



Intel[®] Ethernet 800 Series

Linux Flow Control

Configuration Guide for RDMA Use Cases

Ethernet Products Group (EPG)

April 2021

Revision 1.0
635330-001



Revision History

Revision	Date	Comments
1.0	April 22, 2021	Initial public release.

Contents

1.0	Introduction	5
1.1	QoS/Flow Control Limitations on the 800 Series	5
2.0	Background	6
2.1	Ethernet Flow Control	6
2.2	Flow Control in RDMA Networks	6
2.3	Types of Flow Control: LFC vs. PFC	6
3.0	Link-Level Flow Control	8
3.1	LFC Setup Instructions	8
3.2	Symmetric vs. Asymmetric LFC	9
4.0	Priority Flow Control - Fundamentals	11
4.1	DCB Standards	11
4.1.1	DCB Willing vs. Non-willing Modes	11
4.1.2	Determining PFC Priority: PCP vs. DSCP	12
4.2	Assigning an Application to a Traffic Class	12
4.2.1	Mapping Steps Details	12
4.2.1.1	Step 1: Set ToS in the Application (or for All RoCev2 Traffic)	12
4.2.1.2	Step 2: ToS to Kernel Priority (sk_prio)	13
4.2.1.3	Step 3: Kernel Priority (sk_prio) to UP	14
4.2.1.4	Step 4: UP to TC	14
5.0	Priority Flow Control - Planning and Guidelines	15
5.1	Steps	15
5.1.1	Network Host and Switch Setup	15
5.1.2	Willing vs. Non-willing DCB Mode	16
5.1.3	Firmware vs. Software DCB	16
5.1.4	Separating and Prioritizing Traffic Streams	16
5.1.5	Configuring ETS: Map Priorities to TCs/Allocate Bandwidth	17
5.1.6	Configuring PFC: Set Priorities for Drop or No-drop	18
5.1.7	Run Applications with the Right Priority	18
5.2	Example Configurations	19
5.2.1	Example 1 - 800 Series-800 Series Back-to-Back - PFC with Single TC	19
5.2.2	Example 2 - Adapters Connected Through a Switch - Willing Mode on Adapters, DCB Configured on Switch	21
5.2.3	Example 3 - PFC with Multiple TCs (1 for RDMA, 1 for LAN) - No VLANs	22
5.2.4	Example 4 - PFC with Multiple TCs (1 for RDMA, 1 for LAN) - with VLANs	25
6.0	Priority Flow Control - Verification	31
6.1	Priority Counters	31
6.2	Discard Counters	32
6.2.1	LAN Packet Drops	32
6.2.2	RDMA Discards	32
6.3	tpcdump	33
7.0	Troubleshooting	34



NOTE: *This page intentionally left blank.*

1.0 Introduction

This document contains information on using Ethernet flow control on Intel® Ethernet 800 Series (800 Series) Network Adapters, with a focus on best practices for Linux RDMA traffic.

It includes:

- Background on Ethernet flow control and Data Center Bridging (DCB).
- Differences between Link-level Flow Control (LFC) and Priority Flow Control (PFC).
- Configuration steps for each type on 800 Series Linux hosts.
- Verification tips.

1.1 QoS/Flow Control Limitations on the 800 Series

- L3 DSCP-based QoS is currently not supported in software on the 800 Series. Only L2 VLAN-based QoS is supported.
- Although the 800 Series hardware supports eight Traffic Classes (TCs), the maximum supported configuration is four TCs. Only one TC can have Priority Flow Control enabled.

Number of Adapter Ports	Traffic Class Recommendation	RDMA
1, 2, or 4	Up to four TCs, with one of them enabled with PFC.	Supported
More than 4	No DCB Support	Not Supported

- In RoCEv2 mode, if no flow control is detected (either LFC or PFC), the driver automatically de-tunes. This is an intentional design to allow RoCEv2 to operate without flow control, but with lower performance.
- When the 800 Series is in firmware Link Layer Discovery Protocol (LLDP) mode, only three application priorities are supported. Software LLDP supports 32. This refers to the LLDP APP TLV - see "man lldptool-app" for more info.

2.0 Background

2.1 Ethernet Flow Control

By design, Ethernet is an unreliable protocol, with no guarantee that packets arrive at their destination correctly and in order. Instead, Ethernet relies on upper-layer protocols (such as TCP) or applications to provide reliable service and error correction.

The 802.3x standard introduced flow control to the Ethernet protocol, defining a mechanism for throttling the flow of data between two directly connected full-duplex network devices. If the sender transmits data faster than the receiver can accept it, the overwhelmed receiver can send a pause signal (Xoff, or “transmit off”) to the sender, requesting that the sender stop transmitting data for a specified period of time. The sender resumes transmission either after the timeout period expires or if the receiver indicates that it is ready to accept more data by sending an Xon (“transmit on”) signal.

Without flow control, data might be lost or need to be re-transmitted by a ULP or application, which can significantly affect performance.

2.2 Flow Control in RDMA Networks

The 800 Series supports both iWARP and RoCEv2 RDMA transports. Flow control is strongly recommended for RoCEv2, but iWARP also benefits.

Base Transport	Flow Control Requirements
iWARP TCP	<ul style="list-style-type: none"> iWARP runs over TCP, a reliable protocol that implements its own flow control. TCP's flow control might be relatively slow to respond in a high-performance, low-latency RDMA environment, especially under bursty traffic patterns. Ethernet flow control is optional, but can be beneficial for iWARP.
RoCEv2 UDP	<ul style="list-style-type: none"> RoCEv2 runs over UDP, an unreliable protocol with no built-in flow control. RoCEv2 therefore requires a lossless Ethernet network to ensure packet delivery. If the <i>irdma</i> driver is in RoCEv2 mode and detects no flow control, it automatically de-tunes, causing lower performance. Flow control is always recommended for RoCEv2.

2.3 Types of Flow Control: LFC vs. PFC

Ethernet standards define two types of flow control:

- Link-level Flow Control (LFC)
- Priority Flow Control (PFC)

Both types use Xon/Xoff pause frames to control data transmission. The primary difference is that LFC pauses all traffic on a link, but PFC supports Quality-of-Service (QoS) by defining different traffic priorities that can be individually paused. PFC therefore offers greater flexibility when running multiple traffic streams

Note: Despite LFC being called “link-level” flow control, both LFC and PFC operate at the link level (OSI Layer 2).

Table 1. LFC vs. PFC Comparison

	LFC	PFC
Standard	IEEE 802.3x (1997)	IEEE 802.1Qbb (2011)
Pause Type	Global pause - pauses the entire link, affecting all traffic on that link. If a link carries multiple traffic streams, a high-flow stream can cause the link to pause, thereby blocking ALL streams.	Priority pause - defines eight priorities that can be individually paused. High-bandwidth applications can be paused while allowing low-bandwidth applications to continue running.
Traffic Shaping	None.	Supports traffic classes, priorities, bandwidth allocation, and other QoS features.
Ease of Setup	Straightforward. Turn on Tx/Rx flow control on both the adapter and switch.	More complicated. Priorities, traffic classes, bandwidth allocations, and willing/non-willing mode must be configured on the adapter, switch, or both.

PFC and LFC are mutually exclusive. Only one type at a time can be enabled on a device.

- PFC is generally recommended. It has greater flexibility to handle multiple traffic streams and enhanced QoS capabilities.
- LFC can be used in limited testing situations, when only a single traffic stream is active at any given time.

3.0 Link-Level Flow Control

3.1 LFC Setup Instructions

Configuring LFC on an 800 Series network is relatively straightforward; enable flow control in both directions (Tx and Rx) on both sides of the link.

- If your hosts are connected through a switch, you must also enable flow control on the switch ports.
- If your hosts are connected back-to-back, enable LFC on both adapters.

