

# Enhanced Transmission Selection (ETS)

## Feature Overview and Configuration Guide

### Introduction

This guide describes Enhanced Transmission Selection (ETS) and how to configure it.

Enhanced Transmission Selection (ETS) is a traffic management feature defined in the IEEE 802.1Qaz standard. It helps network devices like switches and routers manage how different types of traffic share bandwidth on a single physical link. ETS works alongside other standards like Data Center Bridging (DCB) to improve performance in environments where multiple traffic classes, such as storage, voice, and video, compete for resources.

At its core, ETS allows administrators to assign minimum bandwidth guarantees to different traffic classes. This ensures that critical applications always get the resources they need, even when the network is busy. ETS also supports bandwidth sharing when some traffic classes are idle, helping to make the most of available capacity.

### Contents

Introduction .....	1
Products and software version that apply to this guide .....	2
Enhanced Transmission Selection (ETS) .....	3
Scheduler .....	3
Enhanced transmission selection (ETS) .....	4
Data Center Bridging eXchange protocol (DCBX) .....	4
Configuration examples .....	6
Example: configuring a scheduler-set with ETS .....	6
Example: configuring DCBX on an interface .....	8
Removing a scheduler-set10 .....	



## Products and software version that apply to this guide

This guide applies to AlliedWare Plus™ products that support Enhanced Transmission Selection (ETS), running version **5.5.5-2.1** or later.

The Data Center Bridging eXchange protocol (DCBX) feature is available from version **5.5.6-0** or later.

To see whether your product supports Enhanced Transmission Selection (ETS), see the following documents:

- The product's [Datasheet](#)
- The product's [Command Reference](#)

## Enhanced Transmission Selection (ETS)

ETS builds on the concept of traffic classes, which group packets based on their type or priority. Each class can be assigned a minimum percentage of the total link bandwidth. For example, you might allocate 40% to storage traffic, 30% to voice, and 30% to general data. These percentages ensure that each class gets a fair share of the link, even during congestion.

ETS works in conjunction with Priority-based Flow Control (PFC), which is part of the IEEE 802.1Qbb standard. PFC helps pause traffic for specific classes during congestion. Together, these features create a more predictable and efficient network.

One of the key benefits of ETS is its ability to support converged networks, where multiple types of traffic share the same infrastructure. This is common in data centers, where storage, compute, and management traffic all flow through the same switches. ETS helps ensure that no single traffic type dominates the link, which improves reliability and performance.

### Scheduler

The scheduler determines how the packets in the eight egress queues of a port are processed.

#### Strict priority

By default, all queues on all ports use strict priority. This means that the highest priority queue is emptied first and when it is completely empty, the next highest priority queue is processed, and so on. Thus, for a strict priority queue to be processed, all higher priority queues must be empty. Note that queue 7 is highest priority, queue 0 is lowest.

#### Weighted round robin (WRR)

The alternative approach is Weighted Round Robin. In this case, you group some queues together and assign weights to each queue within that group. For example, you can group queues 4 and 5 together, and give queue 5 a weight of 20 and queue 4 a weight of 10. That means for the device will transmit 2 packets from queue 5 for every 1 packet from queue 4.

You can also configure multiple WRR groups. For example, you can configure two WRR groups, with all packets in group 1 being processed first, using weights you specify for the queues in group 1. If all queues in group 1 are empty, the queues in group 2 will be processed using the weights specified for the queues in group 2.

If you configure a mixture strict priority and WRR queues on the port, the higher priority queues must still be empty before lower priorities will be processed. For each WRR group, a winning queue will be selected. This will then compete with the strict priority queues, as though it was strict priority as well. For example, imagine you configure a port where queues 7, 6, 1, and 0 are strict priority, and queues 5, 4, 3, and 2 are a WRR group. The WRR group will only be processed once queues 7 and 6 are empty. And queues 1 and 0 will only be processed once the WRR group, as well as queues 7 and 6, are all empty.

