Select ONIE: **Install OS**.

6. The ONIE installer discovery process runs searching for the automatic installation. Press **Enter** to temporarily stop the process.
7. When the discovery process has stopped:

```

ONIE:/ # onie-stop
discover: installer mode detected.
Stopping: discover...start-stop-daemon: warning: killing process
427:
No such process done.

```

8. Configure the IP address, subnet mask, and the default gateway:

```
ifconfig eth0
```

```

ONIE:/ # ifconfig eth0
eth0    Link encap:Ethernet  HWaddr B8:CE:F6:19:1D:F6
        inet addr:10.233.204.71  Bcast:10.233.205.255
Mask:255.255.254.0
        inet6 addr: fe80::bace:f6ff:fe19:1df6/64 Scope:Link
        UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
        RX packets:21344 errors:0 dropped:2135 overruns:0 frame:0
        TX packets:3500 errors:0 dropped:0 overruns:0 carrier:0
        collisions:0 txqueuelen:1000
        RX bytes:6119398 (5.8 MiB)  TX bytes:472975 (461.8 KiB)
        Memory:dfc00000-dfc1ffff

ONIE:/ #
ONIE:/ # ifconfig eth0 10.228.140.27 netmask 255.255.248.0
ONIE:/ # ifconfig eth0
eth0    Link encap:Ethernet HWaddr B8:CE:F6:5E:05:E6
        inet addr:10.228.140.27 Bcast:10.228.143.255
Mask:255.255.248.0
        inet6 addr: fd20:8b1e:b255:822b:bace:f6ff:fe5e:5e6/64
Scope:Global
        inet6 addr: fe80::bace:f6ff:fe5e:5e6/64 Scope:Link
        UP BROADCAST RUNNING MULTICAST MTU:1500 Metric:1
        RX packets:18813 errors:0 dropped:1418 overruns:0 frame:0
        TX packets:491 errors:0 dropped:0 overruns:0 carrier:0
        collisions:0 txqueuelen:1000
        RX bytes:1339596 (1.2 MiB) TX bytes:49379 (48.2 KiB)
        Memory:dfc00000-dfc1ffff

ONIE:/ # route add default gw 10.228.136.1
ONIE:/ # route
Kernel IP routing table
Destination      Gateway          Genmask          Flags Metric Ref
Use Iface

default          10.228.136.1    0.0.0.0          UG    0      0
0 eth0
10.228.136.1     *               255.255.248.0    U    0      0
0 eth0

```

9. Install Cumulus Linux 5.4:

```
# onie-nos-install http://<web-server>/<path>/cumulus-linux-5.4-mlx-amd64.bin
```

```

ONIE:/ # route

Kernel IP routing table

ONIE:/ # onie-nos-install http://<web-server>/<path>/cumulus-
linux-5.4-mlx-amd64.bin

Stopping: discover... done.
Info: Attempting
http://10.60.132.97/x/eng/testbedN,svl/nic/files/cumulus-linux-5.4-
mlx-amd64.bin ...
Connecting to 10.60.132.97 (10.60.132.97:80)
installer          100% |*|    552M  0:00:00 ETA
...
...

```

10. After the installation has completed, log in to the switch.

```

cumulus login: cumulus
Password: cumulus
You are required to change your password immediately (administrator
enforced)
Changing password for cumulus.
Current password: cumulus
New password: <new_password>
Retype new password: <new_password>

```

11. Verify the Cumulus Linux version: `nv show system`

```

cumulus@cumulus:mgmt:~$ nv show system
operational      applied          description
-----
hostname         cumulus         cumulus
build            Cumulus Linux 5.4.0  system build version
uptime           6 days, 13:37:36  system uptime
timezone         Etc/UTC         system time zone

```

12. Create a new user and add this user to the `sudo` group. This user only becomes effective after the console/SSH session is restarted.

```

sudo adduser --ingroup netedit admin

```



```

cumulus@sw1:mgmt:~$ sudo adduser --ingroup netedit admin
[sudo] password for cumulus:
Adding user 'admin' ...
Adding new user 'admin' (1001) with group `netedit' ...
Creating home directory '/home/admin' ...
Copying files from '/etc/skel' ...
New password:
Retype new password:
passwd: password updated successfully
Changing the user information for admin
Enter the new value, or press ENTER for the default
Full Name []:
Room Number []:
Work Phone []:
Home Phone []:
Other []:
Is the information correct? [Y/n] y

cumulus@sw1:mgmt:~$ sudo adduser admin sudo
[sudo] password for cumulus:
Adding user `admin' to group `sudo' ...
Adding user admin to group sudo
Done.
cumulus@sw1:mgmt:~$ exit
logout
Connection to 10.233.204.71 closed.

[admin@cycrh6svl01 ~]$ ssh admin@10.233.204.71
admin@10.233.204.71's password:
Linux sw1 4.19.0-cl-1-amd64 #1 SMP Cumulus 4.19.206-1+cl4.4.1u1
(2021-09-09) x86_64
Welcome to NVIDIA Cumulus (R) Linux (R)

For support and online technical documentation, visit
http://www.cumulusnetworks.com/support

The registered trademark Linux (R) is used pursuant to a sublicense
from LMI, the exclusive licensee of Linus Torvalds, owner of the
mark on a world-wide basis.
admin@sw1:mgmt:~$

```

13. Add additional user groups for the admin user to access `nv` commands:

```
cumulus@cumulus:mgmt:~$ sudo adduser admin nvshow
[sudo] password for cumulus:
Adding user `admin' to group `nvshow' ...
Adding user admin to group nvshow
Done.
```

See [NVIDIA User Accounts](#) for more information.

What's next?

[Install the Reference Configuration File \(RCF\) script.](#)

Install the Reference Configuration File (RCF) script

Follow this procedure to install the RCF script.

What you'll need

Before installing the RCF script, make sure that the following are available on the switch:

- Cumulus Linux is installed. See the [Hardware Universe](#) for supported versions.
- IP address, subnet mask, and default gateway defined via DHCP or manually configured.



You must specify a user in the RCF (in addition to the admin user) to be used specifically for log collection.

Current RCF script versions

There are two RCF scripts available for Cluster and Storage applications. Download RCFs from [here](#). The procedure for each is the same.

- Cluster: **MSN2100-RCF-v1.x-Cluster-HA-Breakout-LLDP**
- Storage: **MSN2100-RCF-v1.x-Storage**

About the examples

The following example procedure shows how to download and apply the RCF script for Cluster switches.

Example command output uses switch management IP address 10.233.204.71, netmask 255.255.254.0 and default gateway 10.233.204.1.

Example 3. Steps

Cumulus Linux 4.4.3

1. Display the available interfaces on the SN2100 switch:

```
admin@sw1:mgmt:~$ net show interface all
```

State	Name	Spd	MTU	Mode	LLDP	Summary
-----	-----	---	-----	-----	-----	-----

...						
...						
ADMDN	swp1	N/A	9216	NotConfigured		
ADMDN	swp2	N/A	9216	NotConfigured		
ADMDN	swp3	N/A	9216	NotConfigured		
ADMDN	swp4	N/A	9216	NotConfigured		
ADMDN	swp5	N/A	9216	NotConfigured		
ADMDN	swp6	N/A	9216	NotConfigured		
ADMDN	swp7	N/A	9216	NotConfigured		
ADMDN	swp8	N/A	9216	NotConfigured		
ADMDN	swp9	N/A	9216	NotConfigured		
ADMDN	swp10	N/A	9216	NotConfigured		
ADMDN	swp11	N/A	9216	NotConfigured		
ADMDN	swp12	N/A	9216	NotConfigured		
ADMDN	swp13	N/A	9216	NotConfigured		
ADMDN	swp14	N/A	9216	NotConfigured		
ADMDN	swp15	N/A	9216	NotConfigured		
ADMDN	swp16	N/A	9216	NotConfigured		

2. Copy the RCF python script to the switch.

```
admin@sw1:mgmt:~$ pwd
/home/cumulus
cumulus@cumulus:mgmt: /tmp$ scp <user>@<host>:/<path>/MSN2100-RCF-
v1.x-Cluster-HA-Breakout-LLDP ./
ssologin@10.233.204.71's password:
MSN2100-RCF-v1.x-Cluster-HA-Breakout-LLDP          100% 8607
111.2KB/s                                00:00
```



While `scp` is used in the example, you can use your preferred method of file transfer.

3. Apply the RCF python script **MSN2100-RCF-v1.x-Cluster-HA-Breakout-LLDP**.

```
cumulus@cumulus:mgmt:/tmp$ sudo python3 MSN2100-RCF-v1.x-Cluster-HA-
Breakout-LLDP
[sudo] password for cumulus:
...
Step 1: Creating the banner file
Step 2: Registering banner message
Step 3: Updating the MOTD file
Step 4: Ensuring passwordless use of cl-support command by admin
Step 5: Disabling apt-get
Step 6: Creating the interfaces
Step 7: Adding the interface config
Step 8: Disabling cdp
Step 9: Adding the lldp config
Step 10: Adding the RoCE base config
Step 11: Modifying RoCE Config
Step 12: Configure SNMP
Step 13: Reboot the switch
```

The RCF script completes the steps listed in the example above.



In step 3 **Updating the MOTD file** above, the command `cat /etc/motd` is run. This allows you to verify the RCF filename, RCF version, ports to use, and other important information in the RCF banner.



For any RCF python script issues that cannot be corrected, contact [NetApp Support](#) for assistance.

4. Verify the configuration after the reboot:

```
admin@sw1:mgmt:~$ net show interface all
```

State	Name	Spd	MTU	Mode	LLDP	Summary
-----	-----	-----	-----	-----	-----	-----
...						
...						
DN	swp1s0	N/A	9216	Trunk/L2		Master:
bridge (UP)						
DN	swp1s1	N/A	9216	Trunk/L2		Master:
bridge (UP)						
DN	swp1s2	N/A	9216	Trunk/L2		Master:
bridge (UP)						
DN	swp1s3	N/A	9216	Trunk/L2		Master:
bridge (UP)						
DN	swp2s0	N/A	9216	Trunk/L2		Master:
bridge (UP)						

DN	swp2s1	N/A	9216	Trunk/L2	Master:
bridge (UP)					
DN	swp2s2	N/A	9216	Trunk/L2	Master:
bridge (UP)					
DN	swp2s3	N/A	9216	Trunk/L2	Master:
bridge (UP)					
UP	swp3	100G	9216	Trunk/L2	Master:
bridge (UP)					
UP	swp4	100G	9216	Trunk/L2	Master:
bridge (UP)					
DN	swp5	N/A	9216	Trunk/L2	Master:
bridge (UP)					
DN	swp6	N/A	9216	Trunk/L2	Master:
bridge (UP)					
DN	swp7	N/A	9216	Trunk/L2	Master:
bridge (UP)					
DN	swp8	N/A	9216	Trunk/L2	Master:
bridge (UP)					
DN	swp9	N/A	9216	Trunk/L2	Master:
bridge (UP)					
DN	swp10	N/A	9216	Trunk/L2	Master:
bridge (UP)					
DN	swp11	N/A	9216	Trunk/L2	Master:
bridge (UP)					
DN	swp12	N/A	9216	Trunk/L2	Master:
bridge (UP)					
DN	swp13	N/A	9216	Trunk/L2	Master:
bridge (UP)					
DN	swp14	N/A	9216	Trunk/L2	Master:
bridge (UP)					
UP	swp15	N/A	9216	BondMember	Master:
bond_15_16 (UP)					
UP	swp16	N/A	9216	BondMember	Master:
bond_15_16 (UP)					
...					
...					

admin@sw1:mgmt:~\$ **net show roce config**

RoCE mode..... lossless

Congestion Control:

Enabled SPs.... 0 2 5

Mode..... ECN

Min Threshold.. 150 KB

Max Threshold.. 1500 KB

PFC:

Status..... enabled

```
Enabled SPs.... 2 5
```

```
Interfaces..... swp10-16,swp1s0-3,swp2s0-3,swp3-9
```

DSCP	802.1p	switch-priority
-----	-----	-----
0 1 2 3 4 5 6 7	0	0
8 9 10 11 12 13 14 15	1	1
16 17 18 19 20 21 22 23	2	2
24 25 26 27 28 29 30 31	3	3
32 33 34 35 36 37 38 39	4	4
40 41 42 43 44 45 46 47	5	5
48 49 50 51 52 53 54 55	6	6
56 57 58 59 60 61 62 63	7	7

switch-priority	TC	ETS
-----	--	-----
0 1 3 4 6 7	0	DWRR 28%
2	2	DWRR 28%
5	5	DWRR 43%

5. Verify information for the transceiver in the interface:

```
admin@sw1:mgmt:~$ net show interface pluggables
```

Interface	Identifier	Vendor Name	Vendor PN	Vendor SN
Vendor Rev				
-----	-----	-----	-----	-----
swp3	0x11 (QSFP28)	Amphenol	112-00574	
APF20379253516	B0			
swp4	0x11 (QSFP28)	AVAGO	332-00440	AF1815GU05Z
A0				
swp15	0x11 (QSFP28)	Amphenol	112-00573	
APF21109348001	B0			
swp16	0x11 (QSFP28)	Amphenol	112-00573	
APF21109347895	B0			

6. Verify that the nodes each have a connection to each switch:

```
admin@sw1:mgmt:~$ net show lldp
```

LocalPort	Speed	Mode	RemoteHost	RemotePort
-----	-----	-----	-----	-----
swp3	100G	Trunk/L2	sw1	e3a
swp4	100G	Trunk/L2	sw2	e3b
swp15	100G	BondMember	sw13	swp15
swp16	100G	BondMember	sw14	swp16

7. Verify the health of cluster ports on the cluster.

a. Verify that e0d ports are up and healthy across all nodes in the cluster:

```
cluster1::*> network port show -role cluster
```

Node: node1

Ignore

						Speed (Mbps)
Health	Health					
Port	IPspace	Broadcast	Domain	Link	MTU	Admin/Oper
Status	Status					
-----	-----	-----	-----	-----	-----	-----
e3a	Cluster	Cluster		up	9000	auto/10000
healthy	false					
e3b	Cluster	Cluster		up	9000	auto/10000
healthy	false					

Node: node2

Ignore

						Speed (Mbps)
Health	Health					
Port	IPspace	Broadcast	Domain	Link	MTU	Admin/Oper
Status	Status					
-----	-----	-----	-----	-----	-----	-----
e3a	Cluster	Cluster		up	9000	auto/10000
healthy	false					
e3b	Cluster	Cluster		up	9000	auto/10000
healthy	false					

b. Verify the switch health from the cluster (this might not show switch sw2, since LIFs are not homed on e0d).

```

cluster1::*> network device-discovery show -protocol lldp
Node/          Local   Discovered
Protocol       Port    Device (LLDP: ChassisID)  Interface Platform
-----
node1/lldp
              e3a     sw1  (b8:ce:f6:19:1a:7e)    swp3      -
              e3b     sw2  (b8:ce:f6:19:1b:96)    swp3      -

node2/lldp
              e3a     sw1  (b8:ce:f6:19:1a:7e)    swp4      -
              e3b     sw2  (b8:ce:f6:19:1b:96)    swp4      -

cluster1::*> system switch ethernet show -is-monitoring-enabled
-operational true
Switch                                     Type                               Address
Model
-----
-----
sw1                                         cluster-network                    10.233.205.90
MSN2100-CB2RC
    Serial Number: MNXXXXXXGD
    Is Monitored: true
    Reason: None
    Software Version: Cumulus Linux version 4.4.3 running on
Mellanox
                                Technologies Ltd. MSN2100
    Version Source: LLDP

sw2                                         cluster-network                    10.233.205.91
MSN2100-CB2RC
    Serial Number: MNCXXXXXXGS
    Is Monitored: true
    Reason: None
    Software Version: Cumulus Linux version 4.4.3 running on
Mellanox
                                Technologies Ltd. MSN2100
    Version Source: LLDP

```

Cumulus Linux 5.x

1. Display the available interfaces on the SN2100 switch:


```

admin@sw1:mgmt:~$ nv show interface
Interface      MTU    Speed State Remote Host      Remote Port-
Type           Summary
-----
+ cluster_isl  9216   200G   up
bond
+ eth0          1500   100M   up    mgmt-sw1          Eth105/1/14
eth            IP Address: 10.231.80 206/22
  eth0
IP Address: fd20:8b1e:f6ff:fe31:4a0e/64
+ lo            65536           up
loopback      IP Address: 127.0.0.1/8
  lo
IP Address: ::1/128
+ swp1s0        9216   10G    up    cluster01         e0b
swp
.
.
.
+ swp15         9216   100G   up    sw2                swp15
swp
+ swp16         9216   100G   up    sw2                swp16
swp

```

2. Copy the RCF python script to the switch.

```

admin@sw1:mgmt:~$ pwd
/home/cumulus
cumulus@cumulus:mgmt: /tmp$ scp <user>@<host>:<path>/MSN2100-RCF-
v1.x-Cluster-HA-Breakout-LLDP ./
ssologin@10.233.204.71's password:
MSN2100-RCF-v1.x-Cluster-HA-Breakout-LLDP          100% 8607
111.2KB/s                                00:00

```



While `scp` is used in the example, you can use your preferred method of file transfer.

3. Apply the RCF python script **MSN2100-RCF-v1.x-Cluster-HA-Breakout-LLDP**.

```
cumulus@cumulus:mgmt:/tmp$ sudo python3 MSN2100-RCF-v1.x-Cluster-HA-Breakout-LLDP
[sudo] password for cumulus:
.
.
Step 1: Creating the banner file
Step 2: Registering banner message
Step 3: Updating the MOTD file
Step 4: Ensuring passwordless use of cl-support command by admin
Step 5: Disabling apt-get
Step 6: Creating the interfaces
Step 7: Adding the interface config
Step 8: Disabling cdp
Step 9: Adding the lldp config
Step 10: Adding the RoCE base config
Step 11: Modifying RoCE Config
Step 12: Configure SNMP
Step 13: Reboot the switch
```

The RCF script completes the steps listed in the example above.



