



Install hardware

Cluster and storage switches

NetApp
April 25, 2024

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

Table of Contents

- Install hardware 1
 - Install the hardware for the NVIDIA SN2100 switch..... 1
 - Configure the NVIDIA SN2100 switch 1
 - Review cabling and configuration considerations 1
 - Cable the NS224 shelves as switch-attached storage..... 9

Install hardware

Install the hardware for the NVIDIA SN2100 switch

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

Steps

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

What's next?

[Configure the switch](#).

Configure the NVIDIA SN2100 switch

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

Steps

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

What's next?

[Review cabling and configuration considerations](#).

Review cabling and configuration considerations

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

NVIDIA port details

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

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

Link-up delays with optical connections

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

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

Show example

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

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

```
Only switchd reload required
```

Support for copper connections

The following configuration changes are required to fix this issue.

Cumulus Linux 4.4.3

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

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

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

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

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

For example:

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

3. Restart the `switchd` service:

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

4. Confirm that the ports are up:

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

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

Cumulus Linux 5.x

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

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

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

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

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

For example:

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

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

Only switchd reload required
```

3. Confirm that the ports are up:

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

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

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

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

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

QSA adapter

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

To resolve this issue, do the following:

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



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

Setting interface speed on breakout ports

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

Cumulus Linux 4.4.3

For example:

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

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

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

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


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

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

Cumulus Linux 5.x

For example:

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

```
link
```

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

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

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

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

What's next?

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

Cable the NS224 shelves as switch-attached storage

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

- Cable NS224 drive shelves through storage switches:

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

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

[NetApp Hardware Universe](#)

What's next?

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

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.