



# **NVIDIA SN2100**

## Cluster and storage switches

NetApp  
April 05, 2024

This PDF was generated from <https://docs.netapp.com/us-en/ontap-systems-switches/switch-nvidia-sn2100/configure-overview-sn2100-storage.html> on April 05, 2024. Always check [docs.netapp.com](https://docs.netapp.com) for the latest.

# Table of Contents

NVIDIA SN2100 .....	1
Overview .....	1
Install hardware .....	3
Configure software .....	6
Migrate switches .....	28
Replace a NVIDIA SN2100 storage switch .....	41

# NVIDIA SN2100

## Overview

### Overview of configuration process for NVIDIA SN2100 storage switches

The NVIDIA SN2100 is a storage switch that allows you to route data between servers and storage arrays in a Storage Area Network (SAN).

#### 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 the NVIDIA documentation.

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

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

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

Follow these 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 script.](#)

There are two RCF scripts available for Clustering and Storage applications.

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](#)

## Configuration requirements for NVIDIA SN2100 switches

For NVIDIA SN2100 switch installation and maintenance, be sure to review all 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/X190106) in the NVIDIA dual/single switch cabinet with the standard brackets that are included with the switch.

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

### ONTAP and Linux support

The NVIDIA SN2100 switch is a 10/25/40/100 Gb Ethernet 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. From ONTAP 9.10.1P3, you can use NVIDIA SN2100 switches to combine storage and cluster functionality into a shared switch configuration.
- Cumulus Linux (CL) OS version 4.4.3. 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/X190106) 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 MSN2100 switches and rail kits:

Part number	Description
X190006-PE	Cluster Switch, NVIDIA SN2100, 16PT 100G, PTSX
X190006-PI	Cluster Switch, NVIDIA SN2100, 16PT 100G, PSIN
X190106-FE-PE	Switch, NVIDIA SN2100, 16PT 100G, PTSX, Front End

Part number	Description
X190106-FE-PI	Switch, NVIDIA SN2100, 16PT 100G, PSIN, Front End
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.

The following table lists the documentation available for the NVIDIA SN2100 switches.

Title	Description
<a href="#">Setup and configure your NVIDIA SN2100 switches</a>	Describes how to setup and configure your NVIDIA SN2100 switches, including installing Cumulus Linux and applicable RCFs.
<a href="#">Migrate from a Cisco cluster switch to a NVIDIA SN2100 cluster switch</a>	Describes how to migrate from environments that use Cisco cluster switches to environments that use NVIDIA SN2100 cluster switches.
<a href="#">Migrate from a Cisco storage switch to a NVIDIA storage switch</a>	Describes how to migrate from environments that use Cisco storage switches to environments that use NVIDIA SN2100 storage switches.
<a href="#">Migrate to a two-node switched cluster with NVIDIA SN2100 cluster switches</a>	Describes how to migrate to a two-node switched environment using NVIDIA SN2100 cluster switches.
<a href="#">Replace a NVIDIA SN2100 cluster switch</a>	Describes the procedure to replace a defective NVIDIA SN2100 switch in a cluster and download Cumulus Linux and reference configuration file.
<a href="#">Replace a NVIDIA SN2100 storage switch</a>	Describes the procedure to replace a defective NVIDIA SN2100 storage switch and download Cumulus Linux and reference configuration file.

## 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	10/40 cluster port nodes
swp2s0-3	25/100 cluster port nodes
swp3-14 40/100 cluster port nodes	swp15-16 40/100 Inter-Switch Link (ISL) ports

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

#### Optical connections

Only optical connections are 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 onboard Intel cluster ports on a platform, not all links come up. Example platforms are: FAS2750, AFF A300, and FAS8200 (all 10G) and AFF A250 (25G).

To resolve this issue, do the following:

1. For Intel 10G, manually set the swp1s0-3 link speed to 10000 and set auto-negotiation to off.
2. For Chelsio 25G, manually set the swp2s0-3 link speed to 25000 and set auto-negotiation to off.



Using 10G/25G QSA, use the non-breakout 40/100G ports. Do not insert the QSA adapter on ports that are configured for breakout.