In step 3 **Updating the MOTD file** above, the command `cat /etc/issue` is run. This allows you to verify the RCF filename, RCF version, ports to use, and other important information in the RCF banner.

For example:

```

admin@sw1:mgmt:~$ cat /etc/issue
*****
*****
*
* NetApp Reference Configuration File (RCF)
* Switch      : Mellanox MSN2100
* Filename     : MSN2100-RCF-1.x-Cluster-HA-Breakout-LLDP
* Release Date : 13-02-2023
* Version      : 1.x-Cluster-HA-Breakout-LLDP
*
* Port Usage:
* Port 1       : 4x10G Breakout mode for Cluster+HA Ports, swp1s0-3
* Port 2       : 4x25G Breakout mode for Cluster+HA Ports, swp2s0-3
* Ports 3-14   : 40/100G for Cluster+HA Ports, swp3-14
* Ports 15-16  : 100G Cluster ISL Ports, swp15-16
*
* NOTE:
*   RCF manually sets swp1s0-3 link speed to 10000 and
*   auto-negotiation to off for Intel 10G
*   RCF manually sets swp2s0-3 link speed to 25000 and
*   auto-negotiation to off for Chelsio 25G
*
*
* IMPORTANT: Perform the following steps to ensure proper RCF
installation:
* - Copy the RCF file to /tmp
* - Ensure the file has execute permission
* - From /tmp run the file as sudo python3 <filename>
*
*****
*****

```



For any RCF python script issues that cannot be corrected, contact [NetApp Support](#) for assistance.

4. Verify the configuration after the reboot:

```

admin@sw1:mgmt:~$ nv show interface
Interface  MTU    Speed State Remote Host Remote Port Type Summary
-----
+ cluster_isl 9216 200G up bond
+ eth0 1500 100M up RTP-LF01-410G38.rtp.eng.netapp.com Eth105/1/14
eth IP Address: 10.231.80.206/22

```

```

eth0 IP Address: fd20:8b1e:b255:85a0:bace:f6ff:fe31:4a0e/64
+ lo 65536 up loopback IP Address: 127.0.0.1/8
lo IP Address: ::1/128
+ swp1s0 9216 10G up cumulus1 e0b swp
.
.
.
+ swp15 9216 100G up cumulus swp15 swp

admin@sw1:mgmt:~$ nv show interface
Interface      MTU    Speed State Remote Host      Remote Port-
Type          Summary
-----
+ cluster_isl 9216  200G  up
bond
+ eth0          1500  100M  up    mgmt-sw1          Eth105/1/14
eth            IP Address: 10.231.80 206/22
  eth0
IP Address: fd20:8b1e:f6ff:fe31:4a0e/64
+ lo            65536      up
loopback IP Address: 127.0.0.1/8
  lo
IP Address: ::1/128
+ swp1s0        9216 10G    up cluster01          e0b
swp
.
.
.
+ swp15         9216 100G    up sw2                swp15
swp
+ swp16         9216 100G    up sw2                swp16
swp

admin@sw1:mgmt:~$ nv show qos roce
                        operational  applied  description
-----
enable                on                Turn feature 'on' or
'off'. This feature is disabled by default.
mode                  lossless    lossless  Roce Mode
congestion-control
  congestion-mode      ECN,RED      Congestion config mode
  enabled-tc           0,2,5        Congestion config enabled
Traffic Class
  max-threshold        195.31 KB    Congestion config max-

```

```

threshold
  min-threshold      39.06 KB                      Congestion config min-
threshold
  probability        100
lldp-app-tlv
  priority            3                          switch-priority of roce
  protocol-id        4791                       L4 port number
  selector            UDP                       L4 protocol
pfc
  pfc-priority        2, 5                      switch-prio on which PFC
is enabled
  rx-enabled          enabled                   PFC Rx Enabled status
  tx-enabled          enabled                   PFC Tx Enabled status
trust
  trust-mode          pcps,dscp                 Trust Setting on the port
for packet classification

```

RoCE PCP/DSCP->SP mapping configurations

```

=====
      pcps  dscps
      ---  -
0  0  0,1,2,3,4,5,6,7  0
1  1  8,9,10,11,12,13,14,15  1
2  2  16,17,18,19,20,21,22,23  2
3  3  24,25,26,27,28,29,30,31  3
4  4  32,33,34,35,36,37,38,39  4
5  5  40,41,42,43,44,45,46,47  5
6  6  48,49,50,51,52,53,54,55  6
7  7  56,57,58,59,60,61,62,63  7

```

RoCE SP->TC mapping and ETS configurations

```

=====
      switch-prio  traffic-class  scheduler-weight
      -
0  0  0  DWRR-28%
1  1  0  DWRR-28%
2  2  2  DWRR-28%
3  3  0  DWRR-28%
4  4  0  DWRR-28%
5  5  5  DWRR-43%
6  6  0  DWRR-28%
7  7  0  DWRR-28%

```

RoCE pool config

```

=====
      name  mode  size  switch-priorities

```

```

traffic-class
--  -----  -----  ----  -----
-----
0   lossy-default-ingress   Dynamic   50%    0,1,3,4,6,7   -
1   roce-reserved-ingress   Dynamic   50%    2,5            -
2   lossy-default-egress    Dynamic   50%    -              0
3   roce-reserved-egress    Dynamic   inf     -              2,5

```

Exception List

```
=====
```

```
description
```

```
--
```

```
-----
```

```
---...
```

- 1 RoCE PFC Priority Mismatch.Expected pfc-priority: 3.
- 2 Congestion Config TC Mismatch.Expected enabled-tc: 0,3.
- 3 Congestion Config mode Mismatch.Expected congestion-mode: ECN.
- 4 Congestion Config min-threshold Mismatch.Expected min-threshold: 150000.
- 5 Congestion Config max-threshold Mismatch.Expected max-threshold: 1500000.
- 6 Scheduler config mismatch for traffic-class mapped to switch-prio0.
Expected scheduler-weight: DWRR-50%.
- 7 Scheduler config mismatch for traffic-class mapped to switch-prio1.
Expected scheduler-weight: DWRR-50%.
- 8 Scheduler config mismatch for traffic-class mapped to switch-prio2.
Expected scheduler-weight: DWRR-50%.
- 9 Scheduler config mismatch for traffic-class mapped to switch-prio3.
Expected scheduler-weight: DWRR-50%.
- 10 Scheduler config mismatch for traffic-class mapped to switch-prio4.
Expected scheduler-weight: DWRR-50%.
- 11 Scheduler config mismatch for traffic-class mapped to switch-prio5.
Expected scheduler-weight: DWRR-50%.
- 12 Scheduler config mismatch for traffic-class mapped to switch-prio6.
Expected scheduler-weight: strict-priority.
- 13 Scheduler config mismatch for traffic-class mapped to switch-prio7.

```
Expected scheduler-weight: DWRR-50%.
14 Invalid reserved config for ePort.TC[2].Expected 0 Got 1024
15 Invalid reserved config for ePort.TC[5].Expected 0 Got 1024
16 Invalid traffic-class mapping for switch-priority 2.Expected
0 Got 2
17 Invalid traffic-class mapping for switch-priority 3.Expected
3 Got 0
18 Invalid traffic-class mapping for switch-priority 5.Expected
0 Got 5
19 Invalid traffic-class mapping for switch-priority 6.Expected
6 Got 0
Incomplete Command: set interface swp3-16 link fast-linkupp3-16 link
fast-linkup
Incomplete Command: set interface swp3-16 link fast-linkupp3-16 link
fast-linkup
Incomplete Command: set interface swp3-16 link fast-linkupp3-16 link
fast-linkup
```



The exceptions listed do not affect performance and can be safely ignored.

5. Verify information for the transceiver in the interface:

```
admin@sw1:mgmt:~$ nv show interface --view=pluggables
```

Interface	Identifier	Vendor Name	Vendor PN	Vendor
SN	Vendor Rev			
swp1s0	0x00	None		
swp1s1	0x00	None		
swp1s2	0x00	None		
swp1s3	0x00	None		
swp2s0	0x11	(QSFP28)	CISCO-LEONI	L45593-D278-D20
LCC2321GTTJ	00			
swp2s1	0x11	(QSFP28)	CISCO-LEONI	L45593-D278-D20
LCC2321GTTJ	00			
swp2s2	0x11	(QSFP28)	CISCO-LEONI	L45593-D278-D20
LCC2321GTTJ	00			
swp2s3	0x11	(QSFP28)	CISCO-LEONI	L45593-D278-D20
LCC2321GTTJ	00			
swp3	0x00	None		
swp4	0x00	None		
swp5	0x00	None		
swp6	0x00	None		
.				
.				
.				
swp15	0x11	(QSFP28)	Amphenol	112-00595
APF20279210117	B0			
swp16	0x11	(QSFP28)	Amphenol	112-00595
APF20279210166	B0			

6. Verify that the nodes each have a connection to each switch:

```
admin@sw1:mgmt:~$ nv show interface --view=lldp
```

LocalPort	Speed	Mode	RemoteHost	RemotePort
eth0	100M	Mgmt	mgmt-sw1	Eth110/1/29
swp2s1	25G	Trunk/L2	node1	e0a
swp15	100G	BondMember	sw2	swp15
swp16	100G	BondMember	sw2	swp16

7. Verify the health of cluster ports on the cluster.

a. Verify that e0d ports are up and healthy across all nodes in the cluster:


```
cluster1::*> network port show -role cluster
```

Node: node1

Ignore

						Speed (Mbps)
Health	Health					
Port	IPspace	Broadcast	Domain	Link	MTU	Admin/Oper
Status	Status					
-----	-----	-----	-----	-----	-----	-----
e3a	Cluster	Cluster		up	9000	auto/10000
healthy	false					
e3b	Cluster	Cluster		up	9000	auto/10000
healthy	false					

Node: node2

Ignore

						Speed (Mbps)
Health	Health					
Port	IPspace	Broadcast	Domain	Link	MTU	Admin/Oper
Status	Status					
-----	-----	-----	-----	-----	-----	-----
e3a	Cluster	Cluster		up	9000	auto/10000
healthy	false					
e3b	Cluster	Cluster		up	9000	auto/10000
healthy	false					

- b. Verify the switch health from the cluster (this might not show switch sw2, since LIFs are not homed on e0d).

```

cluster1::*> network device-discovery show -protocol lldp
Node/      Local  Discovered
Protocol   Port   Device (LLDP: ChassisID)  Interface Platform
-----
node1/lldp
          e3a    sw1  (b8:ce:f6:19:1a:7e)    swp3      -
          e3b    sw2  (b8:ce:f6:19:1b:96)    swp3      -

node2/lldp
          e3a    sw1  (b8:ce:f6:19:1a:7e)    swp4      -
          e3b    sw2  (b8:ce:f6:19:1b:96)    swp4      -

cluster1::*> system switch ethernet show -is-monitoring-enabled
-operational true
Switch                                     Type                Address
Model
-----
-----
sw1                                     cluster-network      10.233.205.90
MSN2100-CB2RC
    Serial Number: MNXXXXXXGD
    Is Monitored: true
    Reason: None
    Software Version: Cumulus Linux version 5.4.0 running on
Mellanox
                                Technologies Ltd. MSN2100
    Version Source: LLDP

sw2                                     cluster-network      10.233.205.91
MSN2100-CB2RC
    Serial Number: MNCXXXXXXGS
    Is Monitored: true
    Reason: None
    Software Version: Cumulus Linux version 5.4.0 running on
Mellanox
                                Technologies Ltd. MSN2100
    Version Source: LLDP

```

What's next?

[Configure switch log collection.](#)

Ethernet Switch Health Monitoring log collection

The Ethernet switch health monitor (CSHM) is responsible for ensuring the operational health of Cluster and Storage network switches and collecting switch logs for debugging purposes. This procedure guides you through the process of setting up and starting the collection of detailed **Support** logs from the switch and starts an hourly collection of **Periodic** data that is collected by AutoSupport.

Before you begin

- The user for log collection must be specified when the Reference Configuration File (RCF) is applied. By default, this user is set to 'admin'. If you wish to use a different user, you must specify this in the `*# SHM User*`s section of the RCF.
- The user must have access to the **nv show** commands. This can be added by running `sudo adduser USER nv show` and replacing `USER` with the user for log collection.
- Switch health monitoring must be enabled for the switch. Verify this by ensuring the `Is Monitored:` field is set to **true** in the output of the `system switch ethernet show` command.

Steps

1. To set up log collection, run the following command for each switch. You are prompted to enter the switch name, username, and password for log collection.

```
system switch ethernet log setup-password
```

Show example

```
cluster1::*> system switch ethernet log setup-password
Enter the switch name: <return>
The switch name entered is not recognized.
Choose from the following list:
cs1
cs2

cluster1::*> system switch ethernet log setup-password

Enter the switch name: cs1
Would you like to specify a user other than admin for log
collection? {y|n}: n

Enter the password: <enter switch password>
Enter the password again: <enter switch password>

cluster1::*> system switch ethernet log setup-password

Enter the switch name: cs2
Would you like to specify a user other than admin for log
collection? {y|n}: n

Enter the password: <enter switch password>
Enter the password again: <enter switch password>
```

2. To start log collection, run the following command, replacing DEVICE with the switch used in the previous command. This starts both types of log collection: the detailed Support logs and an hourly collection of Periodic data.

```
system switch ethernet log modify -device <switch-name> -log-request true
```

Show example

```
cluster1::*> system switch ethernet log modify -device cs1 -log  
-request true
```

```
Do you want to modify the cluster switch log collection  
configuration? {y|n}: [n] y
```

```
Enabling cluster switch log collection.
```

```
cluster1::*> system switch ethernet log modify -device cs2 -log  
-request true
```

```
Do you want to modify the cluster switch log collection  
configuration? {y|n}: [n] y
```

```
Enabling cluster switch log collection.
```

Wait for 10 minutes and then check that the log collection completes:

```
system switch ethernet log show
```



If any of these commands return an error or if the log collection does not complete, contact NetApp support.

Troubleshooting

If you encounter any of the following error statuses reported by the log collection feature (visible in the output of `system switch ethernet log show`), try the corresponding debug steps:

Log collection error status	Resolution
RSA keys not present	Regenerate ONTAP SSH keys. Contact NetApp support.
switch password error	Verify credentials, test SSH connectivity, and regenerate ONTAP SSH keys. Review switch documentation or contact NetApp support for instructions.
ECDSA keys not present for FIPS	If FIPS mode is enabled, ECDSA keys need to be generated on the switch before retrying.
pre-existing log found	Remove the previous log collection directory and '.tar' file located at <code>/tmp/shm_log</code> on the switch.

switch dump log error	Ensure the switch user has log collection permissions. Refer to the prerequisites above.
------------------------------	--

Configure SNMPv3

Follow this procedure to configure SNMPv3, which supports Ethernet switch health monitoring (CSHM).

