



Migrate switches

Cluster and storage switches

NetApp
April 25, 2024

Table of Contents

Migrate switches.	1
Migrate from a Cisco storage switch to a NVIDIA SN2100 storage switch	1

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 LIFs 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
```

Node	Port	Type	Mode	Speed (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 Port ID	Local Intf	Holdtime	Capability
node1 e0c	Eth1/1	121	S
node2 e0c	Eth1/2	121	S
SHFGD1947000186 e0a	Eth1/10	120	S
SHFGD1947000186 e0a	Eth1/11	120	S
SHFGB2017000269 e0a	Eth1/12	120	S
SHFGB2017000269 e0a	Eth1/13	120	S

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

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/ Protocol	Local Port	Discovered 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	-
node2	/lldp			
	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
```

Node	Port	Type	Mode	Speed (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::*>
```

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
```

Node	Port	Type	Mode	Speed (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::*>
```

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
```

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.