



Configuring a Switch Running Cumulus® Linux® for Lossless Networking with OpenFlex™ Platforms

Abstract

This configuration guide provides an overview of how to configure lossless Ethernet settings on an Ethernet switch that is running the Cumulus Linux Operating System for use with Western Digital® OpenFlex platforms.

Table of Contents

Introduction	2
Configuration Process Summary	2
Example Hardware Specifications	2
Configuration Table.....	2
Configure Ports.....	3
Bring Up Interfaces	3
Configure PFC	5
Configure ENC.....	6
Show Pertinent Switch Counters	6

Introduction

NVMe-oF™ based storage offers the promise of low latency shared storage. To obtain the performance potential of this technology, Ethernet Network Adapters in initiator hosts must be configured for lossless networking using standard Data Center Bridging (DCB) technologies.

Configuration Process Summary

The process of enabling lossless networking functionality on Broadcom® based Ethernet Network Adapters can be broken down into the following steps:

1. Configure Ports
2. Configure MTU
3. Configure Primary Bridge
4. Configure Priority Flow Control (PFC)
5. Configure Explicit Congestion Notification (ECN)
6. Show Pertinent Switch Counters

Example Hardware Specifications

In this guide, an SN4600 with Cumulus Linux version 5.6 was used. Terminal output shown in this guide may vary based on the product and firmware version. For additional information, see the [Data24 Compatibility Matrix](#).

The following instructions will be making many changes to the switch. Remember to write memory to ensure no settings are lost during a power outage event. There will be many changes to “nv config apply” and “nv config save” in the switch to ensure no settings are lost during a power outage event.

Configuration Table

Included below for convenience is a table to record the lossless configuration values for deployment.

Description	Variable	Example	Deployment Value
PFC Priority	<ROCE_PRI>	3	
PFC DSCP	<ROCE_DSCP>	24	
CNP Priority	<CNP_PRI>	6	
CNP DSCP	<CNP_DSCP>	48	
Port	<PORT>	swp1 or swp1s0	

Configure Ports

The purpose of this section is to describe how to configure the ports on the switch so that they will successfully connect to the target product.

Bring Up Interfaces

1. Bring up all the interfaces

```
:~$ nv set interface swp1-64
:~$ nv set interface swp1-64 bridge domain br_default
```

2. Show interfaces configuration

```
:~$ nv show interface
```

Example Output:

```
cumulus@OCCLMLSN4600-2:mgmt:~$ nv show interface
Interface      State   Speed   MTU     Type      Remote Host   Remote Port   Summary
-----  -----  -----  -----  -----
br_default    up      9216    bridge
fe80::1e34:daff:fed4:22fd/64                                IP Address:
eth0          up      1G      1500    eth       OCCLHPR01S23U  38      IP Address:
192.168.3.29/24
fe80::bace:f6ff:fed0:9730/64                                IP Address:
lo            up      65536   loopback
127.0.0.1/8
::1/128
mgmt          up      65575   vrf
127.0.0.1/8
::1/128
swp1          up      200G    9216    swp
swp2          up      200G    9216    swp
swp3          up      200G    9216    swp
swp4          up      200G    9216    swp
swp5          up      200G    9216    swp
swp6          up      200G    9216    swp
swp7          up      200G    9216    swp
swp8          up      200G    9216    swp
swp9          up      200G    9216    swp
swp10         up      200G    9216    swp
swp11         down    9216    swp
swp12         down    9216    swp
swp13         down    9216    swp
swp14         down    9216    swp
swp15         down    9216    swp
swp16         down    9216    swp
swp17         down    9216    swp
swp18         down    9216    swp
swp19         down    9216    swp
swp20         down    9216    swp
swp21         down    9216    swp
swp22         down    9216    swp
swp23         down    9216    swp
swp24         down    9216    swp
swp25         down    9216    swp
```

swp26	down	9216	swp
swp27	down	9216	swp
swp28	down	9216	swp
swp29	down	9216	swp
swp30	down	9216	swp
swp31	down	9216	swp
swp32	down	9216	swp
swp33	down	9216	swp
swp34	down	9216	swp
swp35	down	9216	swp
swp36	down	9216	swp
swp37	down	9216	swp
swp38	down	9216	swp
swp39	down	9216	swp
swp40	down	9216	swp
swp41	down	9216	swp
swp42	down	9216	swp
swp43	down	9216	swp
swp44	down	9216	swp
swp45	down	9216	swp
swp46	down	9216	swp
swp47	down	9216	swp
swp48	down	9216	swp
swp49	down	9216	swp
swp50	down	9216	swp
swp51	down	9216	swp
swp52	down	9216	swp
swp53	down	9216	swp
swp54	down	9216	swp
swp55	down	9216	swp
swp56	down	9216	swp
swp57	down	9216	swp
swp58	down	9216	swp
swp59	up	100G	swp
swp60	up	100G	swp
swp61	up	100G	swp
swp62	up	100G	swp
swp63	up	100G	swp
swp64	up	100G	swp