About this task

The following commands configure an SNMPv3 username on NVIDIA SN2100 switches:

- For **no authentication**: `net add snmp-server username SNMPv3_USER auth-none`
- For **MD5/SHA authentication**: `net add snmp-server username SNMPv3_USER [auth-md5|auth-sha] AUTH-PASSWORD`
- For **MD5/SHA authentication with AES/DES encryption**: `net add snmp-server username SNMPv3_USER [auth-md5|auth-sha] AUTH-PASSWORD [encrypt-aes|encrypt-des] PRIV-PASSWORD`

The following command configures an SNMPv3 username on the ONTAP side: `cluster1::*> security login create -user-or-group-name SNMPv3_USER -application snmp -authentication -method usm -remote-switch-ipaddress ADDRESS`

The following command establishes the SNMPv3 username with CSHM: `cluster1::*> system switch ethernet modify -device DEVICE -snmp-version SNMPv3 -community-or-username SNMPv3_USER`

Steps

1. Set up the SNMPv3 user on the switch to use authentication and encryption:

```
net show snmp status
```

Show example

```
cumulus@sw1:~$ net show snmp status
Simple Network Management Protocol (SNMP) Daemon.
-----
Current Status                active (running)
Reload Status                 enabled
Listening IP Addresses        all vrf mgmt
Main snmpd PID                4318
Version 1 and 2c Community String Configured
Version 3 Usernames           Not Configured
-----

cumulus@sw1:~$
cumulus@sw1:~$ net add snmp-server username SNMPv3User auth-md5
<password> encrypt-aes <password>
cumulus@sw1:~$ net commit
--- /etc/snmp/snmpd.conf      2020-08-02 21:09:34.686949282 +0000
+++ /run/nclu/snmp/snmpd.conf 2020-08-11 00:13:51.826126655 +0000
@@ -1,26 +1,28 @@
# Auto-generated config file: do not edit. #
agentaddress udp:@mgmt:161
agentxperms 777 777 snmp snmp
agentxsocket /var/agentx/master
createuser _snmptrapusernameX
+createuser SNMPv3User MD5 <password> AES <password>
ifmib_max_num_ifaces 500
iquerysecname _snmptrapusernameX
master agentx
monitor -r 60 -o laNames -o laErrorMessage "laTable" laErrorFlag != 0
pass -p 10 1.3.6.1.2.1.1.1 /usr/share/snmp/sysDescr_pass.py
pass_persist 1.2.840.10006.300.43
/usr/share/snmp/ieee8023_lag_pp.py
pass_persist 1.3.6.1.2.1.17 /usr/share/snmp/bridge_pp.py
pass_persist 1.3.6.1.2.1.31.1.1.1.18
/usr/share/snmp/snmpifAlias_pp.py
pass_persist 1.3.6.1.2.1.47 /usr/share/snmp/entity_pp.py
pass_persist 1.3.6.1.2.1.99 /usr/share/snmp/entity_sensor_pp.py
pass_persist 1.3.6.1.4.1.40310.1 /usr/share/snmp/resq_pp.py
pass_persist 1.3.6.1.4.1.40310.2
/usr/share/snmp/cl_drop_cntrs_pp.py
pass_persist 1.3.6.1.4.1.40310.3 /usr/share/snmp/cl_poe_pp.py
pass_persist 1.3.6.1.4.1.40310.4 /usr/share/snmp/bgpun_pp.py
pass_persist 1.3.6.1.4.1.40310.5 /usr/share/snmp/cumulus-status.py
pass_persist 1.3.6.1.4.1.40310.6 /usr/share/snmp/cumulus-sensor.py
pass_persist 1.3.6.1.4.1.40310.7 /usr/share/snmp/vrf_bgpun_pp.py
+rocommunity cshml! default
```

```

rouser _snmptrapusernameX
+rouser SNMPv3User priv
sysobjectid 1.3.6.1.4.1.40310
syssservices 72
-rocommunity cshml! default

```

net add/del commands since the last "net commit"

=====

User	Timestamp	Command
-----	-----	-----
-----	-----	-----
SNMPv3User	2020-08-11 00:13:51.826987	net add snmp-server username
SNMPv3User	auth-md5 <password>	encrypt-aes <password>

```

cumulus@sw1:~$
cumulus@sw1:~$ net show snmp status
Simple Network Management Protocol (SNMP) Daemon.
-----
Current Status          active (running)
Reload Status           enabled
Listening IP Addresses  all vrf mgmt
Main snmpd PID          24253
Version 1 and 2c Community String Configured
Version 3 Usernames     Configured    <---- Configured
here
-----
cumulus@sw1:~$

```

2. Set up the SNMPv3 user on the ONTAP side:

```

security login create -user-or-group-name SNMPv3User -application snmp
-authentication-method usm -remote-switch-ipaddress 10.231.80.212

```


Show example

```
cluster1::*> security login create -user-or-group-name SNMPv3User  
-application snmp -authentication-method usm -remote-switch  
-ipaddress 10.231.80.212
```

Enter the authoritative entity's EngineID [remote EngineID]:

Which authentication protocol do you want to choose (none, md5, sha, sha2-256)

[none]: **md5**

Enter the authentication protocol password (minimum 8 characters long):

Enter the authentication protocol password again:

Which privacy protocol do you want to choose (none, des, aes128)

[none]: **aes128**

Enter privacy protocol password (minimum 8 characters long):

Enter privacy protocol password again:

3. Configure CSHM to monitor with the new SNMPv3 user:

```
system switch ethernet show-all -device "sw1 (b8:59:9f:09:7c:22)" -instance
```

Show example

```
cluster1::*> system switch ethernet show-all -device "sw1  
(b8:59:9f:09:7c:22)" -instance  
  
Device Name: sw1  
(b8:59:9f:09:7c:22)  
IP Address: 10.231.80.212  
SNMP Version: SNMPv2c  
Is Discovered: true  
DEPRECATED-Community String or SNMPv3 Username: -  
Community String or SNMPv3 Username: cshml!  
Model Number: MSN2100-CB2FC  
Switch Network: cluster-network  
Software Version: Cumulus Linux  
version 4.4.3 running on Mellanox Technologies Ltd. MSN2100  
Reason For Not Monitoring: None  
Source Of Switch Version: LLDP  
Is Monitored?: true  
Serial Number of the Device: MT2110X06399 <----  
serial number to check  
RCF Version: MSN2100-RCF-v1.9X6-  
Cluster-LLDP Aug-18-2022  
  
cluster1::*>  
cluster1::*> system switch ethernet modify -device "sw1  
(b8:59:9f:09:7c:22)" -snmp-version SNMPv3 -community-or-username  
SNMPv3User
```

4. Verify that the serial number to be queried with the newly created SNMPv3 user is the same as detailed in the previous step once the CSHM polling period has completed.

```
system switch ethernet polling-interval show
```

Show example

```
cluster1::*> system switch ethernet polling-interval show
Polling Interval (in minutes): 5

cluster1::*> system switch ethernet show-all -device "sw1
(b8:59:9f:09:7c:22)" -instance

Device Name: sw1
(b8:59:9f:09:7c:22)
IP Address: 10.231.80.212
SNMP Version: SNMPv3
Is Discovered: true
DEPRECATED-Community String or SNMPv3 Username: -
Community String or SNMPv3 Username: SNMPv3User
Model Number: MSN2100-CB2FC
Switch Network: cluster-network
Software Version: Cumulus Linux
version 4.4.3 running on Mellanox Technologies Ltd. MSN2100
Reason For Not Monitoring: None
Source Of Switch Version: LLDP
Is Monitored?: true
Serial Number of the Device: MT2110X06399 <----
serial number to check
RCF Version: MSN2100-RCF-v1.9X6-
Cluster-LLDP Aug-18-2022
```

Upgrade Cumulus Linux versions

Complete the following procedure to upgrade your Cumulus Linux version as required.

What you'll need

- Intermediate-level Linux knowledge.
- Familiarity with basic text editing, UNIX file permissions, and process monitoring. A variety of text editors are pre-installed, including `vi` and `nano`.
- Access to a Linux or UNIX shell. If you are running Windows, use a Linux environment as your command line tool for interacting with Cumulus Linux.
- The baud rate requirement is set to 115200 on the serial console switch for NVIDIA SN2100 switch console access, as follows:
 - 115200 baud
 - 8 data bits
 - 1 stop bit
 - parity: none

- flow control: none

About this task

Be aware of the following:



Each time Cumulus Linux is upgraded, the entire file system structure is erased and rebuilt. Your existing configuration will be erased. You must save and record your switch configuration before updating Cumulus Linux.



The default password for the cumulus user account is **cumulus**. The first time you log into Cumulus Linux, you must change this default password. You must update any automation scripts before installing a new image. Cumulus Linux provides command line options to change the default password automatically during the installation process.

Example 4. Steps

From Cumulus Linux 4.4.x to Cumulus Linux 5.x

1. Check the current Cumulus Linux version and connected ports:

```
admin@sw1:mgmt:~$ net show system
Hostname..... cumulus
Build..... Cumulus Linux 4.4.3
Uptime..... 0:08:20.860000
Model..... Mlnx X86
CPU..... x86_64 Intel Atom C2558 2.40GHz
Memory..... 8GB
Disk..... 14.7GB
ASIC..... Mellanox Spectrum MT52132
Ports..... 16 x 100G-QSFP28
Part Number..... MSN2100-CB2FC
Serial Number.... MT2105T05177
Platform Name.... x86_64-mlnx_x86-r0
Product Name..... MSN2100
ONIE Version..... 2019.11-5.2.0020-115200
Base MAC Address. 04:3F:72:43:92:80
Manufacturer..... Mellanox

admin@sw1:mgmt:~$ net show interface

State  Name      Spd   MTU    Mode      LLDP
Summary
-----
.
.
UP      swp1      100G  9216   Trunk/L2   node1 (e5b)
Master: bridge(UP)
UP      swp2      100G  9216   Trunk/L2   node2 (e5b)
Master: bridge(UP)
UP      swp3      100G  9216   Trunk/L2   SHFFG1826000112 (e0b)
Master: bridge(UP)
UP      swp4      100G  9216   Trunk/L2   SHFFG1826000112 (e0b)
Master: bridge(UP)
UP      swp5      100G  9216   Trunk/L2   SHFFG1826000102 (e0b)
Master: bridge(UP)
UP      swp6      100G  9216   Trunk/L2   SHFFG1826000102 (e0b)
Master: bridge(UP)
.
.
```

2. Download the Cumulux Linux 5.x image:

```
admin@sw1:mgmt:~$ sudo onie-install -a -i
http://10.60.132.97/x/eng/testbedN,svl/nic/files/NVIDIA/cumulus-
linux-5.4.0-mlx-amd64.bin/
[sudo] password for cumulus:
Fetching installer:
http://10.60.132.97/x/eng/testbedN,svl/nic/files/NVIDIA/cumulus-
linux-5.4.0-mlx-amd64.bin
Downloading URL:
http://10.60.132.97/x/eng/testbedN,svl/nic/files/NVIDIA/cumulus-
linux-5.4.0-mlx-amd64.bin
# 100.0%
Success: HTTP download complete.
EFI variables are not supported on this system
Warning: SecureBoot is not available.
Image is signed.
.
.
.
Staging installer image...done.
WARNING:
WARNING: Activating staged installer requested.
WARNING: This action will wipe out all system data.
WARNING: Make sure to back up your data.
WARNING:
Are you sure (y/N)? y
Activating staged installer...done.
Reboot required to take effect.
```

3. Reboot the switch:

```
admin@sw1:mgmt:~$ sudo onie-install -a -i
http://10.60.132.97/x/eng/testbedN,svl/nic/files/NVIDIA/cumulus-
linux-5.4.0-mlx-amd64.bin/
sudo reboot
```

4. Change the password:

```
cumulus login: cumulus
Password:
You are required to change your password immediately (administrator
enforced)
Changing password for cumulus.
Current password: cumulus
New password: <new_password>
Retype new password: <new_password>
Linux cumulus 5.10.0-cl-1-amd64 #1 SMP Debian 5.10.162-1+cl5.4.0u1
(2023-01-20) x86_64

Welcome to NVIDIA Cumulus (R) Linux (R)

ZTP in progress. To disable, do 'ztp -d'
```

5. Check the Cumulus Linux version: `nv show system`

```
cumulus@cumulus:mgmt:~$ nv show system
```

	operational	applied
hostname	cumulus	cumulus
build	Cumulus Linux 5.4.0	
uptime	14:07:08	
timezone	Etc/UTC	

6. Change the hostname:

```
cumulus@cumulus:mgmt:~$ nv set system hostname sw1
cumulus@cumulus:mgmt:~$ nv config apply
Warning: The following files have been changed since the last save,
and they WILL be overwritten.
- /etc/nsswitch.conf
- /etc/syncd/syncd.conf
.
.
```

7. Logout and log in to the switch again to see the updated switch name at the prompt:

```
cumulus@cumulus:mgmt:~$ exit
logout

Debian GNU/Linux 10 cumulus ttyS0

cumulus login: cumulus
Password:
Last login: Tue Dec 15 21:43:13 UTC 2020 on ttyS0
Linux cumulus 5.10.0-cl-1-amd64 #1 SMP Debian 5.10.162-1+cl5.4.0u1
(2023-01-20) x86_64

Welcome to NVIDIA Cumulus (R) Linux (R)

ZTP in progress. To disable, do 'ztp -d'

cumulus@sw1:mgmt:~$
```

8. Set the IP address:

```
cumulus@sw1:mgmt:~$ nv set interface eth0 ip address 10.231.80.206
cumulus@sw1:mgmt:~$ nv set interface eth0 ip gateway 10.231.80.1
cumulus@sw1:mgmt:~$ nv config apply
applied [rev_id: 2]
cumulus@sw1:mgmt:~$ ip route show vrf mgmt
default via 10.231.80.1 dev eth0 proto kernel
unreachable default metric 4278198272
10.231.80.0/22 dev eth0 proto kernel scope link src 10.231.80.206
127.0.0.0/8 dev mgmt proto kernel scope link src 127.0.0.1
```

9. Create a new user and add this user to the `sudo` group. This user only becomes effective after the console/SSH session is restarted.

```
sudo adduser --ingroup netedit admin
```



```

cumulus@sw1:mgmt:~$ sudo adduser --ingroup netedit admin
[sudo] password for cumulus:
Adding user 'admin' ...
Adding new user 'admin' (1001) with group `netedit' ...
Creating home directory '/home/admin' ...
Copying files from '/etc/skel' ...
New password:
Retype new password:
passwd: password updated successfully
Changing the user information for admin
Enter the new value, or press ENTER for the default
Full Name []:
Room Number []:
Work Phone []:
Home Phone []:
Other []:
Is the information correct? [Y/n] y

cumulus@sw1:mgmt:~$ sudo adduser admin sudo
[sudo] password for cumulus:
Adding user `admin' to group `sudo' ...
Adding user admin to group sudo
Done.
cumulus@sw1:mgmt:~$ exit
logout
Connection to 10.233.204.71 closed.

[admin@cycrh6svl01 ~]$ ssh admin@10.233.204.71
admin@10.233.204.71's password:
Linux sw1 4.19.0-cl-1-amd64 #1 SMP Cumulus 4.19.206-1+cl4.4.1u1
(2021-09-09) x86_64
Welcome to NVIDIA Cumulus (R) Linux (R)

For support and online technical documentation, visit
http://www.cumulusnetworks.com/support

The registered trademark Linux (R) is used pursuant to a sublicense
from LMI, the exclusive licensee of Linus Torvalds, owner of the
mark on a world-wide basis.
admin@sw1:mgmt:~$

```

10. Add additional user groups for the admin user to access `nv` commands:

```
cumulus@sw1:mgmt:~$ sudo adduser admin nvshow
[sudo] password for cumulus:
Adding user `admin' to group `nvshow' ...
Adding user admin to group nvshow
Done.
```

See [NVIDIA User Accounts](#) for more information.

From Cumulus Linux 5.x to Cumulus Linux 5.x

1. Check the current Cumulus Linux version and connected ports:

```
admin@sw1:mgmt:~$ nv show system
```

	operational	applied
hostname	cumulus	cumulus
build	Cumulus Linux 5.3.0	
uptime	6 days, 8:37:36	
timezone	Etc/UTC	

```
admin@sw1:mgmt:~$ nv show interface
```

Interface	MTU	Speed	State	Remote Host	Remote Port-
Type	Summary				

+ cluster_isl	9216	200G	up		
bond					
+ eth0	1500	100M	up	mgmt-sw1	Eth105/1/14
eth	IP Address: 10.231.80 206/22				
eth0					
IP Address: fd20:8b1e:f6ff:fe31:4a0e/64					
+ lo	65536		up		
loopback	IP Address: 127.0.0.1/8				
lo					
IP Address: ::1/128					
+ swp1s0	9216	10G	up	cluster01	e0b
swp					
.					
.					
.					
+ swp15	9216	100G	up	sw2	swp15
swp					
+ swp16	9216	100G	up	sw2	swp16
swp					

2. Download the Cumulux Linux 5.4.0 image:

```
admin@sw1:mgmt:~$ sudo onie-install -a -i
http://10.60.132.97/x/eng/testbedN,svl/nic/files/NVIDIA/cumulus-
linux-5.4.0-mlx-amd64.bin/
[sudo] password for cumulus:
Fetching installer:
http://10.60.132.97/x/eng/testbedN,svl/nic/files/NVIDIA/cumulus-
linux-5.4.0-mlx-amd64.bin
Downloading URL:
http://10.60.132.97/x/eng/testbedN,svl/nic/files/NVIDIA/cumulus-
linux-5.4.0-mlx-amd64.bin
# 100.0%
Success: HTTP download complete.
EFI variables are not supported on this system
Warning: SecureBoot is not available.
Image is signed.
.
.
.
Staging installer image...done.
WARNING:
WARNING: Activating staged installer requested.
WARNING: This action will wipe out all system data.
WARNING: Make sure to back up your data.
WARNING:
Are you sure (y/N)? y
Activating staged installer...done.
Reboot required to take effect.
```

3. Reboot the switch:

```
admin@sw1:mgmt:~$ sudo reboot
```

4. Change the password:

```
cumulus login: cumulus
Password:
You are required to change your password immediately (administrator
enforced)
Changing password for cumulus.
Current password: cumulus
New password: <new_password>
Retype new password: <new_password>
Linux cumulus 5.10.0-cl-1-amd64 #1 SMP Debian 5.10.162-1+cl5.4.0u1
(2023-01-20) x86_64

Welcome to NVIDIA Cumulus (R) Linux (R)

ZTP in progress. To disable, do 'ztp -d'
```

5. Check the Cumulus Linux version: `nv show system`

```
cumulus@cumulus:mgmt:~$ nv show system
operational      applied
-----
hostname         cumulus cumulus
build            Cumulus Linux 5.4.0
uptime           14:07:08
timezone         Etc/UTC
```

6. Change the hostname:

```
cumulus@cumulus:mgmt:~$ nv set system hostname sw1
cumulus@cumulus:mgmt:~$ nv config apply
Warning: The following files have been changed since the last save,
and they WILL be overwritten.
- /etc/nsswitch.conf
- /etc/syncd/syncd.conf
.
.
```

7. Logout and log in again to the switch to see the updated switch name at the prompt:

```
cumulus@cumulus:mgmt:~$ exit
logout

Debian GNU/Linux 10 cumulus ttyS0

cumulus login: cumulus
Password:
Last login: Tue Dec 15 21:43:13 UTC 2020 on ttyS0
Linux cumulus 5.10.0-cl-1-amd64 #1 SMP Debian 5.10.162-1+cl5.4.0u1
(2023-01-20) x86_64

Welcome to NVIDIA Cumulus (R) Linux (R)

ZTP in progress. To disable, do 'ztp -d'

cumulus@sw1:mgmt:~$
```

8. Set the IP address:

```
cumulus@sw1:mgmt:~$ nv set interface eth0 ip address 10.231.80.206
cumulus@sw1:mgmt:~$ nv set interface eth0 ip gateway 10.231.80.1
cumulus@sw1:mgmt:~$ nv config apply
applied [rev_id: 2]
cumulus@sw1:mgmt:~$ ip route show vrf mgmt
default via 10.231.80.1 dev eth0 proto kernel
unreachable default metric 4278198272
10.231.80.0/22 dev eth0 proto kernel scope link src 10.231.80.206
127.0.0.0/8 dev mgmt proto kernel scope link src 127.0.0.1
```

9. Create a new user and add this user to the `sudo` group. This user only becomes effective after the console/SSH session is restarted.

```
sudo adduser --ingroup netedit admin
```

```

cumulus@sw1:mgmt:~$ sudo adduser --ingroup netedit admin
[sudo] password for cumulus:
Adding user 'admin' ...
Adding new user 'admin' (1001) with group `netedit' ...
Creating home directory '/home/admin' ...
Copying files from '/etc/skel' ...
New password:
Retype new password:
passwd: password updated successfully
Changing the user information for admin
Enter the new value, or press ENTER for the default
Full Name []:
Room Number []:
Work Phone []:
Home Phone []:
Other []:
Is the information correct? [Y/n] y

cumulus@sw1:mgmt:~$ sudo adduser admin sudo
[sudo] password for cumulus:
Adding user `admin' to group `sudo' ...
Adding user admin to group sudo
Done.
cumulus@sw1:mgmt:~$ exit
logout
Connection to 10.233.204.71 closed.

[admin@cycrh6svl01 ~]$ ssh admin@10.233.204.71
admin@10.233.204.71's password:
Linux sw1 4.19.0-cl-1-amd64 #1 SMP Cumulus 4.19.206-1+cl4.4.1u1
(2021-09-09) x86_64
Welcome to NVIDIA Cumulus (R) Linux (R)

For support and online technical documentation, visit
http://www.cumulusnetworks.com/support

The registered trademark Linux (R) is used pursuant to a sublicense
from LMI, the exclusive licensee of Linus Torvalds, owner of the
mark on a world-wide basis.
admin@sw1:mgmt:~$

```

10. Add additional user groups for the admin user to access `nv` commands:

```
cumulus@sw1:mgmt:~$ sudo adduser admin nvshow
[sudo] password for cumulus:
Adding user `admin' to group `nvshow' ...
Adding user admin to group nvshow
Done.
```

See [NVIDIA User Accounts](#) for more information.