## Switchport speed

Depending on the transceiver in the switchport, you might need to set the speed on the switchport to fixed speed. If using 10G and 25G breakout ports, make sure that auto-negotiation is off and hard set the port speed on the switch. 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-portadmineedge 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-portadmineedge yes
    mtu 9216

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

## What's next?

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

## Cable 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:

## Information for cabling switch-attached NS224 drive shelves

- Install your storage switches:

[AFF and FAS Switch Documentation](#)

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

[NetApp Hardware Universe](#)

# Configure software

## Software install workflow for NVIDIA SN2100 storage switches

To install and configure the 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 script](#).

There are two RCF scripts available for Clustering and Storage applications.

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 must be 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.

## Steps

1. Log in to the switch.

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

### Show example

```
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
```

## Show example

```
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.

## Show example

```
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.

## Show example

```
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.

- On a terminal, run the following command:

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

- Follow the on-screen menu options to select the geographic area and region.
- To set the time zone for all services and daemons, reboot the switch.
- 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 screens 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
```

**Show example**

```
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
```

## Show example

```
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+cl4.4.3u1
(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:~$
```

## What's next?

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

### About this task

You can install the 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.

### Steps

1. Download the Cumulus Linux installation file to the root directory of the web server. Rename this file `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 displayed, 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
```

**Show example**

```
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:/ # 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  
...  
...
```

- Once the installation has completed, log in to the switch:

**Show example**

```
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>
```

- Verify the Cumulus Linux version:

```
net show version
```

**Show example**

```
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"
```

**What's next?**

[Install RCF script.](#)

## Install the 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 4.4.3 is installed.
- IP address, subnet mask, and default gateway defined via DHCP or manually configured.

### Current RCF script versions

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

- Clustering: **MSN2100-RCF-v1.8-Cluster**
- Storage: **MSN2100-RCF-v1.8-Storage**



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.

### Steps

1. Display the available interfaces on the SN2100 switch:

```
net show interface all
```

## Show example

```
cumulus@cumulus: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:

```
cumulus@cumulus:mgmt:~$ pwd
/home/cumulus
cumulus@cumulus:mgmt: /tmp$ scp <user>@<host:>/<path>/MSN2100-RCF-v1.8-
Cluster
ssologin@10.233.204.71's password:
MSN2100-RCF-v1.8-Cluster                                100% 8607    111.2KB/s
00:00
```

3. Apply the RCF python script **MSN2100-RCF-v1.8-Cluster**:

```
cumulus@cumulus:mgmt:/tmp$ sudo python3 MSN2100-RCF-v1.8-Cluster
[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 above.



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

#### 4. Verify the configuration after the reboot:

```
net show interface all
```

## Show example

cumulus@cumulus:mgmt:~\$ <b>net show interface all</b>						
State	Name	Spd	MTU	Mode	LLDP	Summary
---	---	---	---	---	---	---
...						
...						
DN	swp1s0 bridge(UP)	N/A	9216	Trunk/L2		Master:
DN	swp1s1 bridge(UP)	N/A	9216	Trunk/L2		Master:
DN	swp1s2 bridge(UP)	N/A	9216	Trunk/L2		Master:
DN	swp1s3 bridge(UP)	N/A	9216	Trunk/L2		Master:
DN	swp2s0 bridge(UP)	N/A	9216	Trunk/L2		Master:
DN	swp2s1 bridge(UP)	N/A	9216	Trunk/L2		Master:
DN	swp2s2 bridge(UP)	N/A	9216	Trunk/L2		Master:
DN	swp2s3 bridge(UP)	N/A	9216	Trunk/L2		Master:
UP	swp3 bridge(UP)	100G	9216	Trunk/L2		Master:
UP	swp4 bridge(UP)	100G	9216	Trunk/L2		Master:
DN	swp5 bridge(UP)	N/A	9216	Trunk/L2		Master:
DN	swp6 bridge(UP)	N/A	9216	Trunk/L2		Master:
DN	swp7 bridge(UP)	N/A	9216	Trunk/L2		Master:
DN	swp8 bridge(UP)	N/A	9216	Trunk/L2		Master:
DN	swp9 bridge(UP)	N/A	9216	Trunk/L2		Master:
DN	swp10 bridge(UP)	N/A	9216	Trunk/L2		Master:
DN	swp11 bridge(UP)	N/A	9216	Trunk/L2		Master:
DN	swp12 bridge(UP)	N/A	9216	Trunk/L2		Master:
DN	swp13 bridge(UP)	N/A	9216	Trunk/L2		Master:

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)					
...					
...					
cumulus@cumulus:mgmt:~\$ <b>net show roce config</b>					
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:

```
net show interface pluggables
```

**Show example**

```
cumulus@cumulus: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:

```
net show lldp
```

**Show example**

```
cumulus@cumulus: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.

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

```
network port show -role cluster
```

## Show example

```
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).