Configure PFC

1. Configure the switch for RoCE:

```
:~$ nv set qos roce enable on
:~$ nv set qos roce mode lossless
:~$ nv set qos mapping default-global trust 13
:~$ nv set qos mapping default-global dscp <ROCE _ DSCP> switch-priority <ROCE _ PRI>
```

2. Show the qos mapping settings:

```
nv show qos mapping default-global dscp <ROCE _ DSCP>
```

Example:

```
cumulus@OCCLMLSN4600-2:mgmt:~$ nv show qos mapping default-global dscp 24
                                operational applied
-----
switch-priority 3            3
```

Note: To apply a custom DSCP profile to specific interfaces, “Port Groups” can be used.

```
cumulus@OCCLMLSN4600-2:mgmt:~$ nv set qos mapping myportgroup dscp 34 switch-priority 3
cumulus@OCCLMLSN4600-2:mgmt:~$ nv set qos mapping myportgroup trust 13
```

Note: Set qos portgroup “myportgroup” on ports 1-6 and 9-20.

```
cumulus@OCCLMLSN4600-2:mgmt:~$ nv set interface swp1-6,9-20 qos mapping profile myportgroup
```

Configure ECN

1. Execute the following commands:

```
:~$ nv set qos congestion-control default-global traffic-class <ROCE _ PRI> min-threshold 40000
```

Note: :~\$ nv set qos congestion-control default-global traffic-class <ROCE_PRI> max-threshold 200000To apply a custom ECN profile to specific interfaces, “Port Groups” can be used.

```
:~$ nv set qos congestion-control myecnprofile traffic-class 3 min-threshold 40000
:~$ nv set qos congestion-control myecnprofile traffic-class 3 max-threshold 200000
```

Note: Set qos portgroup “myecnprofile” on ports 1-6 and 9-20.

```
:~$ nv set interface swp1-6,9-20 qos congestion-control profile myecnprofile
```

2. Show the congestion control configuration.

```
:~$ nv show qos congestion-control default-global
cumulus@OCCLMLSN4600-2:mgmt:~$ nv show qos congestion-control default-global
ECN Configurations
=====
traffic-class  ECN      RED      Min Th      Max Th      Probability
-----  -----  -----  -----  -----  -----
0          enable  disable  146.48 KB  1.43 MB  100
3          enable  disable  146.48 KB  1.43 MB  100
```

Show Pertinent Switch Counters

1. Show detailed interface information including counters

```
:~$ net show interface <PORT> detail
```

Example Selections from Output:

cl-netstat Counters: Include Ok, Error, Drop, & Over for the Port

```
cumulus@OCCLMLSN4600-2:mgmt:~$ net show interface swp2 detail
      Name  MAC                  Speed  MTU    Mode
-----  -----  -----  -----  -----
UP    swp2  1c:34:da:c4:22:6a  200G   9216  Access/L2

All VLANs on L2 Port
-----
7
```

```
Untagged
-----
7

cl-netstat counters
-----
RX _OK RX _ERR RX _DRP RX _OVR TX _OK TX _ERR TX _DRP TX _OVR
----- ----- ----- ----- -----
0 0 0 0 952 0 0 0

Routing
-----
Interface swp2 is up, line protocol is up
Link ups: 2 last: 2023/10/19 13:17:03.18
Link downs: 4 last: 2023/10/19 13:16:19.65
PTM status: disabled
vrf: default
index 5 metric 0 mtu 9216 speed 200000
flags: <UP,BROADCAST,RUNNING,MULTICAST>
Ignore all v6 routes with linkdown
Type: Ethernet
HWaddr: 1c:34:da:c4:22:6a
Interface Type Other
Interface Slave Type Bridge
Master interface: br _default
protodown: off

Ethtool
-----
Settings for swp2:
Supported ports: [ FIBRE ]
Supported link modes: 40000baseCR4/Full
40000baseSR4/Full
40000baseLR4/Full
100000baseSR4/Full
100000baseCR4/Full
100000baseLR4 _ ER4/Full
200000baseSR4/Full
200000baseDR4/Full
200000baseCR4/Full

Supported pause frame use: Symmetric
Supports auto-negotiation: Yes
Supported FEC modes: None RS BASER
Advertised link modes: 40000baseCR4/Full
40000baseSR4/Full
40000baseLR4/Full
100000baseSR4/Full
100000baseCR4/Full
100000baseLR4 _ ER4/Full
200000baseSR4/Full
```

```

200000baseDR4/Full
200000baseCR4/Full

Advertised pause frame use: Symmetric
Advertised auto-negotiation: Yes
Advertised FEC modes: None      RS      BASER
Speed: 200000Mb/s
Duplex: Full
Auto-negotiation: on
Port: None
PHYAD: 0
Transceiver: internal
Link detected: yes

Ethtool Pluggables
-----
Identifier : 0x11 (QSFP28)
Extended identifier : 0x00
Extended identifier description : 1.5W max. Power consumption
Extended identifier description : No CDR in TX, No CDR in RX
Extended identifier description : High Power Class (> 3.5 W) not enabled
Power set : Off
Power override : Off
Connector : 0x23 (No separable connector)
Transceiver codes : 0x88 0x00 0x00 0x04 0x00 0x00 0x00 0x00
Transceiver type : 40G Ethernet: 40G Base-CR4
Transceiver type : 50GBASE-CR, 100GBASE-CR2, or 200GBASE-CR4
Transceiver type : Ethernet: 1000BASE-CX
Encoding : 0x08 (PAM4)
BR, Nominal : 25500Mbps
Rate identifier : 0x00
Length (SMF,km) : 0km
Length (OM3 50um) : 0m
Length (OM2 50um) : 0m
Length (OM1 62.5um) : 0m
Length (Copper or Active cable) : 3m
Transmitter technology : 0xa0 (Copper cable unequalized)
Attenuation at 2.5GHz : 4db
Attenuation at 5.0GHz : 6db
Attenuation at 7.0GHz : 8db
Attenuation at 12.9GHz : 11db
Vendor name : Mellanox
Vendor OUI : 00:02:c9
Vendor PN : MCP1650-V003E26
Vendor rev : A2
Vendor SN : MT2220VB30301
Date code : 220519
Revision Compliance : Unallocated
Module temperature : 0.00 degrees C / 32.00 degrees F
Module voltage : 0.0000 V