What's next?

Install the [Reference Configuration File \(RCF\)](#) script.

Migrate switches

Migrate CN1610 cluster switches to NVIDIA SN2100 cluster switches

You can migrate NetApp CN1610 cluster switches for an ONTAP cluster to NVIDIA SN2100 cluster switches. This is a nondisruptive procedure.

Review requirements

You must be aware of certain configuration information, port connections and cabling requirements when you are replacing NetApp CN1610 cluster switches with NVIDIA SN2100 cluster switches. See [Overview of installation and configuration for NVIDIA SN2100 switches](#).

Supported switches

The following cluster switches are supported:

- NetApp CN1610
- NVIDIA SN2100

For details of supported ports and their configurations, see the [Hardware Universe](#).

What you'll need

Verify that you meet the following requirements for you configuration:

- The existing cluster is correctly set up and functioning.
- All cluster ports are in the **up** state to ensure nondisruptive operations.
- The NVIDIA SN2100 cluster switches are configured and operating under the correct version of Cumulus Linux installed with the reference configuration file (RCF) applied.
- The existing cluster network configuration has the following:
 - A redundant and fully functional NetApp cluster using CN1610 switches.
 - Management connectivity and console access to both the CN1610 switches and the new switches.
 - All cluster LIFs in the up state with the cluster Lifs on their home ports.
 - ISL ports enabled and cabled between the CN1610 switches and between the new switches.

- Some of the ports are configured on NVIDIA SN2100 switches to run at 40GbE or 100GbE.
- You have planned, migrated, and documented 40GbE and 100GbE connectivity from nodes to NVIDIA SN2100 cluster switches.

Migrate the switches

About the examples

The examples in this procedure use the following switch and node nomenclature:

- The existing CN1610 cluster switches are *c1* and *c2*.
- The new NVIDIA SN2100 cluster switches are *sw1* and *sw2*.
- The nodes are *node1* and *node2*.
- The cluster LIFs are *node1_clus1* and *node1_clus2* on node 1, and *node2_clus1* and *node2_clus2* on node 2 respectively.
- The `cluster1::*>` prompt indicates the name of the cluster.
- The cluster ports used in this procedure are *e3a* and *e3b*.
- Breakout ports take the format: `swp[port]s[breakout port 0-3]`. For example, four breakout ports on `swp1` are *swp1s0*, *swp1s1*, *swp1s2*, and *swp1s3*.

About this task

This procedure covers the following scenario:

- Switch *c2* is replaced by switch *sw2* first.
 - Shut down the ports to the cluster nodes. All ports must be shut down simultaneously to avoid cluster instability.
 - The cabling between the nodes and *c2* is then disconnected from *c2* and reconnected to *sw2*.
- Switch *c1* is replaced by switch *sw1*.
 - Shut down the ports to the cluster nodes. All ports must be shut down simultaneously to avoid cluster instability.
 - The cabling between the nodes and *c1* is then disconnected from *c1* and reconnected to *sw1*.



No operational inter-switch link (ISL) is needed during this procedure. This is by design because RCF version changes can affect ISL connectivity temporarily. To ensure non-disruptive cluster operations, the following procedure migrates all of the cluster LIFs to the operational partner switch while performing the steps on the target switch.

Step 1: Prepare for migration

1. If AutoSupport is enabled on this cluster, suppress automatic case creation by invoking an AutoSupport message:

```
system node autosupport invoke -node * -type all -message MAINT=xh
```

where *x* is the duration of the maintenance window in hours.

2. Change the privilege level to advanced, entering *y* when prompted to continue:

```
set -privilege advanced
```


The advanced prompt (*>) appears.

3. Disable auto-revert on the cluster LIFs:

```
network interface modify -vserver Cluster -lif * -auto-revert false
```

Step 2: Configure ports and cabling

1. Determine the administrative or operational status for each cluster interface.

Each port should display up for `Link` and `healthy` for `Health Status`.

- a. Display the network port attributes:

```
network port show -ipspace Cluster
```

Show example

```
cluster1::*> network port show -ipspace Cluster
```

Node: node1

Ignore

Health	Health				Speed (Mbps)
Port	IPspace	Broadcast	Domain	Link	MTU
Status	Status				Admin/Oper
-----	-----	-----	----	----	-----
e3a	Cluster	Cluster		up	9000
healthy	false				auto/100000
e3b	Cluster	Cluster		up	9000
healthy	false				auto/100000

Node: node2

Ignore

Health	Health				Speed (Mbps)
Port	IPspace	Broadcast	Domain	Link	MTU
Status	Status				Admin/Oper
-----	-----	-----	----	----	-----
-----	-----				
e3a	Cluster	Cluster		up	9000
healthy	false				auto/100000
e3b	Cluster	Cluster		up	9000
healthy	false				auto/100000

b. Display information about the LIFs and their designated home nodes:

```
network interface show -vserver Cluster
```

Each LIF should display up/up for Status Admin/Oper and true for Is Home.

Show example

```
cluster1::*> network interface show -vserver Cluster
```

	Logical	Status	Network	Current
Current Is				
Vserver	Interface	Admin/Oper	Address/Mask	Node
Port	Home			

Cluster				
	node1_clus1	up/up	169.254.209.69/16	node1
e3a	true			
	node1_clus2	up/up	169.254.49.125/16	node1
e3b	true			
	node2_clus1	up/up	169.254.47.194/16	node2
e3a	true			
	node2_clus2	up/up	169.254.19.183/16	node2
e3b	true			

2. The cluster ports on each node are connected to existing cluster switches in the following way (from the nodes' perspective) using the command:

```
network device-discovery show -protocol
```

Show example

```
cluster1::*> network device-discovery show -protocol cdp
```

Node/	Local	Discovered		
Protocol	Port	Device (LLDP: ChassisID)	Interface	
Platform				

node1	/cdp			
	e3a	c1 (6a:ad:4f:98:3b:3f)	0/1	-
	e3b	c2 (6a:ad:4f:98:4c:a4)	0/1	-
node2	/cdp			
	e3a	c1 (6a:ad:4f:98:3b:3f)	0/2	-
	e3b	c2 (6a:ad:4f:98:4c:a4)	0/2	-

3. The cluster ports and switches are connected in the following way (from the switches' perspective) using the command:

```
show cdp neighbors
```

Show example

A large, empty rectangular box with a dashed border, intended for an example.

c1# **show cdp neighbors**

Capability Codes: R - Router, T - Trans-Bridge, B - Source-Route-Bridge

S - Switch, H - Host, I - IGMP, r - Repeater,
V - VoIP-Phone, D - Remotely-Managed-Device,
s - Supports-STP-Dispute

Device-ID Port ID	Local Intrfce	Hldtme	Capability	Platform
node1 e3a	0/1	124	H	AFF-A400
node2 e3a	0/2	124	H	AFF-A400
c2 0/13	0/13	179	S I s	CN1610
c2 0/14	0/14	175	S I s	CN1610
c2 0/15	0/15	179	S I s	CN1610
c2 0/16	0/16	175	S I s	CN1610

c2# **show cdp neighbors**

Capability Codes: R - Router, T - Trans-Bridge, B - Source-Route-Bridge

S - Switch, H - Host, I - IGMP, r - Repeater,
V - VoIP-Phone, D - Remotely-Managed-Device,
s - Supports-STP-Dispute

Device-ID Port ID	Local Intrfce	Hldtme	Capability	Platform
node1 e3b	0/1	124	H	AFF-A400
node2 e3b	0/2	124	H	AFF-A400
c1 0/13	0/13	175	S I s	CN1610
c1 0/14	0/14	175	S I s	CN1610
c1 0/15	0/15	175	S I s	CN1610
c1 0/16	0/16	175	S I s	CN1610

4. Verify that the cluster network has full connectivity:

```
cluster ping-cluster -node node-name
```

Show example

```
cluster1::*> cluster ping-cluster -node node2

Host is node2
Getting addresses from network interface table...
Cluster node1_clus1 169.254.209.69 node1      e3a
Cluster node1_clus2 169.254.49.125 node1      e3b
Cluster node2_clus1 169.254.47.194 node2      e3a
Cluster node2_clus2 169.254.19.183 node2      e3b
Local = 169.254.47.194 169.254.19.183
Remote = 169.254.209.69 169.254.49.125
Cluster Vserver Id = 4294967293
Ping status:
....
Basic connectivity succeeds on 4 path(s)
Basic connectivity fails on 0 path(s)
.....
Detected 9000 byte MTU on 4 path(s):
    Local 169.254.19.183 to Remote 169.254.209.69
    Local 169.254.19.183 to Remote 169.254.49.125
    Local 169.254.47.194 to Remote 169.254.209.69
    Local 169.254.47.194 to Remote 169.254.49.125
Larger than PMTU communication succeeds on 4 path(s)
RPC status:
2 paths up, 0 paths down (tcp check)
2 paths up, 0 paths down (udp check)
```

5. On switch c2, shut down the ports connected to the cluster ports of the nodes in order to fail over the cluster LIFs.

```
(c2)# configure
(c2)(Config)# interface 0/1-0/12
(c2)(Interface 0/1-0/12)# shutdown
(c2)(Interface 0/1-0/12)# exit
(c2)(Config)# exit
(c2)#
```

6. Move the node cluster ports from the old switch c2 to the new switch sw2, using appropriate cabling supported by NVIDIA SN2100.

7. Display the network port attributes:

```
network port show -ipspace Cluster
```

Show example

```
cluster1::*> network port show -ipspace Cluster
```

Node: node1

Ignore

						Speed (Mbps)	Health
Health							
Port	IPspace	Broadcast	Domain	Link	MTU	Admin/Oper	Status
Status							
-----	-----	-----	-----	-----	-----	-----	-----
e3a	Cluster	Cluster		up	9000	auto/100000	
healthy	false						
e3b	Cluster	Cluster		up	9000	auto/100000	
healthy	false						

Node: node2

Ignore

						Speed (Mbps)	Health
Health							
Port	IPspace	Broadcast	Domain	Link	MTU	Admin/Oper	Status
Status							
-----	-----	-----	-----	-----	-----	-----	-----
e3a	Cluster	Cluster		up	9000	auto/100000	
healthy	false						
e3b	Cluster	Cluster		up	9000	auto/100000	
healthy	false						

8. The cluster ports on each node are now connected to cluster switches in the following way, from the nodes' perspective:

```
network device-discovery show -protocol
```

Show example

```
cluster1::*> network device-discovery show -protocol lldp
```

Node/ Protocol Platform	Local Port	Discovered Device (LLDP: ChassisID)	Interface	

node1	/lldp			
	e3a	c1 (6a:ad:4f:98:3b:3f)	0/1	-
	e3b	sw2 (b8:ce:f6:19:1a:7e)	swp3	-
node2	/lldp			
	e3a	c1 (6a:ad:4f:98:3b:3f)	0/2	-
	e3b	sw2 (b8:ce:f6:19:1b:96)	swp4	-

9. On switch sw2, verify that all node cluster ports are up:

```
net show interface
```

Show example

```
cumulus@sw2:~$ net show interface
```

State	Name	Spd	MTU	Mode	LLDP
Summary					

...					
...					
UP	swp3	100G	9216	Trunk/L2	e3b
Master: bridge(UP)					
UP	swp4	100G	9216	Trunk/L2	e3b
Master: bridge(UP)					
UP	swp15	100G	9216	BondMember	sw1 (swp15)
Master: cluster_isl(UP)					
UP	swp16	100G	9216	BondMember	sw1 (swp16)
Master: cluster_isl(UP)					

10. On switch c1, shut down the ports connected to the cluster ports of the nodes in order to fail over the cluster LIFs.


```
(c1)# configure
(c1) (Config)# interface 0/1-0/12
(c1) (Interface 0/1-0/12)# shutdown
(c1) (Interface 0/1-0/12)# exit
(c1) (Config)# exit
(c1)#
```

11. Move the node cluster ports from the old switch c1 to the new switch sw1, using appropriate cabling supported by NVIDIA SN2100.
12. Verify the final configuration of the cluster:

```
network port show -ipspace Cluster
```

Each port should display up for Link and healthy for Health Status.

Show example

```
cluster1::*> network port show -ipspace Cluster
```

Node: node1

Ignore

						Speed (Mbps)	Health
Health							
Port	IPspace	Broadcast	Domain	Link	MTU	Admin/Oper	Status
Status							
-----	-----	-----	----	-----	-----	-----	
-----	-----						
e3a	Cluster	Cluster		up	9000	auto/100000	
healthy	false						
e3b	Cluster	Cluster		up	9000	auto/100000	
healthy	false						

Node: node2

Ignore

						Speed (Mbps)	Health
Health							
Port	IPspace	Broadcast	Domain	Link	MTU	Admin/Oper	Status
Status							
-----	-----	-----	----	-----	-----	-----	
-----	-----						
e3a	Cluster	Cluster		up	9000	auto/100000	
healthy	false						
e3b	Cluster	Cluster		up	9000	auto/100000	
healthy	false						

13. The cluster ports on each node are now connected to cluster switches in the following way, from the nodes' perspective:

```
network device-discovery show -protocol
```

Show example

```
cluster1::*> network device-discovery show -protocol lldp
```

Node/ Protocol Platform	Local Port	Discovered Device (LLDP: ChassisID)	Interface	

node1	/lldp			
	e3a	sw1 (b8:ce:f6:19:1a:7e)	swp3	-
	e3b	sw2 (b8:ce:f6:19:1b:96)	swp3	-
node2	/lldp			
	e3a	sw1 (b8:ce:f6:19:1a:7e)	swp4	-
	e3b	sw2 (b8:ce:f6:19:1b:96)	swp4	-

14. On switches sw1 and sw2, verify that all node cluster ports are up:

```
net show interface
```

Show example

```
cumulus@sw1:~$ net show interface
```

State	Name	Spd	MTU	Mode	LLDP
Summary					

...					
...					
UP	swp3	100G	9216	Trunk/L2	e3a
Master: bridge(UP)					
UP	swp4	100G	9216	Trunk/L2	e3a
Master: bridge(UP)					
UP	swp15	100G	9216	BondMember	sw2 (swp15)
Master: cluster_isl(UP)					
UP	swp16	100G	9216	BondMember	sw2 (swp16)
Master: cluster_isl(UP)					

```
cumulus@sw2:~$ net show interface
```

State	Name	Spd	MTU	Mode	LLDP
Summary					

...					
...					
UP	swp3	100G	9216	Trunk/L2	e3b
Master: bridge(UP)					
UP	swp4	100G	9216	Trunk/L2	e3b
Master: bridge(UP)					
UP	swp15	100G	9216	BondMember	sw1 (swp15)
Master: cluster_isl(UP)					
UP	swp16	100G	9216	BondMember	sw1 (swp16)
Master: cluster_isl(UP)					

15. Verify that both nodes each have one connection to each switch:

```
net show lldp
```

Show example

The following example shows the appropriate results for both switches:

```
cumulus@sw1:~$ net show lldp
```

LocalPort	Speed	Mode	RemoteHost	RemotePort
-----	-----	-----	-----	-----
swp3	100G	Trunk/L2	node1	e3a
swp4	100G	Trunk/L2	node2	e3a
swp15	100G	BondMember	sw2	swp15
swp16	100G	BondMember	sw2	swp16

```
cumulus@sw2:~$ net show lldp
```

LocalPort	Speed	Mode	RemoteHost	RemotePort
-----	-----	-----	-----	-----
swp3	100G	Trunk/L2	node1	e3b
swp4	100G	Trunk/L2	node2	e3b
swp15	100G	BondMember	sw1	swp15
swp16	100G	BondMember	sw1	swp16

Step 3: Complete the procedure

1. Enable auto-revert on the cluster LIFs:

```
cluster1::*> network interface modify -vserver Cluster -lif * -auto-revert  
true
```

2. Verify that all cluster network LIFs are back on their home ports:

```
network interface show
```

Show example

```
cluster1::*> network interface show -vserver Cluster
```

		Logical	Status	Network	Current
Current Is					
Vserver	Interface	Admin/Oper	Address/Mask		Node
Port	Home				

Cluster					
		node1_clus1	up/up	169.254.209.69/16	node1
e3a	true				
		node1_clus2	up/up	169.254.49.125/16	node1
e3b	true				
		node2_clus1	up/up	169.254.47.194/16	node2
e3a	true				
		node2_clus2	up/up	169.254.19.183/16	node2
e3b	true				

3. To set up log collection, run the following command for each switch. You are prompted to enter the switch name, username, and password for log collection.

```
system switch ethernet log setup-password
```

Show example

```
cluster1::*> system switch ethernet log setup-password
Enter the switch name: <return>
The switch name entered is not recognized.
Choose from the following list:
sw1
sw2

cluster1::*> system switch ethernet log setup-password

Enter the switch name: sw1
Would you like to specify a user other than admin for log
collection? {y|n}: n

Enter the password: <enter switch password>
Enter the password again: <enter switch password>

cluster1::*> system switch ethernet log setup-password

Enter the switch name: sw2
Would you like to specify a user other than admin for log
collection? {y|n}: n

Enter the password: <enter switch password>
Enter the password again: <enter switch password>
```

4. To start log collection, run the following command, replacing **DEVICE** with the switch used in the previous command. This starts both types of log collection: the detailed **Support** logs and an hourly collection of **Periodic** data.

```
system switch ethernet log modify -device <switch-name> -log-request true
```

Show example

```
cluster1::*> system switch ethernet log modify -device cs1 -log  
-request true
```

Do you want to modify the cluster switch log collection
configuration? {y|n}: [n] **y**

Enabling cluster switch log collection.

```
cluster1::*> system switch ethernet log modify -device cs2 -log  
-request true
```

Do you want to modify the cluster switch log collection
configuration? {y|n}: [n] **y**

Enabling cluster switch log collection.