The examples below use **eth0** as the 800 Series net device name (use **ip a** to find the device name on your system).

Switch settings vary by manufacturer. In your switch manual, look for syntax containing words like: flowcontrol, flow-control, tx-pause, or rx-pause.

To enable LFC on your network:

1. Disable firmware-based DCB on the adapter.

```
# ethtool --set-priv-flags <interface> fw-lldp-agent off
```

2. Verify that firmware DCB is disabled.

```
# ethtool --show-priv-flags <interface> | grep fw-lldp-agent
fw-lldp-agent           : off
```

3. Ensure that **lldpad** is not running.

```
# ps -ef | grep lldpad
```

4. Disable PFC on your switch, if applicable (show PFC status per port).

```
switch>show priority-flow-control
```

For example, disable PFC on a given port (like port 31/1) on an Arista 7060CX:

```
switch>enable
switch#configure
switch(config)#interface Ethernet 31/1
switch(config-if-Et31/1)#no priority-flow-control
```

Note: Some switches allow for a range of ports to be specified, for example 1/1-32/1.

5. Enable link-level flow control on the adapter.

```
# ethtool -A eth0 rx on tx on
```

6. Check LFC settings on the adapter.

```
# ethtool -a eth0
Pause parameters for eth0:
Autonegotiate: on
RX: on
TX: on
RX negotiated: on
TX negotiated: on
```


7. Enable flow control on the switch ports.

For example, enable Rx and Tx flow control on switch port 21 on an Arista 7060CX:

```
switch>enable
switch#configure
switch(config)#interface Ethernet 21/1
switch(config-if-Et31/1)#flowcontrol receive on
switch(config-if-Et31/1)#flowcontrol send on
```

8. Check LFC settings on the switch.

For example, show flow control settings on ports 21-22 on an Arista 7060CX:

```
Switch(config-if-Et31/1)#show interfaces ethernet 21/1-22/1 flowcontrol
Port Send FlowControl Receive FlowControl RxPause TxPause
admin oper admin oper
-----
Et21/1 on on off off 170373384 0
Et22/1 on on off off 289143370 0
```

3.2 Symmetric vs. Asymmetric LFC

LFC operates in both the Tx (send) and Rx (receive) directions.

- Tx flow control means that the port generates and sends Ethernet pause frames as needed.
- Rx flow control means that the port accepts and responds to Ethernet pause frames received from the connected peer.

When using LFC on the 800 Series, Intel recommends enabling both Tx and Rx flow control on both sides of the link. Also, configuring asymmetric settings (different Tx or Rx settings on each side) might have non-intuitive results.

For expected behavior, see the pause resolution table of IEEE 802.3bz shown in [Figure 1](#).

From the IEEE Standard: "The PAUSE bit indicates that the device is capable of providing the symmetric PAUSE functions as defined in Annex 31B. The ASM_DIR bit indicates that asymmetric PAUSE operation is supported. The value of the PAUSE bit when the ASM_DIR bit is set indicates the direction PAUSE frames are supported for flow across the link."

Local device		Link partner		Local device resolution	Link partner resolution
PAUSE	ASM_DIR	PAUSE	ASM_DIR		
0	0	Don't care	Don't care	Disable PAUSE Transmit and Receive	Disable PAUSE Transmit and Receive
0	1	0	Don't care	Disable PAUSE Transmit and Receive	Disable PAUSE Transmit and Receive
0	1	1	0	Disable PAUSE Transmit and Receive	Disable PAUSE Transmit and Receive
0	1	1	1	Enable PAUSE transmit Disable PAUSE receive	Enable PAUSE receive Disable PAUSE transmit
1	0	0	Don't care	Disable PAUSE Transmit and Receive	Disable PAUSE Transmit and Receive
1	Don't care	1	Don't care	Enable PAUSE Transmit and Receive	Enable PAUSE Transmit and Receive
1	1	0	0	Disable PAUSE Transmit and Receive	Disable PAUSE Transmit and Receive
1	1	0	1	Enable PAUSE receive Disable PAUSE transmit	Enable PAUSE transmit Disable PAUSE receive

Figure 1. Pause Resolution Table

Note: Some switches might not support Tx flow control. In this case, enable Rx-only flow control on the switch, and either Tx/Rx or Tx-only on the adapter.

4.0 Priority Flow Control - Fundamentals

PFC is defined by IEEE Standard 802.1Qbb and is part of the DCB suite of enhancements designed to make Ethernet a more viable, competitive transport in compute and storage environments.

The following sections provide a brief overview of the DCB standards and the role of PFC.

4.1 DCB Standards

The goal of DCB is to create a completely loss-less Ethernet network that supports bandwidth allocation across links. The features of DCB are applicable to any high-performance Ethernet environment and have significant benefits for both LAN and RDMA traffic.

Several different parts work together to make this happen:

- **PFC:** IEEE 802.1Qbb — Defines eight different traffic priorities that can be paused independently.
- **Enhanced Transmission Selection (ETS):** IEEE 802.1Qaz — Assigns bandwidth percentages to each priority.
- **Congestion Notification:** IEEE 802.1Qau — End-to-end congestion management, further avoiding frame loss.
- **Data Center Bridging Capabilities Exchange Protocol (DCBX):** IEEE 802.1az (same standard as ETS) — Discover and exchange DCB capabilities between link neighbors. Based on functionality provided by Link Layer Discovery Protocol (LLDP) (IEEE 802.1AB).

This document focuses on the PFC and ETS components.

4.1.1 DCB Willing vs. Non-willing Modes

DCB standards have a concept of willing vs. non-willing. This refers to whether the device is willing to receive its DCB settings from its link neighbor.

- In willing mode, a DCB-enabled device can query its neighbor's DCB settings, then apply the same settings to itself.
- In non-willing mode, DCB settings on the device must be explicitly configured.

A common strategy for using willing and non-willing modes in a cluster:

1. Set switches as non-willing.
2. Configure DCB (priority settings, traffic classes, bandwidth allocations, etc.) on the switch ports.
3. Set adapters as willing.
4. Adapters are automatically configured.

This helps simplify DCB cluster configuration by centralizing DCB settings on a switch and pushing the configuration to the adapters (rather than configuring each host individually).

4.1.2 Determining PFC Priority: PCP vs. DSCP

An Ethernet frame's priority can be determined by one of two distinct values: PCP (VLAN) or DSCP.

Value	Reference	Layer	Field Description
PCP	IEEE 802.1Qbb	2	Priority determined by the 3-bit 802.1p Priority Code Point (PCP) field in a frame's VLAN tag. Also sometimes called Class of Service (CoS).
DSCP	RFC 4594	3	Priority determined by the 6-bit Differentiated Services Code Point (DSCP) value in the IPv4 or IPv6 header. DSCP is the upper 6 bits of the Type of Service (ToS) field.

Ethernet devices might choose to use either value when making QoS priority decisions. This setting is usually referred to as trust mode, with options like CoS Trust, DSCP Trust, or Untrusted.

The remainder of this section discusses 802.1Qbb L2 VLAN-based PFC.

Note: Contrary to what its name might suggest, VLAN-based PFC can be used without explicitly creating new VLANs. Instead, when PFC is enabled, "VLAN 0 priority tagging" is automatically and transparently enabled on otherwise untagged interfaces. This creates a VLAN header so that the PCP field exists, but the VLAN ID is 0. Also, the VLAN priority value is influenced by setting the ToS value at the application level.

4.2 Assigning an Application to a Traffic Class

In Linux, there is a chain of mappings that starts with setting the ToS field in the IPv4/6 header that ultimately determines traffic class:

Type of Service (ToS) → Kernel Priority (sk_prio) → User Priority (UP) → Traffic Class (TC)

QoS values, like bandwidth allocations or drop/no-drop characteristics, can then be applied to each TC. This applies whether you're using explicit VLANs or VLAN 0.

