



Migrate switches

Cluster and storage switches

NetApp
April 25, 2024

Table of Contents

- Migrate switches. 1
 - Migrate CN1610 cluster switches to NVIDIA SN2100 cluster switches 1
 - Migrate from a Cisco cluster switch to a NVIDIA SN2100 cluster switch 18
 - Migrate to a two-node switched cluster with NVIDIA SN2100 cluster switches 34

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 your 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  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 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

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 -ip space 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				
	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):

```
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
```

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.