Wait for 10 minutes and then check that the log collection completes:

```
system switch ethernet log show
```

Show example

```
cluster1::*> system switch ethernet log show  
Log Collection Enabled: true
```

Index	Switch	Log Timestamp	Status
-----	-----	-----	-----
1	cs1 (b8:ce:f6:19:1b:42)	4/29/2022 03:05:25	complete
2	cs2 (b8:ce:f6:19:1b:96)	4/29/2022 03:07:42	complete



If any of these commands return an error or if the log collection does not complete, contact NetApp support.

5. Change the privilege level back to admin:

```
set -privilege admin
```

6. If you suppressed automatic case creation, re-enable it by invoking an AutoSupport message:

```
system node autosupport invoke -node * -type all -message MAINT=END
```


Migrate from a Cisco cluster switch to a NVIDIA SN2100 cluster switch

You can migrate Cisco cluster switches for an ONTAP cluster to NVIDIA SN2100 cluster switches. This is a nondisruptive procedure.

Review requirements

You must be aware of certain configuration information, port connections and cabling requirements when you are replacing some older Cisco cluster switches with NVIDIA SN2100 cluster switches. See [Overview of installation and configuration for NVIDIA SN2100 switches](#).

Supported switches

The following Cisco cluster switches are supported:

- Nexus 9336C-FX2
- Nexus 92300YC
- Nexus 5596UP
- Nexus 3232C
- Nexus 3132Q-V

For details of supported ports and their configurations, see the [Hardware Universe](#).

What you'll need

Ensure that:

- The existing cluster is properly set up and functioning.
- All cluster ports are in the **up** state to ensure nondisruptive operations.
- The NVIDIA SN2100 cluster switches are configured and operating under the proper version of Cumulus Linux installed with the reference configuration file (RCF) applied.
- The existing cluster network configuration have the following:
 - A redundant and fully functional NetApp cluster using both older Cisco switches.
 - Management connectivity and console access to both the older Cisco switches and the new switches.
 - All cluster LIFs in the up state with the cluster Lifs are on their home ports.
 - ISL ports enabled and cabled between the older Cisco switches and between the new switches.
- Some of the ports are configured on NVIDIA SN2100 switches to run at 40 GbE or 100 GbE.
- You have planned, migrated, and documented 40 GbE and 100 GbE connectivity from nodes to NVIDIA SN2100 cluster switches.



If you are changing the port speed of the e0a and e1a cluster ports on AFF A800 or AFF C800 systems, you might observe malformed packets being received after the speed conversion. See [Bug 1570339](#) and the Knowledge Base article [CRC errors on T6 ports after converting from 40GbE to 100GbE](#) for guidance.

Migrate the switches

About the examples

In this procedure, Cisco Nexus 3232C cluster switches are used for example commands and outputs.

The examples in this procedure use the following switch and node nomenclature:

- The existing Cisco Nexus 3232C cluster switches are *c1* and *c2*.
- The new NVIDIA SN2100 cluster switches are *sw1* and *sw2*.
- The nodes are *node1* and *node2*.
- The cluster LIFs are *node1_clus1* and *node1_clus2* on node 1, and *node2_clus1* and *node2_clus2* on node 2 respectively.
- The `cluster1: *>` prompt indicates the name of the cluster.
- The cluster ports used in this procedure are *e3a* and *e3b*.
- Breakout ports take the format: `swp[port]s[breakout port 0-3]`. For example, four breakout ports on `swp1` are *swp1s0*, *swp1s1*, *swp1s2*, and *swp1s3*.

About this task

This procedure covers the following scenario:

- Switch *c2* is replaced by switch *sw2* first.
 - Shut down the ports to the cluster nodes. All ports must be shut down simultaneously to avoid cluster instability.
 - Cabling between the nodes and *c2* are then disconnected from *c2* and reconnected to *sw2*.
- Switch *c1* is replaced by switch *sw1*.
 - Shut down the ports to the cluster nodes. All ports must be shut down simultaneously to avoid cluster instability.
 - Cabling between the nodes and *c1* are then disconnected from *c1* and reconnected to *sw1*.

Step 1: Prepare for migration

1. If AutoSupport is enabled on this cluster, suppress automatic case creation by invoking an AutoSupport message:

```
system node autosupport invoke -node * -type all -message MAINT=xh
```

where *x* is the duration of the maintenance window in hours.

2. Change the privilege level to advanced, entering *y* when prompted to continue:

```
set -privilege advanced
```

The advanced prompt (**>*) appears.

3. Disable auto-revert on the cluster LIFs:

```
network interface modify -vserver Cluster -lif * -auto-revert false
```

Step 2: Configure ports and cabling

1. Determine the administrative or operational status for each cluster interface.

Each port should display up for `Link` and healthy for `Health Status`.

- a. Display the network port attributes:

```
network port show -ipspace Cluster
```

Show example

```
cluster1::*> network port show -ipspace Cluster
```

Node: node1

Ignore

Health	Health					Speed (Mbps)
Port	IPspace	Broadcast	Domain	Link	MTU	Admin/Oper
Status	Status					
-----	-----	-----	----	-----	-----	-----
-----	-----					
e3a	Cluster	Cluster		up	9000	auto/100000
healthy	false					
e3b	Cluster	Cluster		up	9000	auto/100000
healthy	false					

Node: node2

Ignore

Health	Health					Speed (Mbps)
Port	IPspace	Broadcast	Domain	Link	MTU	Admin/Oper
Status	Status					
-----	-----	-----	----	-----	-----	-----
-----	-----					
e3a	Cluster	Cluster		up	9000	auto/100000
healthy	false					
e3b	Cluster	Cluster		up	9000	auto/100000
healthy	false					

- b. Display information about the logical interfaces and their designated home nodes:

```
network interface show -vserver Cluster
```

Each LIF should display up/up for Status Admin/Oper and true for Is Home.

Show example

```
cluster1::*> network interface show -vserver Cluster
```

	Logical	Status	Network	Current
Current Is				
Vserver	Interface	Admin/Oper	Address/Mask	Node
Port	Home			

Cluster				
e3a	node1_clus1	up/up	169.254.209.69/16	node1
e3b	node1_clus2	up/up	169.254.49.125/16	node1
e3a	node2_clus1	up/up	169.254.47.194/16	node2
e3b	node2_clus2	up/up	169.254.19.183/16	node2

2. The cluster ports on each node are connected to existing cluster switches in the following way (from the nodes' perspective):

```
network device-discovery show -protocol lldp
```

Show example

```
cluster1::*> network device-discovery show -protocol lldp
```

Node/	Local	Discovered	
Protocol	Port	Device (LLDP: ChassisID)	Interface
Platform			

node1	/lldp		
e3a	c1	(6a:ad:4f:98:3b:3f)	Eth1/1
e3b	c2	(6a:ad:4f:98:4c:a4)	Eth1/1
node2	/lldp		
e3a	c1	(6a:ad:4f:98:3b:3f)	Eth1/2
e3b	c2	(6a:ad:4f:98:4c:a4)	Eth1/2

3. The cluster ports and switches are connected in the following way (from the switches' perspective):

```
show cdp neighbors
```

Show example

```
c1# show cdp neighbors
```

Capability Codes: R - Router, T - Trans-Bridge, B - Source-Route-Bridge

S - Switch, H - Host, I - IGMP, r - Repeater,
V - VoIP-Phone, D - Remotely-Managed-Device,
s - Supports-STP-Dispute

Device-ID Port ID	Local Intrfce	Hldtme	Capability	Platform
node1 e3a	Eth1/1	124	H	AFF-A400
node2 e3a	Eth1/2	124	H	AFF-A400
c2 Eth1/31	Eth1/31	179	S I s	N3K-C3232C
c2 Eth1/32	Eth1/32	175	S I s	N3K-C3232C

```
c2# show cdp neighbors
```

Capability Codes: R - Router, T - Trans-Bridge, B - Source-Route-Bridge

S - Switch, H - Host, I - IGMP, r - Repeater,
V - VoIP-Phone, D - Remotely-Managed-Device,
s - Supports-STP-Dispute

Device-ID Port ID	Local Intrfce	Hldtme	Capability	Platform
node1 e3b	Eth1/1	124	H	AFF-A400
node2 e3b	Eth1/2	124	H	AFF-A400
c1 Eth1/31	Eth1/31	175	S I s	N3K-C3232C
c1 Eth1/32	Eth1/32	175	S I s	N3K-C3232C

4. Ensure that the cluster network has full connectivity:

```
cluster ping-cluster -node node-name
```

Show example

```
cluster1::~*> cluster ping-cluster -node node2

Host is node2
Getting addresses from network interface table...
Cluster node1_clus1 169.254.209.69 node1      e3a
Cluster node1_clus2 169.254.49.125 node1      e3b
Cluster node2_clus1 169.254.47.194 node2      e3a
Cluster node2_clus2 169.254.19.183 node2      e3b
Local = 169.254.47.194 169.254.19.183
Remote = 169.254.209.69 169.254.49.125
Cluster Vserver Id = 4294967293
Ping status:
....
Basic connectivity succeeds on 4 path(s)
Basic connectivity fails on 0 path(s)
.....
Detected 9000 byte MTU on 4 path(s):
    Local 169.254.19.183 to Remote 169.254.209.69
    Local 169.254.19.183 to Remote 169.254.49.125
    Local 169.254.47.194 to Remote 169.254.209.69
    Local 169.254.47.194 to Remote 169.254.49.125
Larger than PMTU communication succeeds on 4 path(s)
RPC status:
2 paths up, 0 paths down (tcp check)
2 paths up, 0 paths down (udp check)
```

5. On switch c2, shut down the ports connected to the cluster ports of the nodes in order to fail over the cluster LIFs.

```
(c2)# configure
Enter configuration commands, one per line. End with CNTL/Z.

(c2) (Config)# interface
(c2) (config-if-range)# shutdown <interface_list>
(c2) (config-if-range)# exit
(c2) (Config)# exit
(c2)#
```

6. Move the node cluster ports from the old switch c2 to the new switch sw2, using appropriate cabling supported by NVIDIA SN2100.
7. Display the network port attributes:

network port show -ipspace Cluster

Show example

```
cluster1::*> network port show -ipspace Cluster
```

Node: node1

Ignore

						Speed (Mbps)	Health
Health							
Port	IPspace	Broadcast	Domain	Link	MTU	Admin/Oper	Status
Status							
-----	-----	-----	-----	-----	-----	-----	-----
-----	-----						
e3a	Cluster	Cluster		up	9000	auto/100000	
healthy	false						
e3b	Cluster	Cluster		up	9000	auto/100000	
healthy	false						

Node: node2

Ignore

						Speed (Mbps)	Health
Health							
Port	IPspace	Broadcast	Domain	Link	MTU	Admin/Oper	Status
Status							
-----	-----	-----	-----	-----	-----	-----	-----
-----	-----						
e3a	Cluster	Cluster		up	9000	auto/100000	
healthy	false						
e3b	Cluster	Cluster		up	9000	auto/100000	
healthy	false						

8. The cluster ports on each node are now connected to cluster switches in the following way, from the nodes' perspective:

Show example

```
cluster1::*> network device-discovery show -protocol lldp
```

Node/ Protocol Platform	Local Port	Discovered Device (LLDP: ChassisID)	Interface	

node1	/lldp			
	e3a	c1 (6a:ad:4f:98:3b:3f)	Eth1/1	-
	e3b	sw2 (b8:ce:f6:19:1a:7e)	swp3	-
node2	/lldp			
	e3a	c1 (6a:ad:4f:98:3b:3f)	Eth1/2	-
	e3b	sw2 (b8:ce:f6:19:1b:96)	swp4	-

9. On switch sw2, verify that all node cluster ports are up:

```
net show interface
```

Show example

```
cumulus@sw2:~$ net show interface
```

State	Name	Spd	MTU	Mode	LLDP
Summary					

...					
...					
UP	swp3	100G	9216	Trunk/L2	e3b
Master: bridge(UP)					
UP	swp4	100G	9216	Trunk/L2	e3b
Master: bridge(UP)					
UP	swp15	100G	9216	BondMember	sw1 (swp15)
Master: cluster_isl(UP)					
UP	swp16	100G	9216	BondMember	sw1 (swp16)
Master: cluster_isl(UP)					

10. On switch c1, shut down the ports connected to the cluster ports of the nodes in order to fail over the cluster LIFs.


```
(c1)# configure
```

Enter configuration commands, one per line. End with CNTL/Z.

```
(c1) (Config)# interface
```

```
(c1) (config-if-range)# shutdown <interface_list>
```

```
(c1) (config-if-range)# exit
```

```
(c1) (Config)# exit
```

```
(c1)#
```

11. Move the node cluster ports from the old switch c1 to the new switch sw1, using appropriate cabling supported by NVIDIA SN2100.
12. Verify the final configuration of the cluster:

```
network port show -ipSpace Cluster
```

Each port should display up for Link and healthy for Health Status.

Show example

```
cluster1::*> network port show -ipspace Cluster
```

Node: node1

Ignore

						Speed (Mbps)	Health
Health							
Port	IPspace	Broadcast	Domain	Link	MTU	Admin/Oper	Status
Status							
-----	-----	-----	-----	-----	-----	-----	
-----	-----						
e3a	Cluster	Cluster		up	9000	auto/100000	
healthy	false						
e3b	Cluster	Cluster		up	9000	auto/100000	
healthy	false						

Node: node2

Ignore

						Speed (Mbps)	Health
Health							
Port	IPspace	Broadcast	Domain	Link	MTU	Admin/Oper	Status
Status							
-----	-----	-----	-----	-----	-----	-----	
-----	-----						
e3a	Cluster	Cluster		up	9000	auto/100000	
healthy	false						
e3b	Cluster	Cluster		up	9000	auto/100000	
healthy	false						

13. The cluster ports on each node are now connected to cluster switches in the following way, from the nodes' perspective:

Show example

```
cluster1::*> network device-discovery show -protocol lldp
```

Node/ Protocol Platform	Local Port	Discovered Device (LLDP: ChassisID)	Interface	

node1	/lldp			
	e3a	sw1 (b8:ce:f6:19:1a:7e)	swp3	-
	e3b	sw2 (b8:ce:f6:19:1b:96)	swp3	-
node2	/lldp			
	e3a	sw1 (b8:ce:f6:19:1a:7e)	swp4	-
	e3b	sw2 (b8:ce:f6:19:1b:96)	swp4	-

14. On switches sw1 and sw2, verify that all node cluster ports are up:

```
net show interface
```

Show example

```
cumulus@sw1:~$ net show interface
```

State	Name	Spd	MTU	Mode	LLDP
Summary					

...					
...					
UP	swp3	100G	9216	Trunk/L2	e3a
Master: bridge(UP)					
UP	swp4	100G	9216	Trunk/L2	e3a
Master: bridge(UP)					
UP	swp15	100G	9216	BondMember	sw2 (swp15)
Master: cluster_isl(UP)					
UP	swp16	100G	9216	BondMember	sw2 (swp16)
Master: cluster_isl(UP)					

```
cumulus@sw2:~$ net show interface
```

State	Name	Spd	MTU	Mode	LLDP
Summary					

...					
...					
UP	swp3	100G	9216	Trunk/L2	e3b
Master: bridge(UP)					
UP	swp4	100G	9216	Trunk/L2	e3b
Master: bridge(UP)					
UP	swp15	100G	9216	BondMember	sw1 (swp15)
Master: cluster_isl(UP)					
UP	swp16	100G	9216	BondMember	sw1 (swp16)
Master: cluster_isl(UP)					

15. Verify that both nodes each have one connection to each switch:

```
net show lldp
```

Show example

The following example shows the appropriate results for both switches:

```
cumulus@sw1:~$ net show lldp
```