4.2.1 Mapping Steps Details

4.2.1.1 Step 1: Set ToS in the Application (or for All RoCEv2 Traffic)

1. ToS is an 8-bit field in the IPv4/6 header. It can be set in two different ways on the 800 Series:

- At the application command line. For example:

LAN applications:

```
ping -Q 16; iperf3 -S 16; netperf -Y 16
```

RDMA applications:

```
ucmatose -t 16; ib_write_bw -t 16
```

- RoCEv2 traffic using **default_roce_tos** in *configs*.

For example, set ToS 16 on device `irdma0`, port 1:

```
# mkdir /sys/kernel/config/rdma_cm/irdma0
# echo 16 > /sys/kernel/config/rdma_cm/irdma0/ports/1/default_roce_tos
```

Notes:

- Common ToS values are 0, 8, 16, and 24. These correspond with priorities 0, 2, 6, and 4, respectively. See [Section 4.2.1.2](#) for details.
- DSCP is the upper six bits of the 8-bit ToS field. For historical reasons, DSCP and ToS are often used somewhat interchangeably, but they are different. Applications tend to set the entire ToS field, as this document references.

4.2.1.2 Step 2: ToS to Kernel Priority (sk_prio)

The ToS value automatically maps to Linux kernel priorities (hard coded in Linux based on **ip_tos2prio** in *net/ipv4/route.c*).

Mappings are:

```
ToS 0 --> sk_prio 0
ToS 8 --> sk_prio 2
ToS 24 --> sk_prio 4
ToS 16 --> sk_prio 6
```

For more detail, refer to:

<http://linux-tc-notes.sourceforge.net/tc/doc/priority.txt>

Also, refer to `man tc-prio`. Although **tc-prio** is about Linux traffic control (note that Intel cannot use this for RDMA traffic), the chart with TOS and Linux Priority lists the hard-coded values.

Linux Source Notes:

Mappings are implemented in the Linux kernel and drivers using `rt_tos2priority`. For example, `prio = rt_tos2priority(tos)` (from *drivers/infiniband/core/cma.c*).

This is how 0, 2, 4, and 6 occurs.

rt_tos2priority from *include/net/route.h*:

```
extern const __u8 ip_tos2prio[16];
static inline char rt_tos2priority(u8 tos)
{
    return ip_tos2prio[IPTOS_TOS(tos)>>1];
}
```

ip_tos2prio from *net/ipv4/route.c*:

```
#define ECN_OR_COST(class) TC_PRIO_##class
const __u8 ip_tos2prio[16] =
{
    TC_PRIO_BESTEFFORT,
    ECN_OR_COST(BESTEFFORT),
    TC_PRIO_BESTEFFORT,
    ECN_OR_COST(BESTEFFORT),
    TC_PRIO_BULK,
    ECN_OR_COST(BULK),
    TC_PRIO_BULK,
    ECN_OR_COST(BULK),
    TC_PRIO_INTERACTIVE,
    ECN_OR_COST(INTERACTIVE),
    TC_PRIO_INTERACTIVE,
    ECN_OR_COST(INTERACTIVE),
    TC_PRIO_INTERACTIVE_BULK,
```

```
ECN_OR_COST (INTERACTIVE_BULK) ,
TC_PRIO_INTERACTIVE_BULK,
ECN_OR_COST (INTERACTIVE_BULK) };
```

TC_PRIO integer defines from `include/uapi/linux/pkt_sched.h`:

```
#define TC_PRIO_BESTEFFORT 0
#define TC_PRIO_FILLER 1
#define TC_PRIO_BULK 2
#define TC_PRIO_INTERACTIVE_BULK 4
#define TC_PRIO_INTERACTIVE 6
#define TC_PRIO_CONTROL 7
#define TC_PRIO_MAX 15
```

4.2.1.3 Step 3: Kernel Priority (sk_prio) to UP

In the 800 Series, this is a hard-coded 1:1 mapping.

Notes:

- The UP value is the priority referenced in the PFC.
- Competing technologies might use **mqprio qdisc** (see `man tc-mqprio`) to adjust this mapping, but Intel's implementation does not.
- The 800 Series ADQ feature uses **mqprio** to direct traffic. ADQ and PFC cannot be used at the same time.

4.2.1.4 Step 4: UP to TC

The ETS aspect of DCB is used to configure this mapping.

ETS has a **up2tc** option, configurable on Linux using the **lldptool** utility. For example, this command assigns packets marked with `prio=2` to TC1, packets with `prio=0` to TC2, and all other packet priorities to TC0.

```
# lldptool -T -i eth0 -V ETS-CFG up2tc=0:2,1:0,2:1,3:0,4:0,5:0,6:0,7:0
```

Linux traffic shaping utilities like **tc**, **cgroups**, or **egress-qos-map** do not work for RDMA applications because RDMA applications bypass the kernel. If running a combination of LAN and RDMA traffic, you can still use Linux traffic shaping for LAN traffic, but you must use the PCP-based mappings described in this section for RDMA traffic.

5.0 Priority Flow Control - Planning and Guidelines

This section covers planning, considerations, and general configuration guidelines for enabling PFC on a network.

5.1 Steps

The steps for enabling PFC on your network include the following:

1. Set up your network hosts and switches. (Section 5.1.1)
2. Decide whether to use willing or non-willing DCB mode on the 800 Series adapter. (Section 5.1.2)
3. Choose firmware or software DCB. (Section 5.1.3)
4. Decide how to separate and prioritize traffic streams. (Section 5.1.4)
5. Configure ETS: Map priorities to traffic classes and allocate bandwidth. (Section 5.1.5)
6. Configure PFC: Set priorities for drop or no-drop. (Section 5.1.6)
7. Run your application with the right priority. (Section 5.1.7)

5.1.1 Network Host and Switch Setup

Note: PFC can be used with or without a switch in the network.

- If using a switch, you must configure PFC on both the adapter ports and the switch ports. Consult the appropriate switch manual for command syntax.
- If using adapters back-to-back, configure PFC on both hosts.

Host prerequisites for RDMA are outside the scope of this guide, but in general, you need at a minimum:

- Two Linux hosts with 800 Series adapters.
- Supporting 800 Series firmware and software (NVM with RDMA support, *ice* (Intel Ethernet) driver, and *irdma* driver).

If using software in DCB mode, you also need **OpenLLDP**, which includes the **lldpad** daemon and **lldptool** configuration utility.

- In RHEL, install it with **yum** (**zypper** or **apt-get** might work as well in other operating systems.):

```
# yum install lldpad
```

- Alternatively, install from source from:

<https://github.com/intel/openlldp>

5.1.2 Willing vs. Non-willing DCB Mode

DCB standards like PFC and ETS must be set to either willing or non-willing mode, which determines whether the port is willing to accept configuration settings from its link partner.

Mode	When to Use
Willing (preferred)	<ul style="list-style-type: none"> If you want to configure DCB on their switch and let adapters accept settings from the switch ports. This is the preferred, most common setup.
Non-willing	<ul style="list-style-type: none"> For back-to-back configurations. For troubleshooting, testing, and manually tweaking the configuration. If preferred, configure DCB on all hosts and set the neighboring switch ports to willing (somewhat uncommon and might not be supported by all switches).

5.1.3 Firmware vs. Software DCB

The 800 Series has two options for using DCB: firmware and software.

- Software DCB runs on the Linux host using **OpenLLDP**. It supports both willing and non-willing modes.
- Firmware DCB runs on the 800 Series adapter firmware. It only supports willing mode.

If you plan on using willing mode, firmware DCB is recommended.

Note: Only one type of DCB might be active at a time. Enabling firmware DCB overrides the software DCB setting.