```

```
Ethtool Statistics
-----
NIC statistics:
    HwIfInOctets: 0
    HwIfInUcastPkts: 0
    HwIfInBcastPkts: 0
    HwIfInMcastPkts: 0
    HwIfOutOctets: 29996433
    HwIfOutUcastPkts: 41
    HwIfOutMcastPkts: 380036
    HwIfOutBcastPkts: 2658
    HwIfInDiscards: 0
    HwIfInL3Drops: 0
    HwIfInBufferDrops: 0
    HwIfInAclDrops: 0
    HwIfInDot3LengthErrors: 0
    HwIfInErrors: 0
    SoftInErrors: 0
    SoftInDrops: 0
    SoftInFrameErrors: 0
    HwIfOutDiscards: 0
    HwIfOutErrors: 0
    HwIfOutQDrops: 0
    HwIfOutNonQDrops: 0
    SoftOutErrors: 0
    SoftOutDrops: 0
    SoftOutTxFifoFull: 0
    HwIfOutQLen: 0
    HwIfInDot3FrameErrors: 0
    HwIfInPausePkt: 0
    HwIfOutPausePkt: 0
    HwIfInPfc0Pkt: 0
    HwIfOutPfc0Pkt: 0
    HwIfInPfc1Pkt: 0
    HwIfOutPfc1Pkt: 0
    HwIfInPfc2Pkt: 0
    HwIfOutPfc2Pkt: 0
    HwIfInPfc3Pkt: 0
    HwIfOutPfc3Pkt: 0
    HwIfInPfc4Pkt: 0
    HwIfOutPfc4Pkt: 0
    HwIfInPfc5Pkt: 0
    HwIfOutPfc5Pkt: 0
    HwIfInPfc6Pkt: 0
    HwIfOutPfc6Pkt: 0
    HwIfInPfc7Pkt: 0
    HwIfOutPfc7Pkt: 0
    HwIfOutWredDrops: 0
    HwIfOutQ0WredDrops: 0
    HwIfOutQ1WredDrops: 0
    HwIfOutQ2WredDrops: 0
```

```
HwIfOutQ3WredDrops: 0
HwIfOutQ4WredDrops: 0
HwIfOutQ5WredDrops: 0
HwIfOutQ6WredDrops: 0
HwIfOutQ7WredDrops: 0
HwIfOutQ8WredDrops: 0
HwIfOutQ9WredDrops: 0
HwIfAlignmentErrors: 0
HwIfSymbolErrors: 0
HwIfOversizeErrors: 0
HwIfUndersizeErrors: 0
HwIfJabberErrors: 0
HwIfIn64OctetsPkts: 0
HwIfIn65To127OctetsPkts: 0
HwIfIn128To255OctetsPkts: 0
HwIfIn256To511OctetsPkts: 0
HwIfIn512To1023OctetsPkts: 0
HwIfIn1024To1518OctetsPkts: 0
HwIfIn1519To2047OctetsPkts: 0
HwIfIn2048To4095OctetsPkts: 0
HwIfIn4096To16383OctetsPkts: 0
HwIfOut64OctetsPkts: 358642
HwIfOut65To127OctetsPkts: 215
HwIfOut128To255OctetsPkts: 24
HwIfOut256To511OctetsPkts: 23852
HwIfOut512To1023OctetsPkts: 0
HwIfOut1024To1518OctetsPkts: 2
HwIfOut1519To2047OctetsPkts: 0
HwIfOut2048To4095OctetsPkts: 0
HwIfOut4096To16383OctetsPkts: 0
HwIfOutQ0BuffDiscards: 0
HwIfOutQ1BuffDiscards: 0
HwIfOutQ2BuffDiscards: 0
HwIfOutQ3BuffDiscards: 0
HwIfOutQ4BuffDiscards: 0
HwIfOutQ5BuffDiscards: 0
HwIfOutQ6BuffDiscards: 0
HwIfOutQ7BuffDiscards: 0
HwIfOutQ8BuffDiscards: 0
HwIfOutQ9BuffDiscards: 0
HwIfOutQ0Pkts: 2707
HwIfOutQ1Pkts: 0
HwIfOutQ2Pkts: 0
HwIfOutQ3Pkts: 0
HwIfOutQ4Pkts: 0
HwIfOutQ5Pkts: 0
HwIfOutQ6Pkts: 0
HwIfOutQ7Pkts: 380028
HwIfOutQ8Pkts: 0
HwIfOutQ9Pkts: 0
HwIfOutEcnMarkedPkts: 0
```

```
HwIfInEtherOctets: 0
HwIfOutQ0Octets: 182589
HwIfOutQ1Octets: 0
HwIfOutQ2Octets: 0
HwIfOutQ3Octets: 0
HwIfOutQ4Octets: 0
HwIfOutQ5Octets: 0
HwIfOutQ6Octets: 0
HwIfOutQ7Octets: 29813844
HwIfOutQ8Octets: 0
HwIfOutQ9Octets: 0
HwIfInPfc0Duration: 0
HwIfInPfc1Duration: 0
HwIfInPfc2Duration: 0
HwIfInPfc3Duration: 0
HwIfInPfc4Duration: 0
HwIfInPfc5Duration: 0
HwIfInPfc6Duration: 0
HwIfInPfc7Duration: 0
HwIfOutPfc0Duration: 0
HwIfOutPfc1Duration: 0
HwIfOutPfc2Duration: 0
HwIfOutPfc3Duration: 0
HwIfOutPfc4Duration: 0
HwIfOutPfc5Duration: 0
HwIfOutPfc6Duration: 0
HwIfOutPfc7Duration: 0
HwIfInQ0Pkts: 0
HwIfInQ1Pkts: 0
HwIfInQ2Pkts: 0
HwIfInQ3Pkts: 0
HwIfInQ4Pkts: 0
HwIfInQ5Pkts: 0
HwIfInQ6Pkts: 0
HwIfInQ7Pkts: 0
HwIfInQ0BuffDiscards: 0
HwIfInQ1BuffDiscards: 0
HwIfInQ2BuffDiscards: 0
HwIfInQ3BuffDiscards: 0
HwIfInQ4BuffDiscards: 0
HwIfInQ5BuffDiscards: 0
HwIfInQ6BuffDiscards: 0
HwIfInQ7BuffDiscards: 0
HwIfInQ0SharedBuffDiscards: 0
HwIfInQ1SharedBuffDiscards: 0
HwIfInQ2SharedBuffDiscards: 0
HwIfInQ3SharedBuffDiscards: 0
HwIfInQ4SharedBuffDiscards: 0
HwIfInQ5SharedBuffDiscards: 0
HwIfInQ6SharedBuffDiscards: 0
HwIfInQ7SharedBuffDiscards: 0
HwIfOutMcBuffDiscards: 0
```