## 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      -

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
```

## What's next?

[Configure switch log collection.](#)

## Configure SNMPv3 for switch log collection

Follow this procedure to configure SNMPv3, which supports switch log collection and Switch Health Monitoring (SHM).

### 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 SHM: `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:~$ 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 laErrMsg "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 cshm1! default
```

```

rouser _snmptrapusernameX
+rouser SNMPv3User priv
sysobjectid 1.3.6.1.4.1.40310
sysservices 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 SHM 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: cshm1!
                                         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-LDP 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 SHM 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-LDP Aug-18-2022
```

## Migrate switches

### Migrate from a Cisco storage switch to a NVIDIA SN2100 storage switch

You can migrate older Cisco switches for an ONTAP cluster to NVIDIA SN2100 storage switches. This is a non-disruptive procedure.

#### Review requirements

The following storage switches are supported:

- Cisco Nexus 9336C-FX2
- Cisco Nexus 3232C
- See the [Hardware Universe](#) for full details of supported ports and their configurations.

#### What you'll need

Ensure that:

- The existing cluster is properly set up and functioning.
- All storage ports are in the up state to ensure nondisruptive operations.
- The NVIDIA SN2100 storage switches are configured and operating under the proper version of Cumulus Linux installed with the reference configuration file (RCF) applied.
- The existing storage network configuration has 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 LLFs are on their home ports.
  - ISL ports enabled and cabled between the older Cisco switches and between the new switches.
- See the [Hardware Universe](#) for full details of supported ports and their configurations.
- Some of the ports are configured on NVIDIA SN2100 switches to run at 100 GbE.
- You have planned, migrated, and documented 100 GbE connectivity from nodes to NVIDIA SN2100 storage switches.

## Migrate the switches

### About the examples

In this procedure, Cisco Nexus 9336C-FX2 storage switches are used for example commands and outputs.

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

- The existing Cisco Nexus 9336C-FX2 storage switches are S1 and S2.
- The new NVIDIA SN2100 storage 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 network ports used in this procedure are *e5a* and *e5b*.
- Breakout ports take the format: *swp1s0-3*. For example four breakout ports on *swp1* are *swp1s0*, *swp1s1*, *swp1s2*, and *swp1s3*.
- Switch S2 is replaced by switch sw2 first and then switch S1 is replaced by switch sw1.
  - Cabling between the nodes and S2 are then disconnected from S2 and reconnected to sw2.
  - Cabling between the nodes and S1 are then disconnected from S1 and reconnected to sw1.

### Step 1: Prepare for migration

1. If AutoSupport is enabled, 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. Determine the administrative or operational status for each storage interface:

Each port should display enabled for status.

## Step 2: Configure cables and ports

1. Display the network port attributes:

```
storage port show
```

### Show example

```
cluster1::*> storage port show
          Speed
          Node   Port Type Mode (Gb/s) State Status VLAN ID
----- -----
node1
        e0c  ENET storage 100  enabled online 30
        e0d  ENET storage 0    enabled offline 30
        e5a  ENET storage 0    enabled offline 30
        e5b  ENET storage 100  enabled online 30
node2
        e0c  ENET storage 100  enabled online 30
        e0d  ENET storage 0    enabled offline 30
        e5a  ENET storage 0    enabled offline 30
        e5b  ENET storage 100  enabled online 30
cluster1::*
```

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

```
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
          e0c    S1 (7c:ad:4f:98:6d:f0)    Eth1/1      -
          e5b    S2 (7c:ad:4f:98:8e:3c)    Eth1/1      -
node2      /lldp
          e0c    S1 (7c:ad:4f:98:6d:f0)    Eth1/2      -
          e5b    S2 (7c:ad:4f:98:8e:3c)    Eth1/2      -
```

3. On switch S1 and S2, make sure that the storage ports and switches are connected in the following way (from the switches' perspective) using the command:

```
show lldp neighbors
```