DCB Type	When to Use	Willing Mode Setup	Non-willing Mode Setup
Firmware	Willing Mode	<code># ethtool --set-priv-flags <iface> fw-lldp-agent on</code>	Not supported in firmware DCB.
Software	Non-willing Mode	<pre># ethtool --set-priv-flags <iface> fw-lldp-agent off # lldptool -Ti <iface> -V PFC willing=yes # lldptool -Ti <iface> -V ETS willing=yes</pre> <p>This is a valid configuration, but firmware willing mode is recommended.</p>	<pre># ethtool --set-priv-flags <iface> fw-lldp-agent off # lldptool -Ti <iface> -V PFC willing=no # lldptool -Ti <iface> -V ETS willing=no</pre>

5.1.4 Separating and Prioritizing Traffic Streams

For networks carrying multiple traffic types, you typically want:

- One loss-less (no-drop) TC for RDMA traffic.
- One or more lossy (drop) TCs for LAN traffic.

This can change depending on specific applications.

Example configuration:

Traffic Stream	Loss-less	TC	Priority	Bandwidth
RDMA Traffic	Yes	0	0	50%
LAN Application #1	No	1	2	25%
LAN Application #2	No	2	4	25%
Unused	No	Any ¹	All Others ¹	None

1. Unused priorities can be mapped to any TC (no traffic is being steered to specific priorities). Leaving them mapped to TC 0 is acceptable.

Notes:

- The 800 Series supports a maximum of four TCs, one of which has PFC enabled.
- Traffic classes must start at zero and be contiguous (0, 1, 2, 3, ...).
- ETS bandwidth allocations must total 100%.
- Multiple priorities might map to the same TC. For example, TC0 can contain prio=0,1,2,3,4,5,6,7. However, a given priority might map to only a single TC (like prio 0 cannot be in both TC0 and TC1).
- Linux PFC defines eight TCs, but if you are steering traffic using ToS, there are only four priorities available: 0, 2, 4, and 6, which correspond with ToS 0, 8, 24, and 16 respectively. (see [Section 4.2.1.2](#)). When planning a QoS configuration, plan for a maximum of four different ToS-based TCs.

5.1.5 Configuring ETS: Map Priorities to TCs/Allocate Bandwidth

The ETS component of DCB is responsible for mapping priorities to TCs and configuring bandwidth allocations.

Note: If your adapter is in non-willing mode, use `lldptool -V ETS-CFG`. Otherwise, configure these mappings on your switch.

Mode	When to Use
Non-willing	<p>Use ETS-CFG to set priority mappings (up2tc), transmission algorithm (tsa), and bandwidth (tcbw) for each traffic class. Continuing the previous example above, set mappings with:</p> <pre># lldptool -Ti <interface> -V ETS-CFG willing=no \ up2tc=0:0,1:0,2:1,3:0,4:2,5:0,6:0,7:0 \ tsa=0:ets,1:ets,2:ets,3:strict,4:strict,5:strict,6:strict,7:strict \ tcbw=50,25,25,0,0,0,0,0</pre> <p>Notes:</p> <ul style="list-style-type: none"> • If setting tcbw for a particular TC, also set tsa=ets for that TC. • Configure the up2tc, tsa and tcbw options for ETS together. Although it's valid syntax to do them individually, the resulting configuration might not be valid. • The other priorities (UPs 1, 3, 5, 6, and 7), which are unused in this example, also map to TC0. Although they are unused, they need to map to somewhere; TC0 is a conventional choice.
Willing	Consult the appropriate switch manual for QoS configurations.

5.1.6 Configuring PFC: Set Priorities for Drop or No-drop

The PFC component of DCB enables drop or no-drop settings for each priority.

- When PFC is enabled, the priority is considered no-drop, or loss-less.
- When PFC is disabled, the priority is considered drop, or lossy.

If an adapter is in non-willing mode, use **lldptool-V PFC**. Otherwise, configure these mappings on the switch.

Mode	When to Use
Non-willing	Enable PFC on your lossless priorities (priority 0 in this example): <pre># lldptool -Ti <interface> -v PFC willing=no enabled=0</pre> Note on syntax: <ul style="list-style-type: none"> • enabled=0 means that PFC is enabled on Priority 0. • enabled=0,1,2,3,4,5,6,7 enables PFC on all priorities. • enabled=none disables PFC on all priorities
Willing	Consult the appropriate switch manual for QoS configurations.

5.1.7 Run Applications with the Right Priority

Set the ToS value in the application to steer it to the right traffic class.

Note: Command line options differ based on application.

Alternatively, for RoCEv2, you can set ToS for all traffic using the **default_roce_tos** parameter in *configs*.

Following are some examples for setting ToS=24 (corresponding to prio 4):

LAN applications:

```
# ping -Q 24
# iperf3 -S 24
# netperf -Y 24
```

RDMA applications:

```
# ucmatose -t 24
# ib_write_bw -t 24
```

Set ToS for all RoCEv2 traffic: (device `irdma0`, port 1):

```
# mkdir /sys/kernel/config/rdma_cm/irdma0
# echo 24 > /sys/kernel/config/rdma_cm/irdma0/ports/1/default_roce_tos
```

5.2 Example Configurations

5.2.1 Example 1 - 800 Series-800 Series Back-to-Back – PFC with Single TC

A common benchmarking setup uses 800 Series adapters back-to-back, running a single traffic type.

Note: When using a single TC, that TC is always TC0. Application traffic runs on TC0 by default, unless explicitly configured otherwise.

Settings in this example:

- Non-willing mode — Configure both hosts in the same way to be compatible.
- Software DCB — Required to use non-willing mode.
- All priorities mapped to TC0.
- 100% bandwidth allocated to TC0.
- PFC enabled on priority 0 — in this configuration, enabling PFC on prio 0 essentially enables it for all traffic.

Perform the following steps on both servers:

1. Disable LFC (LFC and PFC cannot co-exist).

```
# ethtool -A <interface> rx off tx off
```

2. Verify that LFC is disabled.

```
# ethtool -a <interface>
Pause parameters for <interface>:
Autonegotiate: on
RX:            off
TX:            off
RX negotiated: off
TX negotiated: off
```

3. Configure the adapter for software DCB mode by disabling firmware DCB mode.

```
# ethtool --set-priv-flags <interface> fw-lldp-agent off
```

4. Verify that firmware DCB is disabled.

```
# ethtool --show-priv-flags <interface> | grep fw-lldp-agent
fw-lldp-agent          : off
```

5. Install **OpenLLDP** (the software that controls PFC and other DCB settings), if not already installed.

RHEL:

```
# yum install lldpad
```

SLES or Ubuntu:

```
zypper or apt-get might work (untested)
```

All operating systems:

Download and build from source from <https://github.com/intel/openlldp>.

6. Start the LLDP daemon.

```
# lldpad -d
```

7. Verify LLDP is active by showing current LLDP settings on the interface.

These are the default **lspad** settings (some outputs might be different).

```
Chassis ID TLV
  MAC: 68:05:ca:a3:89:78
Port ID TLV
  MAC: 68:05:ca:a3:89:78
Time to Live TLV
  120
IEEE 8021QAZ ETS Configuration TLV
  Willing: yes
  CBS: not supported
  MAX_TCS: 8
  PRIO_MAP: 0:0 1:0 2:0 3:0 4:0 5:0 6:0 7:0
  TC Bandwidth: 0% 0% 0% 0% 0% 0% 0% 0%
  TSA_MAP: 0:strict 1:strict 2:strict 3:strict 4:strict 5:strict 6:strict 7:strict
IEEE 8021QAZ PFC TLV
  Willing: yes
  MACsec Bypass Capable: no
  PFC capable traffic classes: 8
  PFC enabled: none
End of LLDPDU TLV
```

8. Enable PFC on priority 0 in non-willing mode.

Note that `enabled=0` means that PFC is enabled on priority 0, not that PFC is not enabled.

```
# lldptool -Ti <interface> -V PFC willing=no enabled=0
```