Section Ethtool Statistics

```
cumulus@OCCLMLSN4600-2:mgmt:~$ ethtool -S swp1
NIC statistics:
    HwIfInOctets: 107005
    HwIfInUcastPkts: 0
    HwIfInBcastPkts: 13
    HwIfInMcastPkts: 1106
    HwIfOutOctets: 29922181
    HwIfOutUcastPkts: 41
    HwIfOutMcastPkts: 379086
    HwIfOutBcastPkts: 2648
    HwIfInDiscards: 0
    HwIfInL3Drops: 0
    HwIfInBufferDrops: 0
    HwIfInAclDrops: 0
    HwIfInDot3LengthErrors: 0
    HwIfInErrors: 0
    SoftInErrors: 0
    SoftInDrops: 0
    SoftInFrameErrors: 0
    HwIfOutDiscards: 0
    HwIfOutErrors: 0
    HwIfOutQDrops: 0
    HwIfOutNonQDrops: 0
    SoftOutErrors: 0
    SoftOutDrops: 0
    SoftOutTxFifoFull: 0
    HwIfOutQLen: 0
    HwIfInDot3FrameErrors: 0
    HwIfInPausePkt: 0
    HwIfOutPausePkt: 0
    HwIfInPfc0Pkt: 0
    HwIfOutPfc0Pkt: 0
    HwIfInPfc1Pkt: 0
    HwIfOutPfc1Pkt: 0
    HwIfInPfc2Pkt: 0
    HwIfOutPfc2Pkt: 0
    HwIfInPfc3Pkt: 0
    HwIfOutPfc3Pkt: 0
    HwIfInPfc4Pkt: 0
    HwIfOutPfc4Pkt: 0
    HwIfInPfc5Pkt: 0
    HwIfOutPfc5Pkt: 0
    HwIfInPfc6Pkt: 0
    HwIfOutPfc6Pkt: 0
    HwIfInPfc7Pkt: 0
    HwIfOutPfc7Pkt: 0
    HwIfOutWredDrops: 0
    HwIfOutQ0WredDrops: 0
    HwIfOutQ1WredDrops: 0
    HwIfOutQ2WredDrops: 0
```

```
HwIfOutQ3WredDrops: 0
HwIfOutQ4WredDrops: 0
HwIfOutQ5WredDrops: 0
HwIfOutQ6WredDrops: 0
HwIfOutQ7WredDrops: 0
HwIfOutQ8WredDrops: 0
HwIfOutQ9WredDrops: 0
HwIfAlignmentErrors: 0
HwIfSymbolErrors: 0
HwIfOversizeErrors: 0
HwIfUndersizeErrors: 0
HwIfJabberErrors: 0
HwIfIn64OctetsPkts: 13
HwIfIn65To127OctetsPkts: 1106
HwIfIn128To255OctetsPkts: 0
HwIfIn256To511OctetsPkts: 0
HwIfIn512To1023OctetsPkts: 0
HwIfIn1024To1518OctetsPkts: 0
HwIfIn1519To2047OctetsPkts: 0
HwIfIn2048To4095OctetsPkts: 0
HwIfIn4096To16383OctetsPkts: 0
HwIfOut64OctetsPkts: 357742
HwIfOut65To127OctetsPkts: 211
HwIfOut128To255OctetsPkts: 24