## Show example

```
S1# show lldp neighbors

Capability Codes: (R) Router, (B) Bridge, (T) Telephone, (C) DOCSIS
Cable Device,
(W) WLAN Access Point, (P) Repeater, (S) Station
(O) Other

Device-ID          Local Intf   Holdtime  Capability
Port ID
node1             Eth1/1      121        S
e0c
node2             Eth1/2      121        S
e0c
SHFGD1947000186  Eth1/10     120        S
e0a
SHFGD1947000186  Eth1/11     120        S
e0a
SHFGB2017000269  Eth1/12     120        S
e0a
SHFGB2017000269  Eth1/13     120        S
e0a

S2# show lldp neighbors

Capability Codes: (R) Router, (B) Bridge, (T) Telephone, (C) DOCSIS
Cable Device,
(W) WLAN Access Point, (P) Repeater, (S) Station
(O) Other

Device-ID          Local Intf   Holdtime  Capability
Port ID
node1             Eth1/1      121        S
e5b
node2             Eth1/2      121        S
e5b
SHFGD1947000186  Eth1/10     120        S
e0b
SHFGD1947000186  Eth1/11     120        S
e0b
SHFGB2017000269  Eth1/12     120        S
e0b
SHFGB2017000269  Eth1/13     120        S
e0b
```

4. On switch sw2, shut down the ports connected to the storage ports and nodes of the disk shelves.

**Show example**

```
cumulus@sw2:~$ net add interface swp1-16 link down
cumulus@sw2:~$ net pending
cumulus@sw2:~$ net commit
```

5. Move the node storage ports of the controller and disk shelves from the old switch S2 to the new switch sw2, using appropriate cabling supported by NVIDIA SN2100.

6. On switch sw2, bring up the ports connected to the storage ports of the nodes and the disk shelves.

**Show example**

```
cumulus@sw2:~$ net del interface swp1-16 link down
cumulus@sw2:~$ net pending
cumulus@sw2:~$ net commit
```

7. Verify that the storage ports on each node are now connected to the 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
          e0c     S1 (7c:ad:4f:98:6d:f0)    Eth1/1        -
          e5b     sw2 (b8:ce:f6:19:1a:7e)    swp1         -
          e0c     S1 (7c:ad:4f:98:6d:f0)    Eth1/2        -
          e5b     sw2 (b8:ce:f6:19:1a:7e)    swp2         -
```

8. Verify the network port attributes:

```
storage port show
```

**Show example**

```
cluster1::*> storage port show
                                         Speed
                                         (Gb/s)
Node      Port Type Mode      State    Status     VLAN ID
-----  -----  -----  -----  -----  -----
node1
        e0c  ENET storage  100  enabled  online   30
        e0d  ENET storage   0  enabled  offline  30
        e5a  ENET storage   0  enabled  offline  30
        e5b  ENET storage  100  enabled  online   30
node2
        e0c  ENET storage  100  enabled  online   30
        e0d  ENET storage   0  enabled  offline  30
        e5a  ENET storage   0  enabled  offline  30
        e5b  ENET storage  100  enabled  online   30
cluster1::*
```

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

```
net show interface
```

**Show example**

```
cumulus@sw2:~$ 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) )
...
...
```

10. On switch sw1, shut down the ports connected to the storage ports of the nodes and the disk shelves.

**Show example**

```
cumulus@sw1:~$ net add interface swp1-16 link down
cumulus@sw1:~$ net pending
cumulus@sw1:~$ net commit
```

11. Move the node storage ports of the controller and the disk shelves from the old switch S1 to the new switch sw1, using appropriate cabling supported by NVIDIA SN2100.
12. On switch sw1, bring up the ports connected to the storage ports of the nodes and the disk shelves.

**Show example**

```
cumulus@sw1:~$ net del interface swp1-16 link down
cumulus@sw1:~$ net pending
cumulus@sw1:~$ net commit
```

13. Verify that the storage ports on each node are now connected to the 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/Protocol/Platform	Local Port	Discovered Device (LLDP: ChassisID)	Interface	
node1	/lldp	e0c sw1 (b8:ce:f6:19:1b:96)	swp1	-
		e5b sw2 (b8:ce:f6:19:1a:7e)	swp1	-
node2	/lldp	e0c sw1 (b8:ce:f6:19:1b:96)	swp2	-
		e5b sw2 (b8:ce:f6:19:1a:7e)	swp2	-

14. Verify the final configuration:

```
storage port show
```

Each port should display enabled for State and enabled for Status.

**Show example**

```
cluster1::*> storage port show
                                         Speed
                                         (Gb/s)
Node      Port Type Mode      State    Status     VLAN ID
-----  -----  -----  -----  -----  -----
node1
        e0c  ENET storage  100  enabled  online   30
        e0d  ENET storage   0  enabled  offline  30
        e5a  ENET storage   0  enabled  offline  30
        e5b  ENET storage  100  enabled  online   30
node2
        e0c  ENET storage  100  enabled  online   30
        e0d  ENET storage   0  enabled  offline  30
        e5a  ENET storage   0  enabled  offline  30
        e5b  ENET storage  100  enabled  online   30
cluster1::*
```

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

```
net show interface
```

## Show example

```
cumulus@sw2:~$ 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) )
...
...
```

16. 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
-----  -----  -----  -----
...
swp1      100G   Trunk/L2  node1                e0c
swp2      100G   Trunk/L2  node2                e0c
swp3      100G   Trunk/L2  SHFFG1826000112       e0a
swp4      100G   Trunk/L2  SHFFG1826000112       e0a
swp5      100G   Trunk/L2  SHFFG1826000102       e0a
swp6      100G   Trunk/L2  SHFFG1826000102       e0a

cumulus@sw2:~$ net show lldp
LocalPort  Speed  Mode      RemoteHost          RemotePort
-----  -----  -----  -----
...
swp1      100G   Trunk/L2  node1                e5b
swp2      100G   Trunk/L2  node2                e5b
swp3      100G   Trunk/L2  SHFFG1826000112       e0b
swp4      100G   Trunk/L2  SHFFG1826000112       e0b
swp5      100G   Trunk/L2  SHFFG1826000102       e0b
swp6      100G   Trunk/L2  SHFFG1826000102       e0b
```

## Step 3: Complete the procedure

1. Enable the Ethernet switch health monitor log collection feature for collecting switch-related log files, using the two commands:

```
system switch ethernet log setup-password and system switch ethernet log enable-collection
```

Enter: 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
RSA key fingerprint is
e5:8b:c6:dc:e2:18:18:09:36:63:d9:63:dd:03:d9:cc
Do you want to continue? {y|n}::[n] y

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
RSA key fingerprint is
57:49:86:a1:b9:80:6a:61:9a:86:8e:3c:e3:b7:1f:b1
Do you want to continue? {y|n}:: [n] y

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

Followed by:

```
system switch ethernet log enable-collection
```

### Show example

```
cluster1::*> system switch ethernet log enable-collection

Do you want to enable cluster log collection for all nodes in the
cluster?
{y|n}: [n] y

Enabling cluster switch log collection.

cluster1::*
```



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

### 2. Initiate the switch log collection feature:

```
system switch ethernet log collect -device *
```

Wait for 10 minutes and then check that the log collection was successful using the command:

```
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
```

### 3. Change the privilege level back to admin:

```
set -privilege admin
```

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

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

## Replace a NVIDIA SN2100 storage switch

You must be aware of certain configuration information, port connections and cabling

requirements when you replace NVIDIA SN2100 storage switches.

## Before you begin

You must verify that the following conditions exist before installing the Cumulus software and RCFs on a NVIDIA SN2100 storage switch:

- Your system can support NVIDIA SN2100 storage switches.
- You must have downloaded the applicable RCFs.
- The [Hardware Universe](#) provides full details of supported ports and their configurations.

## About this task

The existing network configuration must have the following characteristics:

- Make sure that all troubleshooting steps have been completed to confirm that your switch needs replacing.
- Management connectivity must exist on both switches.



Make sure that all troubleshooting steps have been completed to confirm that your switch needs replacing.

The replacement NVIDIA SN2100 switch must have the following characteristics:

- Management network connectivity must be functional.
- Console access to the replacement switch must be in place.
- The appropriate RCF and Cumulus operating system image must be loaded onto the switch.
- Initial customization of the switch must be complete.

## Procedure summary

This procedure replaces the second NVIDIA SN2100 storage switch sw2 with the new NVIDIA SN2100 switch nsw2. The two nodes are node1 and node2.

Steps to complete:

- Confirm the switch to be replaced is sw2.
- Disconnect the cables from switch sw2.
- Reconnect the cables to switch nsw2.
- Verify all device configurations on switch nsw2.