Output:

```
willing = no
prio = 0
```

9. Enable ETS to map all priorities to TC0 and allocate 100% bandwidth to TC0.

Note: The following is a single long command line:

```
# lldptool -Ti <interface> -V ETS-CFG willing=no \
up2tc=0:0,1:0,2:0,3:0,4:0,5:0,6:0,7:0 \
tsa=0:ets,1:strict,2:strict,3:strict,4:strict,5:strict,6:strict,7:strict \
tcbw=100,0,0,0,0,0,0,0
```

Output:

```
willing = no
up2tc = 0:0,1:0,2:0,3:0,4:0,5:0,6:0,7:0
TSA = 0:ets 1:strict 2:strict 3:strict 4:strict 5:strict 6:strict 7:strict
tcbw = 100% 0% 0% 0% 0% 0% 0% 0%
```

10. Verify new settings.

```
# lldptool -ti <interface>
Chassis ID TLV
  MAC: 68:05:ca:a3:89:78
Port ID TLV
  MAC: 68:05:ca:a3:89:78
Time to Live TLV
  120
IEEE 8021QAZ ETS Configuration TLV
  Willing: no
  CBS: not supported
  MAX_TCS: 8
  PRIO_MAP: 0:0 1:0 2:0 3:0 4:0 5:0 6:0 7:0
  TC Bandwidth: 100% 0% 0% 0% 0% 0% 0% 0%
```

```
TSA_MAP: 0:ets 1:strict 2:strict 3:strict 4:strict 5:strict 6:strict 7:strict
IEEE 8021QAZ PFC TLV
Willing: no
MACsec Bypass Capable: no
PFC capable traffic classes: 8
PFC enabled: 0
End of LLDPDU TLV
```

11. Repeat on other host.

Alternatively, you could configure the other host for willing mode.

12. Run the application.

Run the application normally on the parent interface.

You do not need to specify any ToS or other priority values. The application runs on TC 0, with PFC enabled.

5.2.2 Example 2 - Adapters Connected Through a Switch – Willing Mode on Adapters, DCB Configured on Switch

A common DCB strategy in a cluster is to configure DCB on the switches, then configure the adapters for willing mode. When the willing adapters are connected to a switch with DCB configured, the adapters automatically apply the same DCB configuration.

Note: This example shows how to enable firmware willing mode on an 800 Series adapter. Consult the appropriate switch manual for DCB configuration steps.

Settings in this example:

- Willing mode — Adapters use DCB settings from the link partner.
- Firmware DCB — Recommended for willing mode.

Perform the following steps on all servers:

1. Disable LFC (LFC and PFC cannot co-exist).

```
# ethtool -A <interface> rx off tx off
```

2. Verify that LFC is disabled.

```
# ethtool -a <interface>
Pause parameters for <interface>:
Autonegotiate: on
RX:             off
TX:             off
RX negotiated: off
TX negotiated: off
```

3. Configure the adapter for firmware DCB mode (as opposed to software DCB mode).

```
# ethtool --set-priv-flags <interface> fw-lldp-agent on
```

4. Verify that firmware DCB is enabled.

```
# ethtool --show-priv-flags <interface> | grep fw-lldp-agent
fw-lldp-agent           : on
```

5. Configure DCB on the switch.

Consult the appropriate switch manual for DCB configuration. You can configure PFC, ETS, or any combination.

6. Optional. Verify that the adapter is using the correct settings.

Syntax varies by switch. In general, enable DCBX on the switch, then show the current DCBX information. The switch can then report how the attached adapter is configured.

For example, from an Arista 7060CX:

a. Enable DCBX in IEEE mode on port 21.

```
switch#enable
switch#configure
switch(config)#interface Ethernet 21/1
switch(config-if-Et21/1)#dcbx mode ieee
```

b. Show DCBX settings for port 21.

```
(config-if-Et21/1)#show dcbx Ethernet 21/1
Ethernet21/1:
IEEE DCBX is enabled and active
Last LLDPDU received on Fri Feb 14 15:42:09 2020
- PFC configuration: willing
capable of bypassing MACsec
supports PFC on up to 8 traffic classes
PFC not enabled on any priorities
- Application priority configuration:
1 application priorities configured:
ether 35078 priority 3
```

Note: This output shows a combination of the switch port's own settings and the link partner's settings. It indicates that:

- IEEE DCBX is active on the switch.
- At the given timestamp, the switch port received an LLDPDU message from the link partner describing the link partner's DCB settings.
- Everything after the Last LLDPDU received line describes the contents of the received LLDPDU. These are the adapter's settings.

5.2.3 Example 3 - PFC with Multiple TCs (1 for RDMA, 1 for LAN) – No VLANs

This example describes how to run both RDMA and LAN traffic on the same link using the parent interface (no explicit VLANs, although VLAN 0 are used transparently).

These steps can be used in a back-to-back configuration, or if you are using a switch, be sure to configure the neighboring switch ports for the same configuration (consult the appropriate switch manual for more detail).

Settings in this example:

- Non-willing mode — In this example, adapter settings are configured explicitly using **lldptool** (vs. configuring DCB on a switch and using willing mode on adapters).
- Software DCB — Required to use non-willing mode.
- Two traffic classes:
 - One loss-less TC for RDMA, with 80% bandwidth allocated.
 - One lossy TC for LAN, with 20% bandwidth allocated.

- PFC enabled for only the RDMA traffic class (this makes it loss-less).

Perform the following steps on both servers:

1. Disable LFC (LFC and PFC cannot co-exist).

```
# ethtool -A <interface> rx off tx off
```

2. Verify that LFC is disabled.

```
# ethtool -a <interface>
Pause parameters for <interface>:
Autonegotiate: on
RX:           off
TX:           off
RX negotiated: off
TX negotiated: off
```

3. Configure the adapter for software DCB mode by disabling firmware DCB mode.

```
# ethtool --set-priv-flags <interface> fw-lldp-agent off
```

4. Verify that firmware DCB is disabled:

```
# ethtool --show-priv-flags <interface> | grep fw-lldp-agent
fw-lldp-agent           : off
```

5. Install **OpenLLDP** (the software that controls PFC and other DCB settings), if not already installed:

RHEL:

```
# yum install lldpad
```

SLES or Ubuntu:

```
zypper or apt-get might work (untested)
```

All operating systems:

Download and build from source from <https://github.com/intel/openlldp>.

6. Start the LLDP daemon.

```
# lldpad -d
```

7. Verify LLDP is active by showing current LLDP settings on the interface.

The following example shows the **OpenLLDP** default:

```
# lldptool -ti <interface>
Chassis ID TLV
  MAC: 68:05:ca:a3:89:78
Port ID TLV
  MAC: 68:05:ca:a3:89:78
Time to Live TLV
  120
IEEE 8021QAZ ETS Configuration TLV
  Willing: yes
  CBS: not supported
  MAX_TCS: 8
  PRIO_MAP: 0:0 1:0 2:0 3:0 4:0 5:0 6:0 7:0
  TC Bandwidth: 0% 0% 0% 0% 0% 0% 0% 0%
  TSA_MAP: 0:strict 1:strict 2:strict 3:strict 4:strict 5:strict 6:strict 7:strict
IEEE 8021QAZ PFC TLV
  Willing: yes
  MACsec Bypass Capable: no
```

```

PFC capable traffic classes: 8
PFC enabled: none
End of LLDPDU TLV

```

8. Plan your DCB configuration.

Traffic Stream	Loss-less	TC	Priority	ToS	Bandwidth
RDMA Application	Yes	0	0	0	80%
LAN Application	No	1	4	24	20%
Unused	No	Any ¹	All Others	N/A	0%

1. Unused priorities can be mapped to any TC (no traffic is being steered to specific priorities). Leaving them mapped to TC 0 is acceptable.