HwIfOut256To511OctetsPkts: 23796
HwIfOut512To1023OctetsPkts: 0
HwIfOut1024To1518OctetsPkts: 2
HwIfOut1519To2047OctetsPkts: 0
HwIfOut2048To4095OctetsPkts: 0
HwIfOut4096To16383OctetsPkts: 0
HwIfOutQ0BuffDiscards: 0
HwIfOutQ1BuffDiscards: 0
HwIfOutQ2BuffDiscards: 0
HwIfOutQ3BuffDiscards: 0
HwIfOutQ4BuffDiscards: 0
HwIfOutQ5BuffDiscards: 0
HwIfOutQ6BuffDiscards: 0
HwIfOutQ7BuffDiscards: 0
HwIfOutQ8BuffDiscards: 0
HwIfOutQ9BuffDiscards: 0
HwIfOutQ0Pkts: 2695
HwIfOutQ1Pkts: 0
HwIfOutQ2Pkts: 0
HwIfOutQ3Pkts: 0
HwIfOutQ4Pkts: 0
HwIfOutQ5Pkts: 0
HwIfOutQ6Pkts: 0
HwIfOutQ7Pkts: 379080
HwIfOutQ8Pkts: 0
HwIfOutQ9Pkts: 0
HwIfOutEcnMarkedPkts: 0
```

```
HwIfInEtherOctets: 107005
HwIfOutQ0Octets: 182639
HwIfOutQ1Octets: 0
HwIfOutQ2Octets: 0
HwIfOutQ3Octets: 0
HwIfOutQ4Octets: 0
HwIfOutQ5Octets: 0
HwIfOutQ6Octets: 0
HwIfOutQ7Octets: 29739542
HwIfOutQ8Octets: 0
HwIfOutQ9Octets: 0
HwIfInPfc0Duration: 0
HwIfInPfc1Duration: 0
HwIfInPfc2Duration: 0
HwIfInPfc3Duration: 0
HwIfInPfc4Duration: 0
HwIfInPfc5Duration: 0
HwIfInPfc6Duration: 0
HwIfInPfc7Duration: 0
HwIfOutPfc0Duration: 0
HwIfOutPfc1Duration: 0
HwIfOutPfc2Duration: 0
HwIfOutPfc3Duration: 0
HwIfOutPfc4Duration: 0
HwIfOutPfc5Duration: 0
HwIfOutPfc6Duration: 0
HwIfOutPfc7Duration: 0
HwIfInQ0Pkts: 1119
HwIfInQ1Pkts: 0
HwIfInQ2Pkts: 0
HwIfInQ3Pkts: 0
HwIfInQ4Pkts: 0
HwIfInQ5Pkts: 0
HwIfInQ6Pkts: 0
HwIfInQ7Pkts: 0
HwIfInQ0BuffDiscards: 0
HwIfInQ1BuffDiscards: 0
HwIfInQ2BuffDiscards: 0
HwIfInQ3BuffDiscards: 0
HwIfInQ4BuffDiscards: 0
HwIfInQ5BuffDiscards: 0
HwIfInQ6BuffDiscards: 0
HwIfInQ7BuffDiscards: 0
```

```
HwIfInQ0SharedBuffDiscards: 0
HwIfInQ1SharedBuffDiscards: 0
HwIfInQ2SharedBuffDiscards: 0
HwIfInQ3SharedBuffDiscards: 0
HwIfInQ4SharedBuffDiscards: 0
HwIfInQ5SharedBuffDiscards: 0
HwIfInQ6SharedBuffDiscards: 0
HwIfInQ7SharedBuffDiscards: 0
HwIfOutMcBuffDiscards: 0
```