## Steps

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`  
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`
3. Check on the health status of the storage node ports to make sure that there is connection to storage switch S1:

```
storage port show -port-type ENET
```

#### Show example

```
cluster1::*> storage port show -port-type ENET
                                         Speed
                                         (Gb/s)      VLAN
Node       Port Type Mode      State     Status      ID
-----  -----
node1
        e3a  ENET  storage  100    enabled   online    30
        e3b  ENET  storage   0    enabled   offline    30
        e7a  ENET  storage   0    enabled   offline    30
        e7b  ENET  storage  100    enabled   online    30
node2
        e3a  ENET  storage  100    enabled   online    30
        e3b  ENET  storage   0    enabled   offline    30
        e7a  ENET  storage   0    enabled   offline    30
        e7b  ENET  storage  100    enabled   online    30
cluster1::*
```

4. Verify that storage switch sw1 is available: network device-discovery show

#### 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:f6:19:1b:42)    swp3      -
node2/lldp
        e3a    sw1 (b8:ce:f6:19:1b:42)    swp4      -
cluster1::*
```

5. Run the net show interface command on the working switch to confirm that you can see both nodes and all shelves: net show interface

## Show example

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

State   Name      Spd     MTU      Mode          LLDP
Summary
-----  -----  -----  -----  -----
...
...
UP      swp1     100G    9216    Trunk/L2    node1  (e3a)
Master: bridge (UP)
UP      swp2     100G    9216    Trunk/L2    node2  (e3a)
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) )
...
...
```

6. Verify the shelf ports in the storage system: `storage shelf port show -fields remote-device, remote-port`

### Show example

```
cluster1::*> storage shelf port show -fields remote-device, remote-port
shelf    id  remote-port   remote-device
-----  --  -----
3.20     0    swp3        sw1
3.20     1    -           -
3.20     2    swp4        sw1
3.20     3    -           -
3.30     0    swp5        sw1
3.20     1    -           -
3.30     2    swp6        sw1
3.20     3    -           -
cluster1::*
```

7. Remove all cables attached to storage switch sw2.
8. Reconnect all cables to the replacement switch nsw2.
9. Recheck the health status of the storage node ports: `storage port show -port-type ENET`

### Show example

```
cluster1::*> storage port show -port-type ENET
                                         Speed
                                         (Gb/s)
Node      Port Type  Mode      State   Status   VLAN ID
-----  -----  -----  -----  -----  -----  -----
node1
          e3a  ENET  storage  100  enabled  online   30
          e3b  ENET  storage   0  enabled  offline   30
          e7a  ENET  storage   0  enabled  offline   30
          e7b  ENET  storage  100  enabled  online   30
node2
          e3a  ENET  storage  100  enabled  online   30
          e3b  ENET  storage   0  enabled  offline   30
          e7a  ENET  storage   0  enabled  offline   30
          e7b  ENET  storage  100  enabled  online   30
cluster1::*
```

10. Verify that both switches are available: `net device-discovery show`

**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:f6:19:1b:96)    swp1        -
          e7b    nsw2  (b8:ce:f6:19:1a:7e)    swp1        -
node2/lldp
          e3a    sw1  (b8:ce:f6:19:1b:96)    swp2        -
          e7b    nsw2  (b8:ce:f6:19:1a:7e)    swp2        -
cluster1::*
```

11. Verify the shelf ports in the storage system: `storage shelf port show -fields remote-device, remote-port`

**Show example**

```
cluster1::*> storage shelf port show -fields remote-device, remote-
port
shelf   id     remote-port     remote-device
-----  --     -----
3.20    0      swp3           sw1
3.20    1      swp3           nsw2
3.20    2      swp4           sw1
3.20    3      swp4           nsw2
3.30    0      swp5           sw1
3.20    1      swp5           nsw2
3.30    2      swp6           sw1
3.20    3      swp6           nsw2
cluster1::*
```

12. 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:
sw1
nsw2

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

Enter the switch name: csw1
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: nsw2
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>
```

13. 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	sw1 (b8:ce:f6:19:1b:42)	4/29/2022 03:05:25	complete
2	nsw2 (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.

14. Change the privilege level back to admin: `set -privilege admin`
15. If you suppressed automatic case creation, re-enable it by invoking an AutoSupport message: `system node autosupport invoke -node * -type all -message MAINT=END`

## **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—with 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.