9. Configure ETS.

- Map priorities to traffic classes.
- Allocate bandwidths.

Note: The following is a single long command line:

```

# lldptool -Ti <interface> -V ETS-CFG willing=no \
up2tc=0:0,1:0,2:0,3:0,4:1,5:0,6:0,7:0 \
tsa=0:ets,1:ets,2:strict,3:strict,4:strict,5:strict,6:strict,7:strict \
tcbw=80,20,0,0,0,0,0,0

```

Output:

```

willing = no
up2tc = 0:0,1:0,2:0,3:0,4:1,5:0,6:0,7:0
TSA = 0:ets 1:ets 2:strict 3:strict 4:strict 5:strict 6:strict 7:strict
tcbw = 80% 20% 0% 0% 0% 0% 0% 0%

```

10. Enable PFC on priority 0 (non-willing).

```

# lldptool -Ti <interface> -V PFC willing=no enabled=0

```

Output:

```

willing = no
prio = 0

```

11. Verify new settings.

```

# lldptool -ti <interface>
Chassis ID TLV
MAC: 68:05:ca:a3:89:78
Port ID TLV
MAC: 68:05:ca:a3:89:78
Time to Live TLV
120
IEEE 8021QAZ ETS Configuration TLV
Willing: no
CBS: not supported
MAX_TCS: 8
PRIO_MAP: 0:0 1:0 2:0 3:0 4:1 5:0 6:0 7:0
TC Bandwidth: 80% 20% 0% 0% 0% 0% 0% 0%
TSA_MAP: 0:ets 1:ets 2:strict 3:strict 4:strict 5:strict 6:strict 7:strict

```



```
IEEE 8021QAZ PFC TLV
    Willing: no
    MACsec Bypass Capable: no
    PFC capable traffic classes: 8
    PFC enabled: 0 ← This means PFC is enabled on prio 0 (not that PFC is disabled)
End of LLDPDU TLV
```

12. Repeat on neighbor node:

- If using a back-to-back configuration, either repeat the configuration on the other host or enable willing mode on that host.
- If using a switch, configure the same DCB scheme on the switch port. Consult the appropriate switch manual for details.

13. Run the application.

RDMA:

Run the RDMA application normally. Since RDMA is running on the default TC0 with prio 0 in this example, you do not need any command line options to set ToS. ToS (and therefore priority) is 0 by default.

LAN:

Run the LAN application with a command line option to set ToS=24. In Linux, this maps to prio=4.

5.2.4 Example 4 - PFC with Multiple TCs (1 for RDMA, 1 for LAN) – with VLANs

This example describes how to run both RDMA and LAN traffic on the same link using VLANs.

These steps can be used in a back-to-back configuration, or if you are using a switch, be sure to configure the neighboring switch ports for the same configuration (consult the appropriate switch manual for more detail).

Settings in this example:

- Non-willing mode — In this example, adapter settings are configured explicitly using **lldptool** (vs. configuring DCB on a switch and using willing mode on adapters).
- Software DCB — Required to use non-willing mode.
- Three traffic classes:
 - 1 lossy TC for general LAN traffic on the parent interface.
 - 1 loss-less TC for RDMA traffic on VLAN 100.
 - 1 lossy TC for LAN traffic on VLAN 200.
- PFC enabled for only the RDMA traffic class (this makes it loss-less).

Perform the following steps on both servers:

1. Disable LFC (LFC and PFC cannot co-exist).

```
# ethtool -A <interface> rx off tx off
```

2. Verify that LFC is disabled.

```
# ethtool -a <interface>
Pause parameters for <interface>:
Autonegotiate: on
RX:            off
TX:            off
RX negotiated: off
TX negotiated: off
```

3. Configure the adapter for software DCB mode by disabling firmware DCB mode.

```
# ethtool --set-priv-flags <interface> fw-lldp-agent off
```

4. Verify that firmware DCB is disabled.

```
# ethtool --show-priv-flags <interface> | grep fw-lldp-agent
fw-lldp-agent      : off
```

5. Install **OpenLLDP** (the software that controls PFC and other DCB settings), if not already installed.

RHEL:

```
# yum install lldpad
```

SLES or Ubuntu:

zypper or apt-get might work (untested)

All operating systems:

Download and build from source from <https://github.com/intel/openlldp>.

6. Start the LLDP daemon.

```
# lldpad -d
```

7. Verify LLDP is active by showing current LLDP settings on the interface.

The following example shows the **OpenLLDP** default:

```
# lldptool -ti <interface>
Chassis ID TLV
    MAC: 68:05:ca:a3:89:78
Port ID TLV
    MAC: 68:05:ca:a3:89:78
Time to Live TLV
    120
IEEE 8021QAZ ETS Configuration TLV
    Willing: yes
    CBS: not supported
    MAX_TCS: 8
    PRIO_MAP: 0:0 1:0 2:0 3:0 4:0 5:0 6:0 7:0
    TC Bandwidth: 0% 0% 0% 0% 0% 0% 0% 0%
    TSA_MAP: 0:strict 1:strict 2:strict 3:strict 4:strict 5:strict 6:strict 7:strict
IEEE 8021QAZ PFC TLV
    Willing: yes
    MACsec Bypass Capable: no
    PFC capable traffic classes: 8
    PFC enabled: none
End of LLDPDU TLV
```

8. Plan your DCB configuration

- RDMA is loss-less (PFC enabled).
- LAN traffic is lossy (PFC disabled).

Traffic Stream	Loss-less	TC	Priority	ToS	Bandwidth	Interface
General Traffic	No	0	0	0	10%	Parent
RDMA Application	Yes	1	2	8	50%	VLAN 100
LAN Application	No	2	3	0 ¹	40%	VLAN 200
Unused	No	Any ²	All Others	N/A	0%	N/A

1. LAN traffic can set VLAN priority directly using `egress-qos-map` when configuring the interface, so ToS mappings are not required. RDMA traffic cannot use `egress-qos-map` because `egress-qos-map` is controlled by the kernel, and RDMA traffic bypasses the kernel.
2. Unused priorities can be mapped to any TC (no traffic is being steered to specific priorities). Leaving them mapped to TC 0 is acceptable.

9. Configure ETS.

- Map priorities to traffic classes.
- Allocate bandwidth.

Note: The following is a single long command line:

```
# lldptool -Ti <interface> -V ETS-CFG willing=no \
up2tc=0:0,1:0,2:1,3:2,4:0,5:0,6:0,7:0 \
tsa=0:ets,1:ets,2:ets,3:strict,4:strict,5:strict,6:strict,7:strict \
tcbw=10,50,40,0,0,0,0,0
```

Output:

```
willing = no
up2tc = 0:0,1:0,2:1,3:2,4:0,5:0,6:0,7:0
TSA = 0:ets 1:ets 2:ets 3:strict 4:strict 5:strict 6:strict 7:strict
tcbw = 10% 50% 40% 0% 0% 0% 0%
```

10. Enable PFC on Priority 2 (non-willing).

```
# lldptool -Ti <interface> -V PFC willing=no enabled=2
```

Output:

```
willing = no
prio = 2
```

11. Verify new settings.

```
# lldptool -ti <interface>
Chassis ID TLV
    MAC: 68:05:ca:a3:89:78
Port ID TLV
    MAC: 68:05:ca:a3:89:78
Time to Live TLV
    120
```

```
IEEE 8021QAZ ETS Configuration TLV
    Willing: no
    CBS: not supported
    MAX_TCS: 8
    PRIO_MAP: 0:0 1:0 2:1 3:2 4:0 5:0 6:0 7:0
    TC Bandwidth: 10% 50% 40% 0% 0% 0% 0% 0%
    TSA_MAP: 0:ets 1:ets 2:ets 3:strict 4:strict 5:strict 6:strict 7:strict
IEEE 8021QAZ PFC TLV
    Willing: no
    MACsec Bypass Capable: no
    PFC capable traffic classes: 8
    PFC enabled: 0x4 ← This is a mask. 0x4 = 0b0000_0100, meaning PFC is enabled on TC 2.
End of LLDPDU TLV
```

12. Repeat DCB settings on the neighbor node:

- If using a back-to-back configuration, either repeat the DCB configuration on the other host or enable willing mode on that host.
- If using a switch, configure the same DCB scheme on the switch port. Consult the appropriate switch manual for details.