2. Get information on QOS for a port.

```
nv show interface <PORT> qos
```

Example:

```
cumulus@OCCLMLSN4600-2:mgmt:~$ nv show interface swp1 qos
                                     operational applied
-----
egress-queue-mapping
[switch-priority]      0
[switch-priority]      1
[switch-priority]      2
[switch-priority]      3
[switch-priority]      4
[switch-priority]      5
[switch-priority]      6
[switch-priority]      7
congestion-control
[traffic-class]        0
[traffic-class]        3
egress-scheduler
[traffic-class]        0
[traffic-class]        1
[traffic-class]        2
[traffic-class]        3
[traffic-class]        4
[traffic-class]        5
[traffic-class]        6
[traffic-class]        7
egress-shaper
port-max-rate          2.147e+09
port-max-shaper-state  enable
[traffic-class]         0
[traffic-class]         1
[traffic-class]         2
[traffic-class]         3
[traffic-class]         4
[traffic-class]         5
[traffic-class]         6
[traffic-class]         7
link-pause
port-buffer            150.19 KB
xoff-threshold        18.98 KB
xon-threshold         18.98 KB
```

```
rx          disable
tx          disable
mapping
port-default-sp    0
trust           13
[dscp]          0
[dscp]          1
[dscp]          2
[dscp]          3
[dscp]          4
[dscp]          5
[dscp]          6
[dscp]          7
[dscp]          8
[dscp]          9
[dscp]         10
[dscp]         11
[dscp]         12
[dscp]         13
[dscp]         14
[dscp]         15
[dscp]         16
[dscp]         17
[dscp]         18
[dscp]         19
[dscp]         20
[dscp]         21
[dscp]         22
[dscp]         23
[dscp]         24
[dscp]         25
[dscp]         26
[dscp]         27
[dscp]         28
[dscp]         29
[dscp]         30
[dscp]         31
[dscp]         32
[dscp]         33
[dscp]         34
[dscp]         35
[dscp]         36
[dscp]         37
[dscp]         38
[dscp]         39
[dscp]         40
[dscp]         41
[dscp]         42
[dscp]         43
[dscp]         44
[dscp]         45
```

```
[dscp]          46
[dscp]          47
[dscp]          48
[dscp]          49
[dscp]          50
[dscp]          51
[dscp]          52
[dscp]          53
[dscp]          54
[dscp]          55
[dscp]          56
[dscp]          57
[dscp]          58
[dscp]          59
[dscp]          60
[dscp]          61
[dscp]          62
[dscp]          63

pfc
  cable-length      100
  port-buffer       150.19 KB
  xoff-threshold   18.98 KB
  xon-threshold    18.98 KB
  rx                enable
  tx                enable
  switch-priority  3

remark
  [switch-priority] 0
  [switch-priority] 1
  [switch-priority] 2
  [switch-priority] 3
  [switch-priority] 4
  [switch-priority] 5
  [switch-priority] 6
  [switch-priority] 7
```

3. Get information on RoCE counter for a port.

```
nv show interface <PORT> qos roce counters
```

Example:

```
cumulus@OCCLMLS4600-2:mgmt:~$ nv show interface swp1 qos roce counters
                                         operational applied
-----
rx-stats
rx-non-roce-stats
  buffer-max-usage      0 Bytes
  buffer-usage          0 Bytes
  no-buffer-discard    0
  non-roce-bytes        0 Bytes
  non-roce-packets      0
  pg-max-usage          0 Bytes
  pg-usage              0 Bytes
```

```

rx-pfc-stats
  pause-duration          0
  pause-packets           0

rx-roce-stats
  buffer-max-usage       0 Bytes
  buffer-usage            0 Bytes
  no-buffer-discard      0
  pg-max-usage            0 Bytes
  pg-usage                0 Bytes
  roce-bytes              0 Bytes
  roce-packets             0

tx-stats
  tx-cnp-stats
    cnp-bytes              0 Bytes
    cnp-packets             0
    tc-max-usage            0 Bytes
    tc-usage                0 Bytes
    unicast-no-buffer-discard 0

  tx-ecn-stats
    ecn-marked-packets     0

  tx-non-roce-stats
    buffer-max-usage       0 Bytes
    buffer-usage            0 Bytes

  tx-pfc-stats
    pause-duration          0
    pause-packets           0

  tx-roce-stats
    buffer-max-usage       0 Bytes
    buffer-usage            0 Bytes
    roce-bytes              0 Bytes
    roce-packets             0
    tc-max-usage            0 Bytes
    tc-usage                0 Bytes
    unicast-no-buffer-discard 0