## Enhanced transmission selection (ETS)

ETS is a traffic management feature defined in the IEEE 802.1Qaz standard. It helps network switches manage how different types of traffic share bandwidth on a single physical link. ETS works alongside Data Center Bridging (DCB) technologies to improve performance in data center environments.

At its core, ETS allows administrators to assign minimum bandwidth guarantees to different traffic classes as a percentage of link bandwidth. This ensures that critical applications always get the resources they need, even when the network is busy. ETS also supports bandwidth sharing when some traffic classes are idle, helping to make the most of available capacity.

On AlliedWare Plus, you can configure WRR groups using percentage-based weights. This is part of work to implement ETS in our devices.

### Percentage-based weights

To configure a WRR group with percentage-based weights, use the following command:

```
awplus(config)# mls qos scheduler-set <1-12> wrr-queue group <1-2> percent
<1-100> queues [0] [1] [2] [3] [4] [5] [6] [7]
```

## Data Center Bridging eXchange protocol (DCBX)

DCBX is used by Data Center Bridging (DCB) devices to exchange configuration information with directly connected peers. The protocol may also be used for misconfiguration detection and for configuration of the peer.

DCBX uses LLDP to exchange attributes between two links peers. When DCBX and LLDP are enabled on an interface, the following Type-Length-Value (TLV) elements will be advertised:

**Table 1-1: ETS Configuration TLV — D.2.9 of IEEE Std 802.1Q-2018**

Field	Description	AlliedWare Plus Support
Willing	Indicates if the device is willing to accept configuration from neighbors	Always 0 to indicate unwilling
Credit Based Shaper	Indicates if the device supports the Credit-based Shaper transmission selection algorithm	Always 0 to indicate unsupported
Max Traffic Classes	Indicates the maximum number of traffic classes the device supports.	Always 0 to indicate support for 8 traffic classes
Priority Assignment Table	Mapping of priority to traffic classes	Mappings as configured by <code>mls qos map cos-queue &lt;0-7&gt; to &lt;0-7&gt;</code>
Traffic Class Bandwidth Table	Indicates the current bandwidth percentage configured for each traffic class	Percentages as configured by <code>mls qos scheduler-set &lt;1-12&gt; wrr-queue group 1 percent &lt;1-100&gt; queue &lt;0-7&gt;</code>
TSA Assignment Table	Indicates the Transmission selection algorithm to be used for each traffic class	ETS (2) for traffic classes configured with <code>wrr-queue</code> and <code>percent</code> . Otherwise Strict Priority (0)

Additionally, if PFC is enabled on an interface, the following TLV elements will be advertised:

**Table 1-2: PFC TLV — D.2.11 of IEEE Std 802.1Q-2018**

Item	Default profile	Profile1
Willing	Indicates if the device is willing to accept configuration from neighbors	Always 0 to indicate unwilling
MACsec Bypass Capability	Indicates if the device is capable of bypassing MACsec processing when MACsec is disabled	Always 0 to indicate capable
PFC Capability	Indicates the maximum number of traffic classes that simultaneously support PFC on the device	Always 8
PFC Enable	Indicates if PFC is enabled for each the priority	As configured by pfc priority <0-7>

## Configuration examples

In this section, you can find examples of how to configure the features described above.

### Example: configuring a scheduler-set with ETS

In the following example, you want to configure a scheduler-set with two WRR groups, and set port1.0.1 to use that scheduler-set.

In our example, port1.0.1 will have the following scheduler-set:

- Queues 5, 6, and 7 are strict priority.
- Queues 2, 3, and 4 are WRR group 1, with percentages of 50%, 25% and 25% respectively.
- Queues 0, and 1 are WRR group 2, both with percentages of 50%.

You can configure this example with the following process.