LocalPort	Speed	Mode	RemoteHost	RemotePort
-----	-----	-----	-----	-----
swp3	100G	Trunk/L2	node1	e3a
swp4	100G	Trunk/L2	node2	e3a
swp15	100G	BondMember	sw2	swp15
swp16	100G	BondMember	sw2	swp16

```
cumulus@sw2:~$ net show lldp
```

LocalPort	Speed	Mode	RemoteHost	RemotePort
-----	-----	-----	-----	-----
swp3	100G	Trunk/L2	node1	e3b
swp4	100G	Trunk/L2	node2	e3b
swp15	100G	BondMember	sw1	swp15
swp16	100G	BondMember	sw1	swp16

Step 3: Complete the procedure

1. Enable auto-revert on the cluster LIFs:

```
cluster1::*> network interface modify -vserver Cluster -lif * -auto-revert true
```

2. Verify that all cluster network LIFs are back on their home ports:

```
network interface show
```

Show example

```
cluster1::*> network interface show -vserver Cluster
```

		Logical	Status	Network	Current
Current Is					
Vserver	Interface	Admin/Oper	Address/Mask		Node
Port	Home				

Cluster					
		node1_clus1	up/up	169.254.209.69/16	node1
e3a	true				
		node1_clus2	up/up	169.254.49.125/16	node1
e3b	true				
		node2_clus1	up/up	169.254.47.194/16	node2
e3a	true				
		node2_clus2	up/up	169.254.19.183/16	node2
e3b	true				

3. To set up log collection, run the following command for each switch. You are prompted to enter the switch name, username, and password for log collection.

```
system switch ethernet log setup-password
```

Show example

```
cluster1::*> system switch ethernet log setup-password
Enter the switch name: <return>
The switch name entered is not recognized.
Choose from the following list:
sw1
sw2

cluster1::*> system switch ethernet log setup-password

Enter the switch name: sw1
Would you like to specify a user other than admin for log
collection? {y|n}: n

Enter the password: <enter switch password>
Enter the password again: <enter switch password>

cluster1::*> system switch ethernet log setup-password

Enter the switch name: sw2
Would you like to specify a user other than admin for log
collection? {y|n}: n

Enter the password: <enter switch password>
Enter the password again: <enter switch password>
```

4. To start log collection, run the following command, replacing **DEVICE** with the switch used in the previous command. This starts both types of log collection: the detailed **Support** logs and an hourly collection of **Periodic** data.

```
system switch ethernet log modify -device <switch-name> -log-request true
```

Show example

```
cluster1::*> system switch ethernet log modify -device sw1 -log  
-request true
```

Do you want to modify the cluster switch log collection
configuration? {y|n}: [n] **y**

Enabling cluster switch log collection.

```
cluster1::*> system switch ethernet log modify -device sw2 -log  
-request true
```

Do you want to modify the cluster switch log collection
configuration? {y|n}: [n] **y**

Enabling cluster switch log collection.

Wait for 10 minutes and then check that the log collection completes:

```
system switch ethernet log show
```

Show example

```
cluster1::*> system switch ethernet log show  
Log Collection Enabled: true
```

Index	Switch	Log Timestamp	Status
-----	-----	-----	-----
1	sw1 (b8:ce:f6:19:1b:42)	4/29/2022 03:05:25	complete
2	sw2 (b8:ce:f6:19:1b:96)	4/29/2022 03:07:42	complete



If any of these commands return an error or if the log collection does not complete, contact NetApp support.

5. Change the privilege level back to admin:

```
set -privilege admin
```

6. If you suppressed automatic case creation, re-enable it by invoking an AutoSupport message:

```
system node autosupport invoke -node * -type all -message MAINT=END
```


Migrate to a two-node switched cluster with NVIDIA SN2100 cluster switches

If you have an existing two-node switchless cluster environment, you can migrate to a two-node switched cluster environment using NVIDIA SN2100 switches to enable you to scale beyond two nodes in the cluster.

The procedure you use depends on whether you have two dedicated cluster-network ports on each controller or a single cluster port on each controller. The process documented works for all nodes using optical or Twinax ports but is not supported on this switch if nodes are using onboard 10GBASE-T RJ45 ports for the cluster-network ports.

Review requirements

Two-node switchless configuration

Ensure that:

- The two-node switchless configuration are properly set up and functioning.
- The nodes are running ONTAP 9.10.1P3 and later.
- All cluster ports are in the **up** state.
- All cluster logical interfaces (LIFs) are in the **up** state and on their home ports.

NVIDIA SN2100 cluster switch configuration

Ensure that:

- Both switches have management network connectivity.
- There is console access to the cluster switches.
- NVIDIA SN2100 node-to-node switch and switch-to-switch connections use Twinax or fiber cables.



See [Review cabling and configuration considerations](#) for caveats and further details. The [Hardware Universe - Switches](#) also contains more information about cabling.

- Inter-Switch Link (ISL) cables are connected to ports swp15 and swp16 on both NVIDIA SN2100 switches.
- Initial customization of both the SN2100 switches are completed, so that:
 - SN2100 switches are running the latest version of Cumulus Linux
 - Reference Configuration Files (RCFs) are applied to the switches
 - Any site customization, such as SMTP, SNMP, and SSH are configured on the new switches.

The [Hardware Universe](#) contains the latest information about the actual cluster ports for your platforms.

Migrate the switches

About the examples

The examples in this procedure use the following cluster switch and node nomenclature:

- The names of the SN2100 switches are *sw1* and *sw2*.
- The names of the cluster SVMs are *node1* and *node2*.
- The names of the LIFs are *node1_clus1* and *node1_clus2* on node 1, and *node2_clus1* and *node2_clus2*

on node 2 respectively.

- The `cluster1: :*>` prompt indicates the name of the cluster.
- The cluster ports used in this procedure are *e3a* and *e3b*.
- Breakout ports take the format: `swp[port]s[breakout port 0-3]`. For example, four breakout ports on `swp1` are *swp1s0*, *swp1s1*, *swp1s2*, and *swp1s3*.

Step 1: Prepare for migration

1. If AutoSupport is enabled on this cluster, suppress automatic case creation by invoking an AutoSupport message: `system node autosupport invoke -node * -type all -message MAINT=xh`

where *x* is the duration of the maintenance window in hours.

2. Change the privilege level to advanced, entering *y* when prompted to continue: `set -privilege advanced`

The advanced prompt (**>*) appears.

Step 2: Configure ports and cabling

Cumulus Linux 4.4.x

1. Disable all node-facing ports (not ISL ports) on both the new cluster switches sw1 and sw2.

You must not disable the ISL ports.

The following commands disable the node-facing ports on switches sw1 and sw2:

```
cumulus@sw1:~$ net add interface swp1s0-3, swp2s0-3, swp3-14 link
down
cumulus@sw1:~$ net pending
cumulus@sw1:~$ net commit

cumulus@sw2:~$ net add interface swp1s0-3, swp2s0-3, swp3-14 link
down
cumulus@sw2:~$ net pending
cumulus@sw2:~$ net commit
```

2. Verify that the ISL and the physical ports on the ISL between the two SN2100 switches sw1 and sw2 are up on ports swp15 and swp16:

```
net show interface
```

The following commands show that the ISL ports are up on switches sw1 and sw2:

```
cumulus@sw1:~$ net show interface
```

State	Name	Spd	MTU	Mode	LLDP	Summary
UP	swp15	100G	9216	BondMember	sw2 (swp15)	Master: cluster_isl (UP)
UP	swp16	100G	9216	BondMember	sw2 (swp16)	Master: cluster_isl (UP)

```
cumulus@sw2:~$ net show interface
```

State	Name	Spd	MTU	Mode	LLDP	Summary
UP	swp15	100G	9216	BondMember	sw1 (swp15)	Master: cluster_isl (UP)
UP	swp16	100G	9216	BondMember	sw1 (swp16)	Master: cluster_isl (UP)

Cumulus Linux 5.x

1. Disable all node-facing ports (not ISL ports) on both new cluster switches sw1 and sw2.

You must not disable the ISL ports.

The following commands disable the node-facing ports on switches sw1 and sw2:

```
cumulus@sw1:~$ nv set interface swp1s0-3,swp2s0-3,swp3-14 link state  
down  
cumulus@sw1:~$ nv config apply  
cumulus@sw1:~$ nv save  
  
cumulus@sw2:~$ nv set interface swp1s0-3,swp2s0-3,swp3-14 link state  
down  
cumulus@sw2:~$ nv config apply  
cumulus@sw2:~$ nv save
```

2. Verify that the ISL and the physical ports on the ISL between the two SN2100 switches sw1 and sw2 are up on ports swp15 and swp16:

```
nv show interface
```

The following examples show that the ISL ports are up on switches sw1 and sw2:

```
cumulus@sw1:~$ nv show interface
```

Interface	MTU	Speed	State	Remote Host	Remote Port
Type	Summary				

...					
...					
+ swp14	9216		down		
swp					
+ swp15	9216	100G	up	oss-g-rcf1	Intra-Cluster Switch
ISL Port swp15 swp					
+ swp16	9216	100G	up	oss-g-rcf2	Intra-Cluster Switch
ISL Port swp16 swp					

```
cumulus@sw2:~$ nv show interface
```

Interface	MTU	Speed	State	Remote Host	Remote Port
Type	Summary				

...					
...					
+ swp14	9216		down		
swp					
+ swp15	9216	100G	up	oss-g-rcf1	Intra-Cluster Switch
ISL Port swp15 swp					
+ swp16	9216	100G	up	oss-g-rcf2	Intra-Cluster Switch
ISL Port swp16 swp					

3. Verify that all cluster ports are up:

```
network port show
```

Each port should display up for Link and healthy for Health Status.

Show example

```
cluster1::*> network port show
```

```
Node: node1
```

```
Ignore
```

						Speed (Mbps)
Health	Health					
Port	IPspace	Broadcast	Domain	Link	MTU	Admin/Oper
Status	Status					
-----	-----	-----	-----	-----	-----	-----
e3a	Cluster	Cluster		up	9000	auto/100000
healthy	false					
e3b	Cluster	Cluster		up	9000	auto/100000
healthy	false					

```
Node: node2
```

```
Ignore
```

						Speed (Mbps)
Health	Health					
Port	IPspace	Broadcast	Domain	Link	MTU	Admin/Oper
Status	Status					
-----	-----	-----	-----	-----	-----	-----
e3a	Cluster	Cluster		up	9000	auto/100000
healthy	false					
e3b	Cluster	Cluster		up	9000	auto/100000
healthy	false					

4. Verify that all cluster LIFs are up and operational:

```
network interface show
```

Each cluster LIF should display true for Is Home and have a Status Admin/Oper of up/up.

Show example

```
cluster1::*> network interface show -vserver Cluster
```

	Logical	Status	Network	Current
Current Is				
Vserver	Interface	Admin/Oper	Address/Mask	Node
Port	Home			

Cluster				
	node1_clus1	up/up	169.254.209.69/16	node1
e3a	true			
	node1_clus2	up/up	169.254.49.125/16	node1
e3b	true			
	node2_clus1	up/up	169.254.47.194/16	node2
e3a	true			
	node2_clus2	up/up	169.254.19.183/16	node2
e3b	true			

5. Disable auto-revert on the cluster LIFs:

```
network interface modify -vserver Cluster -lif * -auto-revert false
```

Show example

```
cluster1::*> network interface modify -vserver Cluster -lif * -auto-revert false
```

	Logical	
Vserver	Interface	Auto-revert

Cluster		
	node1_clus1	false
	node1_clus2	false
	node2_clus1	false
	node2_clus2	false

6. Disconnect the cable from cluster port e3a on node1, and then connect e3a to port 3 on cluster switch sw1, using the appropriate cabling supported by the SN2100 switches.

The [Hardware Universe - Switches](#) contains more information about cabling.

7. Disconnect the cable from cluster port e3a on node2, and then connect e3a to port 4 on cluster switch sw1,

using the appropriate cabling supported by the SN2100 switches.

Cumulus Linux 4.4.x

8. On switch sw1, enable all node-facing ports.

The following commands enable all node-facing ports on switch sw1.

```
cumulus@sw1:~$ net del interface swp1s0-3, swp2s0-3, swp3-14 link  
down  
cumulus@sw1:~$ net pending  
cumulus@sw1:~$ net commit
```

9. On switch sw1, verify that all ports are up:

```
net show interface all
```

```
cumulus@sw1:~$ net show interface all
```

State	Name	Spd	MTU	Mode	LLDP	Summary
-----	-----	-----	-----	-----	-----	-----
...						
DN	swp1s0	10G	9216	Trunk/L2		Master: br_default(UP)
DN	swp1s1	10G	9216	Trunk/L2		Master: br_default(UP)
DN	swp1s2	10G	9216	Trunk/L2		Master: br_default(UP)
DN	swp1s3	10G	9216	Trunk/L2		Master: br_default(UP)
DN	swp2s0	25G	9216	Trunk/L2		Master: br_default(UP)
DN	swp2s1	25G	9216	Trunk/L2		Master: br_default(UP)
DN	swp2s2	25G	9216	Trunk/L2		Master: br_default(UP)
DN	swp2s3	25G	9216	Trunk/L2		Master: br_default(UP)
UP	swp3	100G	9216	Trunk/L2	node1 (e3a)	Master: br_default(UP)
UP	swp4	100G	9216	Trunk/L2	node2 (e3a)	Master: br_default(UP)
...						
...						
UP	swp15	100G	9216	BondMember	swp15	Master: cluster_isl(UP)
UP	swp16	100G	9216	BondMember	swp16	Master: cluster_isl(UP)
...						

Cumulus Linux 5.x

8. On switch sw1, enable all node-facing ports.

The following commands enable all node-facing ports on switch sw1.

```
cumulus@sw1:~$ nv unset interface swp1s0-3,swp2s0-3,swp3-14 link  
state down  
cumulus@sw1:~$ nv config apply  
cumulus@sw1:~$ nv config save
```

9. On switch sw1, verify that all ports are up:

```
nv show interface
```

```
cumulus@sw1:~$ nv show interface
```

Interface	State	Speed	MTU	Type	Remote Host
Remote Port	Summary				
-----	-----	-----	-----	-----	
-----	-----	-----	-----	-----	
...					
...					
swp1s0	up	10G	9216	swp	odq-a300-1a
e0a					
swp1s1	up	10G	9216	swp	odq-a300-1b
e0a					
swp1s2	down	10G	9216	swp	
swp1s3	down	10G	9216	swp	
swp2s0	down	25G	9216	swp	
swp2s1	down	25G	9216	swp	
swp2s2	down	25G	9216	swp	
swp2s3	down	25G	9216	swp	
swp3	down		9216	swp	
swp4	down		9216	swp	
...					
...					
swp14	down		9216	swp	
swp15	up	100G	9216	swp	oss-g-int-rcf10
swp15					
swp16	up	100G	9216	swp	oss-g-int-rcf10
swp16					

10. Verify that all cluster ports are up:

```
network port show -ip space Cluster
```

Show example

The following example shows that all of the cluster ports are up on node1 and node2:

```
cluster1::*> network port show -ipspace Cluster
```

Node: node1

Ignore

Health	Health				Speed (Mbps)	
Port	IPspace	Broadcast	Domain	Link	MTU	Admin/Oper
Status	Status					
-----	-----	-----	-----	-----	-----	-----
e3a	Cluster	Cluster		up	9000	auto/100000
healthy	false					
e3b	Cluster	Cluster		up	9000	auto/100000
healthy	false					

Node: node2

Ignore

Health	Health				Speed (Mbps)	
Port	IPspace	Broadcast	Domain	Link	MTU	Admin/Oper
Status	Status					
-----	-----	-----	-----	-----	-----	-----
e3a	Cluster	Cluster		up	9000	auto/100000
healthy	false					
e3b	Cluster	Cluster		up	9000	auto/100000
healthy	false					

11. Display information about the status of the nodes in the cluster:

```
cluster show
```

Show example

The following example displays information about the health and eligibility of the nodes in the cluster:

```
cluster1::*> cluster show
```

Node	Health	Eligibility	Epsilon
-----	-----	-----	-----
node1	true	true	false
node2	true	true	false

12. Disconnect the cable from cluster port e3b on node1, and then connect e3b to port 3 on cluster switch sw2, using the appropriate cabling supported by the SN2100 switches.
13. Disconnect the cable from cluster port e3b on node2, and then connect e3b to port 4 on cluster switch sw2, using the appropriate cabling supported by the SN2100 switches.

Cumulus Linux 4.4.x

14. On switch sw2, enable all node-facing ports.

The following commands enable the node-facing ports on switch sw2:

```
cumulus@sw2:~$ net del interface swp1s0-3, swp2s0-3, swp3-14 link  
down  
cumulus@sw2:~$ net pending  
cumulus@sw2:~$ net commit
```

15. On switch sw2, verify that all ports are up:

```
net show interface all
```

```
cumulus@sw2:~$ net show interface all
```

State	Name	Spd	MTU	Mode	LLDP	Summary
...						
DN	swp1s0	10G	9216	Trunk/L2		Master:
br_default(UP)						
DN	swp1s1	10G	9216	Trunk/L2		Master:
br_default(UP)						
DN	swp1s2	10G	9216	Trunk/L2		Master:
br_default(UP)						
DN	swp1s3	10G	9216	Trunk/L2		Master:
br_default(UP)						
DN	swp2s0	25G	9216	Trunk/L2		Master:
br_default(UP)						
DN	swp2s1	25G	9216	Trunk/L2		Master:
br_default(UP)						
DN	swp2s2	25G	9216	Trunk/L2		Master:
br_default(UP)						
DN	swp2s3	25G	9216	Trunk/L2		Master:
br_default(UP)						
UP	swp3	100G	9216	Trunk/L2	node1 (e3b)	Master:
br_default(UP)						
UP	swp4	100G	9216	Trunk/L2	node2 (e3b)	Master:
br_default(UP)						
...						
...						
UP	swp15	100G	9216	BondMember	swp15	Master:
cluster_isl(UP)						
UP	swp16	100G	9216	BondMember	swp16	Master:
cluster_isl(UP)						
...						