```

4. Get RoCE QoS configuration.

```
nv show qos roce
```

Example:

```

cumulus@OCCLMLSN4600-2:mgmt:~$ nv show qos roce
                                operational applied
-----
enable                      on
mode                        lossless   lossless
congestion-control
  congestion-mode        ECN
  enabled-tc              0,3
  max-threshold          1.43 MB
  min-threshold          146.48 KB
  probability             100

```

```

lldp-app-tlv
  priority      3
  protocol-id  4791
  selector      UDP

pfc
  pfc-priority 3
  rx-enabled    enabled
  tx-enabled    enabled

trust
  trust-mode   dscp

```

RoCE PCP/DSCP->SP mapping configurations

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

RoCE SP->TC mapping and ETS configurations

	switch-prio	traffic-class	scheduler-weight
-	-	-	-
0	0	0	DWRR-50%
1	1	0	DWRR-50%
2	2	0	DWRR-50%
3	3	3	DWRR-50%
4	4	0	DWRR-50%
5	5	0	DWRR-50%
6	6	6	strict-priority
7	7	0	DWRR-50%

RoCE pool config

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

```
Exception List
=====
description
-----
1 Port-Trust Config Mismatch.Expected Trust: pcp,dscp. Got dscp Trust
2 PCP mapping config invalid for switch-prio 0.Expected PCP: 0.
3 PCP mapping config invalid for switch-prio 1.Expected PCP: 1.
4 PCP mapping config invalid for switch-prio 2.Expected PCP: 2.
5 PCP mapping config invalid for switch-prio 3.Expected PCP: 3.
6 PCP mapping config invalid for switch-prio 4.Expected PCP: 4.
7 PCP mapping config invalid for switch-prio 5.Expected PCP: 5.
8 PCP mapping config invalid for switch-prio 6.Expected PCP: 6.
9 PCP mapping config invalid for switch-prio 7.Expected PCP: 7.
```

5. Get Roce QoS status for a port.

```
nv show interface <PORT> qos roce status
```

Example:

```
cumulus@OCCLMLSN4600-2:mgmt:~$ nv show interface swp1 qos roce status
              operational     applied
-----
congestion-control
  congestion-mode    ecn, absolute
  enabled-tc          0,3
  max-threshold      1.43 MB
  min-threshold      153.00 KB
pfc
  pfc-priority        3
  rx-enabled           yes
  tx-enabled           yes
trust
  trust-mode           dscp
mode
  mode                 lossless
```

RoCE PCP/DSCP->SP mapping configurations

```
=====
      pcp  dscp  switch-prio
-----
  cnp   6    48    6
  roce  3    26    3
```

RoCE SP->TC mapping and ETS configurations

```
=====
      switch-prio  traffic-class  scheduler-weight
-----
  cnp     6            6          strict priority
  roce   3            3          dwrr-50%
```

RoCE Pool Status							
name	mode	pool-id	switch-priorities	traffic-class	size	current-usage	
max-usage							
-	-	-	-	-	-	-	-
0 Bytes	lossy-default-ingress	DYNAMIC	2	0,1,2,4,5,6,7	-	25.69 MB	0 Bytes
0 Bytes	1 roce-reserved-ingress	DYNAMIC	3	3	-	25.69 MB	0 Bytes
144 Bytes	2 lossy-default-egress	DYNAMIC	13	-	0,6	25.69 MB	0 Bytes
0 Bytes	3 roce-reserved-egress	DYNAMIC	14	-	3	inf	0 Bytes