13. Create VLAN 100 for RDMA traffic:

- a. Create VLAN 100 as a part of the parent interface (like parent interface eth0, with an IP Address of 192.168.0.3).

```
# ip link add eth0.100 link eth0 type vlan id 100
```

- b. Bring the new interface up.

```
# ip link set eth0.100 up
```

- c. Set the IP Address on the new interface. In the example address, the third octet is 100, same as the VLAN ID, but the values do not need to match. However, the VLAN IP Address does need to be in a different subnet than the parent address.

```
# ip address add dev eth0.100 192.168.100.3/24
```

14. Create VLAN 200 (for LAN traffic) with `qos-egress-map` set:

- a. Create VLAN 200 as a part of the same parent interface (still using eth0 in the example).
- b. Use `egress-qos-map` to map all VLAN 200 LAN traffic to priority 3 in the VLAN header (see `man ip-link` for documentation).

Note: The `egress-qos-map` method works only for LAN traffic and not RDMA traffic.

```
# ip link add eth0.200 link eth0 type vlan id 200 egress-qos-map 0:3 1:3 2:3
3:3 4:3 5:3 6:3 7:3
```

- c. Bring the new interface up.

```
# ip link set eth0.200 up
```

- d. Set the IP Address on the new interface. In the example address, the third octet is 200, like the VLAN ID, but it does not need to match.

```
# ip address add dev eth0.200 192.168.200.3/24
```

15. Verify new interfaces:

- a. Examine the output of `ip link show` and verify both new VLANs are up and have the right IP Address.

```
10: enp175s0f0.100@enp175s0f0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500
qdisc noqueue state UP group default qlen 1000
link/ether 68:05:ca:a3:89:78 brd ff:ff:ff:ff:ff:ff
inet 192.168.100.1/24 scope global enp175s0f0.100
valid_lft forever preferred_lft forever
inet6 fe80::6a05:caff:fea3:8978/64 scope link
valid_lft forever preferred_lft forever
12: enp175s0f0.200@enp175s0f0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500
qdisc noqueue state UP group default qlen 1000
link/ether 68:05:ca:a3:89:78 brd ff:ff:ff:ff:ff:ff
inet 192.168.200.3/24 scope global enp175s0f0.200
valid_lft forever preferred_lft forever
inet6 fe80::6a05:caff:fea3:8978/64 scope link
valid_lft forever preferred_lft forever
```

- b. Verify egress mappings for VLAN 200 here.

```
# cat /proc/net/vlan/eth0.200 | grep EGRESS
EGRESS priority mappings: 0:3 1:3 2:3 3:3 4:3 5:3 6:3 7:3
```

16. Repeat VLAN settings on the neighbor node:

- If using a back-to-back configuration, configure the same VLANs on the other host.
- If using a switch, consult the appropriate switch manual for details.

Sample commands from an Arista 7060CX:

- a. Create VLANs 100 and 200.

```
switch>enable
switch>config
switch(config)>vlan 100
switch(config)>vlan 200
```

- b. Set the switch ports where adapters are connected to trunk mode. Example, for port 21/1.

```
switch(config)>#interface Et21/1
switch(config-if-Et21/1)#switchport mode trunk
```

- c. Add the VLANs to the switch ports where adapters are connected.

```
switch(config-if-Et21/1)#switchport trunk allowed vlan 1,100,200
```

- d. Show current VLANs (VLAN 1 always exists by default).

```
switch> show vlan
VLAN Name Status Ports
-----
1 default active Et21/1, Et23/1
100 VLAN0100 active Et21/1, Et23/1
200 VLAN0200 active Et21/1, Et23/1
```

Note: If needed, undo settings, preface them with "no".

- To delete a VLAN: `switch(config)>no vlan 100`
- To remove trunk mode: `switch(config-if-Et23/1)#no switchport mode trunk`

17. Run the applications.

Traffic Stream	Interface	Example IP Address	TC	Priority	ToS	Set Application Priority
General Traffic	Parent	192.168.0.1	0	0	0	Run normally on 192.168.0.1, no ToS options needed. prio 0 and TC 0 are the defaults.
RDMA Application	VLAN100	192.168.100.1	1	2	8	Run on 192.168.100.1 and set ToS=8 on the application command line. Alternatively, if using RoCEv2: Set default_roce_tos=8 (Ctrl-F this article for syntax). This sets ToS=8 for all RoCEv2 traffic, so you do not need the application command line option.
LAN Application	VLAN 200	192.168.200.1	2	3	0	Run on 192.168.200.1 normally. No command line ToS options needed because egress-qos-map is set to use priority 3.

6.0 Priority Flow Control - Verification

6.1 Priority Counters

Priority flow control counters for each interface are available in **ethtool**. They measure the number of Xon and Xoff (transmit on and off) frames sent and received by that interface.

To view them:

```
# ethtool -S <interface> | grep prio
```

Counters are named with tx/rx, priority number, and either xon or xoff.

For example:

```
# ethtool -S enp175s0f0 | grep prio
tx_priority_0_xon.nic: 1
tx_priority_0_xoff.nic: 6434
tx_priority_1_xon.nic: 1
tx_priority_1_xoff.nic: 6434
tx_priority_2_xon.nic: 2
tx_priority_2_xoff.nic: 14864
tx_priority_3_xon.nic: 1
tx_priority_3_xoff.nic: 6434
tx_priority_4_xon.nic: 0
tx_priority_4_xoff.nic: 0
tx_priority_5_xon.nic: 1
tx_priority_5_xoff.nic: 6434
tx_priority_6_xon.nic: 1
tx_priority_6_xoff.nic: 6434
tx_priority_7_xon.nic: 1
tx_priority_7_xoff.nic: 6434
rx_priority_0_xon.nic: 0
rx_priority_0_xoff.nic: 0
rx_priority_1_xon.nic: 0
rx_priority_1_xoff.nic: 0
rx_priority_2_xon.nic: 0
rx_priority_2_xoff.nic: 0
rx_priority_3_xon.nic: 0
rx_priority_3_xoff.nic: 0
rx_priority_4_xon.nic: 0
rx_priority_4_xoff.nic: 0
rx_priority_5_xon.nic: 0
rx_priority_5_xoff.nic: 0
rx_priority_6_xon.nic: 0
rx_priority_6_xoff.nic: 0
rx_priority_7_xon.nic: 0
rx_priority_7_xoff.nic: 0
```

Note that the Rx counters all 0.

When adapters are connected through a switch, the `rx_priority_*` counters might be 0, indicating that the adapter has not received any pause frames from the switch. Depending on the level of stress in the network, this is acceptable if the switch has enough buffering to keep up with the host demand. However, for high stress traffic such as HPC applications at larger scale, often the switch sends pause frames to the host. In general, it is expected to see both tx and rx_priority counters.

Note that some of the Tx counters have the same value.

In the 800 Series QoS implementation, if PFC is enabled for any priority in a traffic class, all priorities in that traffic class get pause frames. This means that the counters for all priorities in the same TC are incremented in unison, regardless of the particular single priority that is causing PFC to trigger. If all priorities are mapped to the same TC, they all increment in unison.

This implementation is in line with 802.1Q recommendations.