16. On both switches sw1 and sw2, verify that both nodes each have one connection to each switch:

```
net show lldp
```

The following example shows the appropriate results for both switches sw1 and sw2:

```
cumulus@sw1:~$ net show lldp
```

LocalPort	Speed	Mode	RemoteHost	RemotePort
-----	-----	-----	-----	-----
swp3	100G	Trunk/L2	node1	e3a
swp4	100G	Trunk/L2	node2	e3a
swp15	100G	BondMember	sw2	swp15
swp16	100G	BondMember	sw2	swp16

```
cumulus@sw2:~$ net show lldp
```

LocalPort	Speed	Mode	RemoteHost	RemotePort
-----	-----	-----	-----	-----
swp3	100G	Trunk/L2	node1	e3b
swp4	100G	Trunk/L2	node2	e3b
swp15	100G	BondMember	sw1	swp15
swp16	100G	BondMember	sw1	swp16

Cumulus Linux 5.x

14. On switch sw2, enable all node-facing ports.

The following commands enable the node-facing ports on switch sw2:

```
cumulus@sw2:~$ nv unset interface swp1s0-3,swp2s0-3,swp3-14 link  
state down  
cumulus@sw2:~$ nv config apply  
cumulus@sw2:~$ nv config save
```

15. On switch sw2, verify that all ports are up:

```
nv show interface
```



```
cumulus@sw2:~$ nv show interface
```

Interface	State	Speed	MTU	Type	Remote Host
Remote Port	Summary				
-----	-----	-----	-----	-----	-----
-----	-----	-----	-----	-----	-----
...					
...					
swp1s0	up	10G	9216	swp	odq-a300-1a
e0a					
swp1s1	up	10G	9216	swp	odq-a300-1b
e0a					
swp1s2	down	10G	9216	swp	
swp1s3	down	10G	9216	swp	
swp2s0	down	25G	9216	swp	
swp2s1	down	25G	9216	swp	
swp2s2	down	25G	9216	swp	
swp2s3	down	25G	9216	swp	
swp3	down		9216	swp	
swp4	down		9216	swp	
...					
...					
swp14	down		9216	swp	
swp15	up	100G	9216	swp	ossq-int-rcf10
swp15					
swp16	up	100G	9216	swp	ossq-int-rcf10
swp16					

16. On both switches sw1 and sw2, verify that both nodes each have one connection to each switch:

```
nv show interface --view=lldp
```

The following examples show the appropriate results for both switches sw1 and sw2:

```
cumulus@sw1:~$ nv show interface --view=lldp
```

Interface	Speed	Type	Remote Host
Remote Port			
-----	-----	-----	-----
-----	-----	-----	-----
...			
...			
swp1s0	10G	swp	odq-a300-1a
e0a			
swp1s1	10G	swp	odq-a300-1b

```

e0a
swp1s2      10G    swp
swp1s3      10G    swp
swp2s0      25G    swp
swp2s1      25G    swp
swp2s2      25G    swp
swp2s3      25G    swp
swp3                swp
swp4                swp
...
...
swp14                swp
swp15      100G    swp      ossg-int-rcf10
swp15
swp16      100G    swp      ossg-int-rcf10
swp16

```

```
cumulus@sw2:~$ nv show interface --view=lldp
```

Interface	Speed	Type	Remote Host
Remote Port			
-----	-----	-----	-----

...			
...			
swp1s0	10G	swp	odq-a300-1a
e0a			
swp1s1	10G	swp	odq-a300-1b
e0a			
swp1s2	10G	swp	
swp1s3	10G	swp	
swp2s0	25G	swp	
swp2s1	25G	swp	
swp2s2	25G	swp	
swp2s3	25G	swp	
swp3		swp	
swp4		swp	
...			
...			
swp14		swp	
swp15	100G	swp	ossg-int-rcf10
swp15			
swp16	100G	swp	ossg-int-rcf10
swp16			

17. Display information about the discovered network devices in your cluster:

```
network device-discovery show -protocol lldp
```

Show example

```
cluster1::*> network device-discovery show -protocol lldp
Node/      Local  Discovered
Protocol   Port   Device (LLDP: ChassisID)  Interface  Platform
-----
node1      /lldp
           e3a    sw1 (b8:ce:f6:19:1a:7e)    swp3       -
           e3b    sw2 (b8:ce:f6:19:1b:96)    swp3       -
node2      /lldp
           e3a    sw1 (b8:ce:f6:19:1a:7e)    swp4       -
           e3b    sw2 (b8:ce:f6:19:1b:96)    swp4       -
```

18. Verify that all cluster ports are up:

```
network port show -ipSpace Cluster
```

Show example

The following example shows that all of the cluster ports are up on node1 and node2:

```
cluster1::*> network port show -ipspace Cluster
```

Node: node1

Ignore

						Speed(Mbps)	Health
Health							
Port	IPspace	Broadcast	Domain	Link	MTU	Admin/Oper	Status
Status							
-----	-----	-----	-----	-----	-----	-----	-----
e3a	Cluster	Cluster		up	9000	auto/10000	
healthy	false						
e3b	Cluster	Cluster		up	9000	auto/10000	
healthy	false						

Node: node2

Ignore

						Speed(Mbps)	Health
Health							
Port	IPspace	Broadcast	Domain	Link	MTU	Admin/Oper	Status
Status							
-----	-----	-----	-----	-----	-----	-----	-----
e3a	Cluster	Cluster		up	9000	auto/10000	
healthy	false						
e3b	Cluster	Cluster		up	9000	auto/10000	
healthy	false						

Step 3: Complete the procedure

1. Enable auto-revert on all cluster LIFs:

```
net interface modify -vserver Cluster -lif * -auto-revert true
```

Show example

```
cluster1::*> net interface modify -vserver Cluster -lif * -auto  
-revert true
```

Vserver	Logical Interface	Auto-revert
Cluster	node1_clus1	true
	node1_clus2	true
	node2_clus1	true
	node2_clus2	true

2. Verify that all interfaces display true for Is Home:

```
net interface show -vserver Cluster
```



This might take a minute to complete.

Show example

The following example shows that all LIFs are up on node1 and node2 and that Is Home results are true:

```
cluster1::*> net interface show -vserver Cluster
```

	Logical	Status	Network	Current	
Current Is					
Vserver	Interface	Admin/Oper	Address/Mask	Node	Port
Home					
Cluster					
	node1_clus1	up/up	169.254.209.69/16	node1	e3a
true					
	node1_clus2	up/up	169.254.49.125/16	node1	e3b
true					
	node2_clus1	up/up	169.254.47.194/16	node2	e3a
true					
	node2_clus2	up/up	169.254.19.183/16	node2	e3b
true					

3. Verify that the settings are disabled:

```
network options switchless-cluster show
```

Show example

The false output in the following example shows that the configuration settings are disabled:

```
cluster1::*> network options switchless-cluster show  
Enable Switchless Cluster: false
```

4. Verify the status of the node members in the cluster:

```
cluster show
```

Show example

The following example shows information about the health and eligibility of the nodes in the cluster:

```
cluster1::*> cluster show
```

Node	Health	Eligibility	Epsilon
node1	true	true	false
node2	true	true	false

5. Verify that the cluster network has full connectivity:

```
cluster ping-cluster -node node-name
```

Show example

```
cluster1::*> cluster ping-cluster -node node1
Host is node1
Getting addresses from network interface table...
Cluster node1_clus1 169.254.209.69 node1 e3a
Cluster node1_clus2 169.254.49.125 node1 e3b
Cluster node2_clus1 169.254.47.194 node2 e3a
Cluster node2_clus2 169.254.19.183 node2 e3b
Local = 169.254.47.194 169.254.19.183
Remote = 169.254.209.69 169.254.49.125
Cluster Vserver Id = 4294967293
Ping status:

Basic connectivity succeeds on 4 path(s)
Basic connectivity fails on 0 path(s)

Detected 9000 byte MTU on 4 path(s):
Local 169.254.47.194 to Remote 169.254.209.69
Local 169.254.47.194 to Remote 169.254.49.125
Local 169.254.19.183 to Remote 169.254.209.69
Local 169.254.19.183 to Remote 169.254.49.125
Larger than PMTU communication succeeds on 4 path(s)
RPC status:
2 paths up, 0 paths down (tcp check)
2 paths up, 0 paths down (udp check)
```

6. To set up log collection, run the following command for each switch. You are prompted to enter the switch name, username, and password for log collection.

```
system switch ethernet log setup-password
```

Show example

```
cluster1::*> system switch ethernet log setup-password
Enter the switch name: <return>
The switch name entered is not recognized.
Choose from the following list:
cs1
cs2

cluster1::*> system switch ethernet log setup-password

Enter the switch name: cs1
Would you like to specify a user other than admin for log
collection? {y|n}: n

Enter the password: <enter switch password>
Enter the password again: <enter switch password>

cluster1::*> system switch ethernet log setup-password

Enter the switch name: cs2
Would you like to specify a user other than admin for log
collection? {y|n}: n

Enter the password: <enter switch password>
Enter the password again: <enter switch password>
```

7. To start log collection, run the following command, replacing **DEVICE** with the switch used in the previous command. This starts both types of log collection: the detailed **Support** logs and an hourly collection of **Periodic** data.

```
system switch ethernet log modify -device <switch-name> -log-request true
```


Show example

```
cluster1::*> system switch ethernet log modify -device sw1 -log  
-request true
```

Do you want to modify the cluster switch log collection
configuration? {y|n}: [n] **y**

Enabling cluster switch log collection.

```
cluster1::*> system switch ethernet log modify -device sw2 -log  
-request true
```

Do you want to modify the cluster switch log collection
configuration? {y|n}: [n] **y**

Enabling cluster switch log collection.

Wait for 10 minutes and then check that the log collection completes:

```
system switch ethernet log show
```

Show example

```
cluster1::*> system switch ethernet log show  
Log Collection Enabled: true
```

Index	Switch	Log Timestamp	Status
1	sw1 (b8:ce:f6:19:1b:42)	4/29/2022 03:05:25	complete
2	sw2 (b8:ce:f6:19:1b:96)	4/29/2022 03:07:42	complete



If any of these commands return an error, contact NetApp support.

8. Change the privilege level back to admin:

```
set -privilege admin
```

9. If you suppressed automatic case creation, reenable it by invoking an AutoSupport message:

```
system node autosupport invoke -node * -type all -message MAINT=END
```

Replace switches

Replace a NVIDIA SN2100 cluster switch

Follow this procedure to replace a defective NVIDIA SN2100 switch in a cluster network. This is a nondisruptive procedure (NDU).

Review requirements

Existing cluster and network infrastructure

Ensure that:

- The existing cluster are verified as completely functional, with at least one fully connected cluster switch.
- All cluster ports are up.
- All cluster logical interfaces (LIFs) are up and on their home ports.
- The ONTAP `cluster ping-cluster -node node1` command indicates that basic connectivity and larger than PMTU communication are successful on all paths.

NVIDIA SN2100 replacement switch

Ensure that:

- Management network connectivity on the replacement switch are functional.
- Console access to the replacement switch are in place.
- The node connections are ports swp1 through swp14.
- All Inter-Switch Link (ISL) ports are disabled on ports swp15 and swp16.
- The desired reference configuration file (RCF) and Cumulus operating system image switch are loaded onto the switch.
- Initial customization of the switch is complete.

Also make sure that any previous site customizations, such as STP, SNMP, and SSH, are copied to the new switch.



You must execute the command for migrating a cluster LIF from the node where the cluster LIF is hosted.

Replace the switch

About the examples

The examples in this procedure use the following switch and node nomenclature:

- The names of the existing NVIDIA SN2100 switches are *sw1* and *sw2*.
- The name of the new NVIDIA SN2100 switch is *nsw2*.
- The node names are *node1* and *node2*.
- The cluster ports on each node are named *e3a* and *e3b*.
- The cluster LIF names are *node1_clus1* and *node1_clus2* for node1, and *node2_clus1* and *node2_clus2* for node2.

- The prompt for changes to all cluster nodes is `cluster1:*>`
- Breakout ports take the format: `swp[port]s[breakout port 0-3]`. For example, four breakout ports on `swp1` are *swp1s0*, *swp1s1*, *swp1s2*, and *swp1s3*.

About the cluster network topology

This procedure is based on the following cluster network topology:

Show example topology

```
cluster1::*> network port show -ipspace Cluster
```

Node: node1

Ignore

						Speed(Mbps)	Health
Health							
Port	IPspace	Broadcast	Domain	Link	MTU	Admin/Oper	Status
Status							
-----	-----	-----	----	----	----	-----	-----

e3a	Cluster	Cluster		up	9000	auto/100000	healthy
false							
e3b	Cluster	Cluster		up	9000	auto/100000	healthy
false							

Node: node2

Ignore

						Speed(Mbps)	Health
Health							
Port	IPspace	Broadcast	Domain	Link	MTU	Admin/Oper	Status
Status							
-----	-----	-----	----	----	----	-----	-----

e3a	Cluster	Cluster		up	9000	auto/100000	healthy
false							
e3b	Cluster	Cluster		up	9000	auto/100000	healthy
false							

```
cluster1::*> network interface show -vserver Cluster
```

	Logical	Status	Network		Current
Current Is					
Vserver	Interface	Admin/Oper	Address/Mask	Node	Port
Home					
-----	-----	-----	-----	-----	-----

Cluster					
	node1_clus1	up/up	169.254.209.69/16	node1	e3a
true					
	node1_clus2	up/up	169.254.49.125/16	node1	e3b
true					

```

node2_clus1 up/up 169.254.47.194/16 node2 e3a
true
node2_clus2 up/up 169.254.19.183/16 node2 e3b
true

```

```
cluster1::*> network device-discovery show -protocol lldp
```

Node/	Local	Discovered			
Protocol	Port	Device (LLDP: ChassisID)	Interface	Platform	
node1	/lldp				
	e3a	sw1 (b8:ce:f6:19:1a:7e)	swp3	-	
	e3b	sw2 (b8:ce:f6:19:1b:96)	swp3	-	
node2	/lldp				
	e3a	sw1 (b8:ce:f6:19:1a:7e)	swp4	-	
	e3b	sw2 (b8:ce:f6:19:1b:96)	swp4	-	

+

```
cumulus@sw1:~$ net show lldp
```

LocalPort	Speed	Mode	RemoteHost	RemotePort
swp3	100G	Trunk/L2	sw2	e3a
swp4	100G	Trunk/L2	sw2	e3a
swp15	100G	BondMember	sw2	swp15
swp16	100G	BondMember	sw2	swp16

```
cumulus@sw2:~$ net show lldp
```

LocalPort	Speed	Mode	RemoteHost	RemotePort
swp3	100G	Trunk/L2	sw1	e3b
swp4	100G	Trunk/L2	sw1	e3b
swp15	100G	BondMember	sw1	swp15
swp16	100G	BondMember	sw1	swp16

Step 1: Prepare for replacement

1. If AutoSupport is enabled on this cluster, suppress automatic case creation by invoking an AutoSupport message:

```
system node autosupport invoke -node * -type all -message MAINT=xh
```

where x is the duration of the maintenance window in hours.

2. Change the privilege level to advanced, entering **y** when prompted to continue:

```
set -privilege advanced
```

The advanced prompt (***>**) appears.

3. Install the appropriate RCF and image on the switch, nsw2, and make any necessary site preparations.

If necessary, verify, download, and install the appropriate versions of the RCF and Cumulus software for the new switch.

- a. You can download the applicable Cumulus software for your cluster switches from the *NVIDIA Support* site. Follow the steps on the Download page to download the Cumulus Linux for the version of ONTAP software you are installing.
- b. The appropriate RCF is available from the [NVIDIA Cluster and Storage Switches](#) page. Follow the steps on the Download page to download the correct RCF for the version of ONTAP software you are installing.

Step 2: Configure ports and cabling

1. On the new switch nsw2, log in as admin and shut down all of the ports that will be connected to the node cluster interfaces (ports swp1 to swp14).

The LIFs on the cluster nodes should have already failed over to the other cluster port for each node.

Show example

```
cumulus@nsw2:~$ net add interface swp1s0-3, swp2s0-3, swp3-14 link
down
cumulus@nsw2:~$ net pending
cumulus@nsw2:~$ net commit
```

2. Disable auto-revert on the cluster LIFs:

```
network interface modify -vserver Cluster -lif * -auto-revert false
```

Show example

```
cluster1::*> network interface modify -vserver Cluster -lif * -auto
-revert false
```

```
Warning: Disabling the auto-revert feature of the cluster logical
interface may effect the availability of your cluster network. Are
you sure you want to continue? {y|n}: y
```

3. Verify that all cluster LIFs have auto-revert enabled:

```
net interface show -vserver Cluster -fields auto-revert
```

4. Shut down the ISL ports swp15 and swp16 on the SN2100 switch sw1.

Show example

```
cumulus@sw1:~$ net add interface swp15-16 link down
cumulus@sw1:~$ net pending
cumulus@sw1:~$ net commit
```

5. Remove all the cables from the SN2100 sw1 switch, and then connect them to the same ports on the SN2100 nsw2 switch.
6. Bring up the ISL ports swp15 and swp16 between the sw1 and nsw2 switches.