6. Get general port counters for all ports.

```
cl-netstat
```

Example:

```
cumulus@OCCLMLS4600-2:mgmt:~$ cl-netstat
```

Kernel Interface table

Iface Flg	MTU	RX _OK	RX _ERR	RX _DRP	RX _OVR	TX _OK	TX _ERR	TX _DRP	TX _OVR
lo	65536	522445	0	0	0	522445	0	0	0 LRU
eth0	1500	70878	0	0	0	44761	0	0	0 BMRU
swp1	9216	1191	0	0	0	382883	0	0	0 BMRU
swp2	9216	0	0	0	0	382896	0	0	0 BMRU
swp3	9216	0	0	0	0	382861	0	0	0 BMRU
swp4	9216	1175	0	0	0	382835	0	0	0 BMRU
swp5	9216	0	0	0	0	382876	0	0	0 BMRU
swp6	9216	1238	0	0	0	382885	0	0	0 BMRU
swp7	9216	0	0	0	0	382853	0	0	0 BMRU
swp8	9216	1193	0	0	0	382838	0	0	0 BMRU
swp9	9216	0	0	0	0	382872	0	0	0 BMRU
swp10	9216	49696	0	98	0	422745	0	0	0 BMU
swp11	9216	0	0	0	0	0	0	0	0 BMU
swp12	9216	0	0	0	0	0	0	0	0 BMU
swp13	9216	0	0	0	0	0	0	0	0 BMU
swp14	9216	0	0	0	0	0	0	0	0 BMU
swp15	9216	0	0	0	0	0	0	0	0 BMU
swp16	9216	0	0	0	0	0	0	0	0 BMU
swp17	9216	0	0	0	0	0	0	0	0 BMU
swp18	9216	0	0	0	0	0	0	0	0 BMU
swp19	9216	0	0	0	0	0	0	0	0 BMU
swp20	9216	0	0	0	0	0	0	0	0 BMU
swp21	9216	0	0	0	0	0	0	0	0 BMU
swp22	9216	0	0	0	0	0	0	0	0 BMU
swp23	9216	0	0	0	0	0	0	0	0 BMU
swp24	9216	0	0	0	0	0	0	0	0 BMU
swp25	9216	0	0	0	0	0	0	0	0 BMU
swp26	9216	0	0	0	0	0	0	0	0 BMU
swp27	9216	0	0	0	0	0	0	0	0 BMU
swp28	9216	0	0	0	0	0	0	0	0 BMU

swp29	9216	0	0	0	0	0	0	0	0	BMU
swp30	9216	0	0	0	0	0	0	0	0	BMU
swp31	9216	0	0	0	0	0	0	0	0	BMU
swp32	9216	0	0	0	0	0	0	0	0	BMU
swp33	9216	0	0	0	0	0	0	0	0	BMU
swp34	9216	0	0	0	0	0	0	0	0	BMU
swp35	9216	0	0	0	0	0	0	0	0	BMU
swp36	9216	0	0	0	0	0	0	0	0	BMU
swp37	9216	0	0	0	0	0	0	0	0	BMU
swp38	9216	0	0	0	0	0	0	0	0	BMU
swp39	9216	0	0	0	0	0	0	0	0	BMU
swp40	9216	0	0	0	0	0	0	0	0	BMU
swp41	9216	0	0	0	0	0	0	0	0	BMU
swp42	9216	0	0	0	0	0	0	0	0	BMU
swp43	9216	0	0	0	0	0	0	0	0	BMU
swp44	9216	0	0	0	0	0	0	0	0	BMU
swp45	9216	0	0	0	0	0	0	0	0	BMU
swp46	9216	0	0	0	0	0	0	0	0	BMU
swp47	9216	0	0	0	0	0	0	0	0	BMU
swp48	9216	0	0	0	0	0	0	0	0	BMU
swp49	9216	0	0	0	0	0	0	0	0	BMU
swp50	9216	0	0	0	0	0	0	0	0	BMU
swp51	9216	0	0	0	0	0	0	0	0	BMU
swp52	9216	0	0	0	0	0	0	0	0	BMU
swp53	9216	0	0	0	0	0	0	0	0	BMU
swp54	9216	0	0	0	0	0	0	0	0	BMU
swp55	9216	0	0	0	0	0	0	0	0	BMU
swp56	9216	0	0	0	0	0	0	0	0	BMU
swp57	9216	0	0	0	0	0	0	0	0	BMU
swp58	9216	0	0	0	0	0	0	0	0	BMU
swp59	9216	74625	0	9815	0	472960	0	0	0	BMRU
swp60	9216	32122	0	9936	0	430331	0	0	0	BMRU
swp61	9216	32175	0	9986	0	430283	0	0	0	BMRU
swp62	9216	32172	0	9990	0	430282	0	0	0	BMRU
swp63	9216	32176	0	9990	0	430272	0	0	0	BMRU
swp64	9216	32208	0	9969	0	430294	0	0	0	BMRU
br_default	9216	62325	0	0	0	290	0	0	0	BMRU
mgmt	65575	26178	0	0	0	0	0	2643	0	OmRU
mirror	9216	0	0	0	0	0	0	0	0	BMU
swid0_eth	9216	0	0	0	0	0	0	0	0	BMRU

7. To clear general counters on all port.

```
cl-netstat -c
```

8. To clear RoCE counters on a port.

```
:~$ nv action clear interface <PORT> qos roce counters
```