- 802.1Q Section 37.3: NOTE 2 — All priorities within a traffic class typically share similar traffic handling requirements (e.g., loss and bandwidth).
- 802.1Q Section 8.6.8: NOTE 1 — Two or more priorities can be combined in a single queue. In this case if one or more of the priorities in the queue are paused, it is possible for frames in that queue not belonging to the paused priority to not be scheduled for transmission.
- 802.1Q Section 8.6.8: NOTE 2 — Mixing PFC and non-PFC priorities in the same queue results in non-PFC traffic being paused causing congestion spreading, and therefore is not recommended.

Tip: Run a **watch** command in a separate terminal window to see priority counters moving in real time:

```
# watch -d -n 1 "ethtool -S <interface> | grep prio"
```

6.2 Discard Counters

Enabling flow control should eliminate drops and discards in the network.

6.2.1 LAN Packet Drops

```
# ethtool -S enp175s0f0 | grep drop
rx_dropped: 0
tx_dropped_link_down.nic: 0
rx_dropped.nic: 0
```

6.2.2 RDMA Discards

```
# cd /sys/class/infiniband/irdma-enp175s0f0/hw_counters
# for f in *Discards; do echo -n "$f: "; cat "$f"; done
ip4InDiscards: 0
ip6InDiscards: 0
```

If you see non-zero counters for packet discards, double check that both tx and rx_priority counters are being sent and received by the NIC as described in [Section 6.1](#). If any are zero, PFC might not be fully enabled.

6.3 tpcdump

tcpdump can be used to verify ToS, DSCP, PCP, or VLAN ID values.

RDMA traffic requires switch port mirroring or an inline protocol analyzer to capture RDMA traffic.

To find each value in tcpdump:

Value	tcpdump Option	tcpdump Output
ToS	-v	<p>Run ping with -Q 24 to set ToS=24 (0x18):</p> <pre># ping -I enp175s0f0 -Q 24 192.168.0.3</pre> <p>Use tcpdump with -v and look for tos 0x18 in the IP header:</p> <pre># tcpdump -nXX -v -i enp175s0f0 15:05:53.197406 IP (tos 0x18, ttl 64, id 20320, offset 0, flags [DF], proto ICMP (1), length 84) 192.168.0.1 > 192.168.0.3: ICMP echo request, id 14718, seq 3, length 64 0x0000: 6805 caa3 8778 6805 caa3 8978 0800 4518 h...xh...x...E. 0x0010: 0054 4f60 4000 4001 69dc c0a8 0001 c0a8 .TO`@.@.i..... 0x0020: 0003 0800 d918 397e 0003 312f 585e 00009~..1/X^.. 0x0030: 0000 9a05 0300 0000 0000 1011 1213 1415 0x0040: 1617 1819 1a1b 1c1d 1e1f 2021 2223 2425!"#\$\$% 0x0050: 2627 2829 2a2b 2c2d 2e2f 3031 3233 3435 &'()*+,-./012345 0x0060: 3637</pre>
DSCP	-v	<p>DSCP is not shown explicitly in tcpdump. DSCP is based on ToS value. Example: ToS 0x18 maps to DSCP 0x6.</p> <ol style="list-style-type: none"> 1. Start with ToS 0x18. 2. Convert 0x18 from hex to decimal. 0x18 = 0b0001_1000 3. Look at the upper 6 bits of the ToS field to find DSCP. 0x18 = 0b00011000 0b000110 = 0x06 4. DSCP = 0x6
VLAN ID and VLAN priority (PCP)	-e vlan	<p>Run ping on VLAN 200 with -Q 24 to set ToS=24 (0x18):</p> <pre># ping -I enp175s0f0.200 -Q 24 192.168.200.3</pre> <p>Run tcpdump with -e vlan and look for vlan 200 and p 4 in the VLAN header:</p> <pre># tcpdump -nXX -v -i enp175s0f0 -e vlan 15:23:37.554911 68:05:ca:a3:89:78 > 68:05:ca:a3:87:78, ethertype 802.1Q (0x8100), length 102: vlan 200, p 4, ethertype IPv4, (tos 0x18, ttl 64, id 53320, offset 0, flags [DF], proto ICMP (1), length 84) 192.168.200.1 > 192.168.200.3: ICMP echo request, id 15739, seq 3, length 64 0x0000: 6805 caa3 8778 6805 caa3 8978 8100 80c8 h...xh...x... # PCP is 3 bits: 0x8 -> 0b1000 -> 4 0x0010: 0800 4518 0054 d048 4000 4001 58ea c0a8 ..E..T.H@.@.X... 0x0020: c801 c0a8 c803 0800 82a4 3d7b 0003 5933={..Y3 0x0030: 585e 0000 0000 bf78 0800 0000 0000 1011 X^.....x..... 0x0040: 1213 1415 1617 1819 1a1b 1c1d 1e1f 2021! 0x0050: 2223 2425 2627 2829 2a2b 2c2d 2e2f 3031 "#\$%&'()*+,-./01 0x0060: 3233 3435 3637</pre>

7.0 Troubleshooting

Problem	Solution
"Agent instance for device not found 00007f" when running lldptool .	"ifup" your netdev interface.
"Failed to connect to lldpad - clif_open: Connection refused" when running lldptool .	Make sure lldpad is running. <pre>ps -ef grep lldpad)</pre>
"Invalid command 00007f" when running lldptool	Check your command syntax and arguments. Possible missing argument or bad values.
"kernel: ice 0000:af:00.0: Set DCB Config failed" in the system log or dmesg .	Likely an invalid DCB configuration: <ul style="list-style-type: none"> • Verify in ETS that TCs are contiguous and start at TC0. • Verify in ETS that bandwidth allocations total 100%. • Verify that you are not trying to allocate bandwidth to TCs that do not exist in up2tc. If necessary, reset the OpenLLDP configuration to the default by removing the configuration file. <pre># killall lldpad && rm /var/lib/lldpad/lldpad.conf</pre> When you restart lldpad , it regenerated the default configuration file.
"[4353.872101] ice 0000:af:00.0: application TOS[0] and remote client TOS[24] mismatch" in the system log or dmesg .	This happens because an RDMA application with, for example, 16 QP with ToS 24 might still create a 17th QP for application setup using ToS 0. This is a quirk of the application and not controlled by the 800 Series <i>ice</i> or <i>irdma</i> drivers.
lldpad or lldptool problems.	Enable verbose lldpad debug logging: <ol style="list-style-type: none"> 1. Edit <code>/usr/lib/systemd/system/lldpad.service</code>. <pre>ExecStart=/usr/sbin/lldpad -t -V9</pre> 2. Save and reload lldpad service. <pre>systemctl daemon-reload && systemctl restart lldpad</pre> 3. Save log to a file. <pre>journalctl -u lldpad -f tee lldpad.log</pre>

NOTE: *This page intentionally left blank.*



LEGAL

No license (express or implied, by estoppel or otherwise) to any intellectual property rights is granted by this document.

This document (and any related software) is Intel copyrighted material, and your use is governed by the express license under which it is provided to you. Unless the license provides otherwise, you may not use, modify, copy, publish, distribute, disclose or transmit this document (and related materials) without Intel's prior written permission. This document (and related materials) is provided as is, with no express or implied warranties, other than those that are expressly stated in the license.

Intel disclaims all express and implied warranties, including without limitation, the implied warranties of merchantability, fitness for a particular purpose, and non-infringement, as well as any warranty arising from course of performance, course of dealing, or usage in trade.

This document contains information on products, services and/or processes in development. All information provided here is subject to change without notice. Contact your Intel representative to obtain the latest forecast, schedule, specifications and roadmaps.

The products and services described may contain defects or errors which may cause deviations from published specifications.

Copies of documents which have an order number and are referenced in this document may be obtained by calling 1-800-548-4725 or by visiting www.intel.com/design/literature.htm.

Intel and the Intel logo are trademarks of Intel Corporation in the U.S. and/or other countries.

Other names and brands may be claimed as the property of others.

© 2021 Intel Corporation.