Show example

The following commands enable ISL ports swp15 and swp16 on switch sw1:

```
cumulus@sw1:~$ net del interface swp15-16 link down
cumulus@sw1:~$ net pending
cumulus@sw1:~$ net commit
```

The following example shows that the ISL ports are up on switch sw1:

```
cumulus@sw1:~$ net show interface
```

State	Name	Spd	MTU	Mode	LLDP	Summary
UP	swp15	100G	9216	BondMember	nsw2 (swp15)	Master: cluster_isl (UP)
UP	swp16	100G	9216	BondMember	nsw2 (swp16)	Master: cluster_isl (UP)

+ The following example shows that the ISL ports are up on switch nsw2:

+

```
cumulus@nsw2:~$ net show interface
```

State	Name	Spd	MTU	Mode	LLDP	Summary
UP	swp15	100G	9216	BondMember	sw1 (swp15)	Master: cluster_isl (UP)
UP	swp16	100G	9216	BondMember	sw1 (swp16)	Master: cluster_isl (UP)

7. Verify that port e3b is up on all nodes:

```
network port show -ipSpace Cluster
```


Show example

The output should be similar to the following:

```
cluster1::*> network port show -ipspace Cluster
```

Node: node1

Ignore

						Speed (Mbps)
Health	Health					
Port	IPspace	Broadcast	Domain	Link	MTU	Admin/Oper
Status	Status					
-----	-----	-----	-----	-----	-----	-----
e3a	Cluster	Cluster		up	9000	auto/100000
healthy	false					
e3b	Cluster	Cluster		up	9000	auto/100000
healthy	false					

Node: node2

Ignore

						Speed (Mbps)
Health	Health					
Port	IPspace	Broadcast	Domain	Link	MTU	Admin/Oper
Status	Status					
-----	-----	-----	-----	-----	-----	-----
e3a	Cluster	Cluster		up	9000	auto/100000
healthy	false					
e3b	Cluster	Cluster		up	9000	auto/100000
healthy	false					

8. The cluster ports on each node are now connected to cluster switches in the following way, from the nodes' perspective:

Show example

```
cluster1::~*> network device-discovery show -protocol lldp
```

Node/	Local	Discovered			
Protocol	Port	Device	(LLDP: ChassisID)	Interface	Platform
-----	-----	-----	-----	-----	
node1	/lldp				
	e3a	sw1	(b8:ce:f6:19:1a:7e)	swp3	-
	e3b	nsw2	(b8:ce:f6:19:1b:b6)	swp3	-
node2	/lldp				
	e3a	sw1	(b8:ce:f6:19:1a:7e)	swp4	-
	e3b	nsw2	(b8:ce:f6:19:1b:b6)	swp4	-

9. Verify that all node cluster ports are up:

```
net show interface
```

Show example

```
cumulus@nsw2::~~$ net show interface
```

State	Name	Spd	MTU	Mode	LLDP
Summary	-----	-----	-----	-----	-----

...					
...					
UP	swp3	100G	9216	Trunk/L2	
Master: bridge(UP)					
UP	swp4	100G	9216	Trunk/L2	
Master: bridge(UP)					
UP	swp15	100G	9216	BondMember	sw1 (swp15)
Master: cluster_isl(UP)					
UP	swp16	100G	9216	BondMember	sw1 (swp16)
Master: cluster_isl(UP)					

10. Verify that both nodes each have one connection to each switch:

```
net show lldp
```

Show example

The following example shows the appropriate results for both switches:

```
cumulus@sw1:~$ net show lldp
```

LocalPort	Speed	Mode	RemoteHost	RemotePort
-----	-----	-----	-----	-----
swp3	100G	Trunk/L2	node1	e3a
swp4	100G	Trunk/L2	node2	e3a
swp15	100G	BondMember	nsw2	swp15
swp16	100G	BondMember	nsw2	swp16

```
cumulus@nsw2:~$ net show lldp
```

LocalPort	Speed	Mode	RemoteHost	RemotePort
-----	-----	-----	-----	-----
swp3	100G	Trunk/L2	node1	e3b
swp4	100G	Trunk/L2	node2	e3b
swp15	100G	BondMember	sw1	swp15
swp16	100G	BondMember	sw1	swp16

11. Enable auto-revert on the cluster LIFs:

```
cluster1::*> network interface modify -vserver Cluster -lif * -auto-revert  
true
```

12. On switch nsw2, bring up the ports connected to the network ports of the nodes.

Show example

```
cumulus@nsw2:~$ net del interface swp1-14 link down  
cumulus@nsw2:~$ net pending  
cumulus@nsw2:~$ net commit
```

13. Display information about the nodes in a cluster:

```
cluster show
```

Show example

This example shows that the node health for node1 and node2 in this cluster is true:

```
cluster1::*> cluster show
```

Node	Health	Eligibility
-----	-----	-----
node1	true	true
node2	true	true

14. Verify that all physical cluster ports are up:

```
network port show ipspace Cluster
```

Show example

```
cluster1::*> network port show -ipspace Cluster
```

Node node1

Ignore

Health	Health					Speed (Mbps)
Port	IPspace	Broadcast	Domain	Link	MTU	Admin/Oper
Status	Status					
-----	-----	-----	-----	-----	-----	-----
e3a	Cluster	Cluster		up	9000	auto/10000
healthy	false					
e3b	Cluster	Cluster		up	9000	auto/10000
healthy	false					

Node: node2

Ignore

Health	Health					Speed (Mbps)
Port	IPspace	Broadcast	Domain	Link	MTU	Admin/Oper
Status	Status					
-----	-----	-----	-----	-----	-----	-----
e3a	Cluster	Cluster		up	9000	auto/10000
healthy	false					
e3b	Cluster	Cluster		up	9000	auto/10000
healthy	false					

Step 3: Complete the procedure

1. Verify that the cluster network is healthy.

Show example

```
cumulus@sw1:~$ net show lldp
```

LocalPort	Speed	Mode	RemoteHost	RemotePort
-----	-----	-----	-----	-----
swp3	100G	Trunk/L2	node1	e3a
swp4	100G	Trunk/L2	node2	e3a
swp15	100G	BondMember	nsw2	swp15
swp16	100G	BondMember	nsw2	swp16

2. Create a password for the Ethernet switch health monitor log collection feature:

```
system switch ethernet log setup-password
```

Show example

```
cluster1::*> system switch ethernet log setup-password
```

```
Enter the switch name: <return>
```

```
The switch name entered is not recognized.
```

```
Choose from the following list:
```

```
cs1
```

```
cs2
```

```
cluster1::*> system switch ethernet log setup-password
```

```
Enter the switch name: cs1
```

```
Would you like to specify a user other than admin for log  
collection? {y|n}: n
```

```
Enter the password: <enter switch password>
```

```
Enter the password again: <enter switch password>
```

```
cluster1::*> system switch ethernet log setup-password
```

```
Enter the switch name: cs2
```

```
Would you like to specify a user other than admin for log  
collection? {y|n}: n
```

```
Enter the password: <enter switch password>
```

```
Enter the password again: <enter switch password>
```

3. Enable the Ethernet switch health monitor log collection feature.

```
system switch ethernet log modify -device <switch-name> -log-request true
```

Show example

```
cluster1::*> system switch ethernet log modify -device cs1 -log  
-request true
```

```
Do you want to modify the cluster switch log collection  
configuration? {y|n}: [n] y
```

```
Enabling cluster switch log collection.
```

```
cluster1::*> system switch ethernet log modify -device cs2 -log  
-request true
```

```
Do you want to modify the cluster switch log collection  
configuration? {y|n}: [n] y
```

```
Enabling cluster switch log collection.
```

Wait for 10 minutes and then check that the log collection completes:

```
system switch ethernet log show
```

Show example

```
cluster1::*> system switch ethernet log show  
Log Collection Enabled: true
```

Index	Switch	Log Timestamp	Status
-----	-----	-----	-----
1	cs1 (b8:ce:f6:19:1b:42)	4/29/2022 03:05:25	complete
2	cs2 (b8:ce:f6:19:1b:96)	4/29/2022 03:07:42	complete



If any of these commands return an error or if the log collection does not complete, contact NetApp support.

4. Change the privilege level back to admin:

```
set -privilege admin
```

5. If you suppressed automatic case creation, re-enable it by invoking an AutoSupport message:

```
system node autosupport invoke -node * -type all -message MAINT=END
```

Replace NVIDIA SN2100 cluster switches with switchless connections

You can migrate from a cluster with a switched cluster network to one where two nodes are directly connected for ONTAP 9.3 and later.

Review requirements

Guidelines

Review the following guidelines:

- Migrating to a two-node switchless cluster configuration is a nondisruptive operation. Most systems have two dedicated cluster interconnect ports on each node, but you can also use this procedure for systems with a larger number of dedicated cluster interconnect ports on each node, such as four, six or eight.
- You cannot use the switchless cluster interconnect feature with more than two nodes.
- If you have an existing two-node cluster that uses cluster interconnect switches and is running ONTAP 9.3 or later, you can replace the switches with direct, back-to-back connections between the nodes.


What you'll need

- A healthy cluster that consists of two nodes connected by cluster switches. The nodes must be running the same ONTAP release.
- Each node with the required number of dedicated cluster ports, which provide redundant cluster interconnect connections to support your system configuration. For example, there are two redundant ports for a system with two dedicated cluster interconnect ports on each node.

Migrate the switches

About this task

The following procedure removes the cluster switches in a two-node cluster and replaces each connection to the switch with a direct connection to the partner node.



About the examples

The examples in the following procedure show nodes that are using "e0a" and "e0b" as cluster ports. Your

nodes might be using different cluster ports as they vary by system.

Step 1: Prepare for migration

1. Change the privilege level to advanced, entering `y` when prompted to continue:

```
set -privilege advanced
```

The advanced prompt `*>` appears.

2. ONTAP 9.3 and later supports automatic detection of switchless clusters, which is enabled by default.

You can verify that detection of switchless clusters is enabled by running the advanced privilege command:

```
network options detect-switchless-cluster show
```

Show example

The following example output shows if the option is enabled.

```
cluster::*> network options detect-switchless-cluster show
(network options detect-switchless-cluster show)
Enable Switchless Cluster Detection: true
```

If "Enable Switchless Cluster Detection" is `false`, contact NetApp support.

3. If AutoSupport is enabled on this cluster, suppress automatic case creation by invoking an AutoSupport message:

```
system node autosupport invoke -node * -type all -message
MAINT=<number_of_hours>h
```

where `h` is the duration of the maintenance window in hours. The message notifies technical support of this maintenance task so that they can suppress automatic case creation during the maintenance window.

In the following example, the command suppresses automatic case creation for two hours:

Show example

```
cluster::*> system node autosupport invoke -node * -type all
-message MAINT=2h
```


Step 2: Configure ports and cabling

1. Organize the cluster ports on each switch into groups so that the cluster ports in group1 go to cluster switch1 and the cluster ports in group2 go to cluster switch2. These groups are required later in the procedure.

2. Identify the cluster ports and verify link status and health:

```
network port show -ipspace Cluster
```

In the following example for nodes with cluster ports "e0a" and "e0b", one group is identified as "node1:e0a" and "node2:e0a" and the other group as "node1:e0b" and "node2:e0b". Your nodes might be using different cluster ports because they vary by system.



Verify that the ports have a value of `up` for the "Link" column and a value of `healthy` for the "Health Status" column.

Show example

```
cluster::> network port show -ipspace Cluster
Node: node1

Ignore
Speed (Mbps) Health
Health
Port IPspace Broadcast Domain Link MTU Admin/Oper Status
Status
-----
-----
e0a Cluster Cluster up 9000 auto/10000 healthy
false
e0b Cluster Cluster up 9000 auto/10000 healthy
false

Node: node2

Ignore
Speed (Mbps) Health
Health
Port IPspace Broadcast Domain Link MTU Admin/Oper Status
Status
-----
-----
e0a Cluster Cluster up 9000 auto/10000 healthy
false
e0b Cluster Cluster up 9000 auto/10000 healthy
false
4 entries were displayed.
```

3. Confirm that all the cluster LIFs are on their home ports.

Verify that the “is-home” column is `true` for each of the cluster LIFs:

```
network interface show -vserver Cluster -fields is-home
```

Show example

```
cluster::*> net int show -vserver Cluster -fields is-home
(network interface show)
vserver  lif           is-home
-----  -
Cluster  node1_clus1   true
Cluster  node1_clus2   true
Cluster  node2_clus1   true
Cluster  node2_clus2   true
4 entries were displayed.
```

If there are cluster LIFs that are not on their home ports, revert those LIFs to their home ports:

```
network interface revert -vserver Cluster -lif *
```

4. Disable auto-revert for the cluster LIFs:

```
network interface modify -vserver Cluster -lif * -auto-revert false
```

5. Verify that all ports listed in the previous step are connected to a network switch:

```
network device-discovery show -port cluster_port
```

The “Discovered Device” column should be the name of the cluster switch that the port is connected to.

Show example

The following example shows that cluster ports "e0a" and "e0b" are correctly connected to cluster switches "cs1" and "cs2".

```
cluster::> network device-discovery show -port e0a|e0b
(network device-discovery show)
Node/      Local  Discovered
Protocol  Port   Device (LLDP: ChassisID)  Interface  Platform
-----  -
node1/cdp
          e0a    cs1                      0/11       BES-53248
          e0b    cs2                      0/12       BES-53248
node2/cdp
          e0a    cs1                      0/9        BES-53248
          e0b    cs2                      0/9        BES-53248
4 entries were displayed.
```

6. Verify the cluster connectivity:

```
cluster ping-cluster -node local
```

7. Verify that the cluster is healthy:

```
cluster ring show
```

All units must be either master or secondary.


8. Set up the switchless configuration for the ports in group 1.



To avoid potential networking issues, you must disconnect the ports from group1 and reconnect them back-to-back as quickly as possible, for example, **in less than 20 seconds**.


a. Disconnect all the cables from the ports in group1 at the same time.

In the following example, the cables are disconnected from port "e0a" on each node, and cluster traffic continues through the switch and port "e0b" on each node:



b. Cable the ports in group1 back-to-back.

In the following example, "e0a" on node1 is connected to "e0a" on node2:



9. The switchless cluster network option transitions from `false` to `true`. This might take up to 45 seconds. Confirm that the switchless option is set to `true`:

```
network options switchless-cluster show
```

The following example shows that the switchless cluster is enabled:

```
cluster::*> network options switchless-cluster show
Enable Switchless Cluster: true
```

10. Verify that the cluster network is not disrupted:

```
cluster ping-cluster -node local
```



Before proceeding to the next step, you must wait at least two minutes to confirm a working back-to-back connection on group 1.


11. Set up the switchless configuration for the ports in group 2.



To avoid potential networking issues, you must disconnect the ports from group2 and reconnect them back-to-back as quickly as possible, for example, **in less than 20 seconds**.

- a. Disconnect all the cables from the ports in group2 at the same time.

In the following example, the cables are disconnected from port "e0b" on each node, and cluster traffic continues through the direct connection between the "e0a" ports:



b. Cable the ports in group2 back-to-back.

In the following example, "e0a" on node1 is connected to "e0a" on node2 and "e0b" on node1 is connected to "e0b" on node2:



Step 3: Verify the configuration

1. Verify that the ports on both nodes are correctly connected:

```
network device-discovery show -port cluster_port
```

Show example

The following example shows that cluster ports "e0a" and "e0b" are correctly connected to the corresponding port on the cluster partner:

```
cluster::> net device-discovery show -port e0a|e0b
(network device-discovery show)
Node/      Local  Discovered
Protocol   Port   Device (LLDP: ChassisID)  Interface  Platform
-----
node1/cdp
          e0a    node2                      e0a        AFF-A300
          e0b    node2                      e0b        AFF-A300
node1/lldp
          e0a    node2 (00:a0:98:da:16:44) e0a        -
          e0b    node2 (00:a0:98:da:16:44) e0b        -
node2/cdp
          e0a    node1                      e0a        AFF-A300
          e0b    node1                      e0b        AFF-A300
node2/lldp
          e0a    node1 (00:a0:98:da:87:49) e0a        -
          e0b    node1 (00:a0:98:da:87:49) e0b        -
8 entries were displayed.
```

2. Re-enable auto-revert for the cluster LIFs:

```
network interface modify -vserver Cluster -lif * -auto-revert true
```

3. Verify that all LIFs are home. This might take a few seconds.

```
network interface show -vserver Cluster -lif lif_name
```


Show example

The LIFs have been reverted if the “Is Home” column is `true`, as shown for `node1_clus2` and `node2_clus2` in the following example:

```
cluster::> network interface show -vserver Cluster -fields curr-  
port,is-home  
vserver  lif                curr-port is-home  
-----  -  
Cluster  node1_clus1          e0a      true  
Cluster  node1_clus2          e0b      true  
Cluster  node2_clus1          e0a      true  
Cluster  node2_clus2          e0b      true  
4 entries were displayed.
```

If any cluster LIFS have not returned to their home ports, revert them manually from the local node:

```
network interface revert -vserver Cluster -lif lif_name
```

4. Check the cluster status of the nodes from the system console of either node:

```
cluster show
```

Show example

The following example shows `epsilon` on both nodes to be `false`:

```
Node  Health  Eligibility Epsilon  
-----  
node1 true    true       false  
node2 true    true       false  
2 entries were displayed.
```

5. Confirm connectivity between the cluster ports:

```
cluster ping-cluster local
```

6. If you suppressed automatic case creation, reenable it by invoking an AutoSupport message:

```
system node autosupport invoke -node * -type all -message MAINT=END
```

For more information, see [NetApp KB Article 1010449: How to suppress automatic case creation during scheduled maintenance windows](#).

7. Change the privilege level back to admin:

```
set -privilege admin
```

Copyright information

Copyright © 2024 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.

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

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.