#### Step 1: Enter configuration mode

Enter configuration mode for the device. Use the command:

```
awplus# configure terminal
```

#### Step 2: Add queue 2 to group 1

Put queue 2 into group 1 and allocate its percentage (50%). Use the command:

```
awplus(config)# mls qos scheduler-set 1 wrr-queue group 1 percent 50 queues  
2
```

#### Step 3: Add queues 3 and 4 to group 1

Put queues 3 and 4 into group 1 and allocate their percentages (25% each). Use the command:

```
awplus(config)# mls qos scheduler-set 1 wrr-queue group 1 percent 25 queues  
3 4
```

This means that in group 1, the switch will empty twice as many packets from queue 2 as it will from queues 3 and 4.

**Note:** The percentages in a group must add up to 100%. In this group, we have:

- queue 2 - 50%
- queue 3 - 25%
- queue 4 - 25%

This adds up to a total of 100% for group 1.

**Step 4: Add queues 0 and 1 to group 2**

Put queues 0 and 1 into group 2 and allocate their percentages (50% each). Use the command:

```
awplus(config)# mls qos scheduler-set 1 wrr-queue group 2 percent 50 queues
0 1
```

This means that in group 2, the switch will empty packets from queue 0 and queue 1 at the same rate.

**Note:** Again, for group 2, the percentages add up to 100%.

**Step 5: Enter interface configuration mode**

Enter interface configuration mode for the device. Use the command:

```
awplus(config)# interface port1.0.1
```

**Step 6: Set the scheduler-set on the interface**

Set port 1.0.1 to use scheduler-set 1. Use the command:

```
awplus(config-if)# mls qos scheduler-set 1
```

So now, in our example, on port1.0.1, the queues are processed as follows:

- Queue 7 is processed first.
- If queue 7 is empty, queue 6 is processed next.
- If queues 6 and 7 are empty, queue 5 is processed next.
- If queues 5 to 7 are empty, queues 2, 3 and 4 are processed with a ratio of 2:1:1
- If queues 2 to 7 are empty, queues 0 and 1 are processed with equal weighting.

**Step 7: Check scheduler-set settings**

You can use the following show command to display information about scheduler sets:

```
show mls qos scheduler-set
```

This show command will display WRR percentage-based values with a % suffix to differentiate from regular WRR weights.

```
awplus#show mls qos scheduler-set
Key: SP = Strict Priority
      WRR1 = Weighted Round Robin arbitration group 1
      WRR2 = Weighted Round Robin arbitration group 2

egress queue:          0      1      2      3      4      5      6      7
Scheduler-set 1 algorithm: WRR2  WRR2  WRR1  WRR1  WRR1  SP    SP    SP
WRR weight:      50%   50%   50%   25%   25%
```

The show command will also display a warning if:

- the WRR percentage-based values of a group do not sum to 100%.
- a WRR group contains a combination of weight and percentage values.

```
awplus#show mls qos scheduler-set
Key: SP = Strict Priority
      WRR1 = Weighted Round Robin arbitration group 1
      WRR2 = Weighted Round Robin arbitration group 2

egress queue:          0      1      2      3      4      5      6      7
Scheduler-set 1 algorithm: SP    SP    WRR1  WRR1  SP    WRR2  SP    SP
                       WRR weight:          30%  30          50%
Scheduler-set 2 algorithm: SP    SP    SP    SP    SP    WRR1  WRR1  SP
                       WRR weight:          80%  80%
Scheduler-set 3...
...
% Warning: Scheduler-set 1 WRR1: unsupported combination of weight and
percent.
% Warning: Scheduler-set 1 WRR1: total allocation of 30% is less than 100%.
% Warning: Scheduler-set 1 WRR2: total allocation of 50% is less than 100%.
% Warning: Scheduler-set 2 WRR1: total allocation of 160% exceeds 100%.
```

## Example: configuring DCBX on an interface

In the following example, you have completed the previous example, and configured ETS on port1.0.1. Now, you want to enable DCBX on port1.0.1. This will allow the interface to advertise its ETS settings.

You can configure this example with the following process.

### Step 1: Enter configuration mode

Enter configuration mode for the device. Use the command:

```
awplus#configure terminal
```

### Step 2: Enable the DCBX service

Enable the DCBX service on the device. Use the command:

```
awplus(config)#service dcbx
```

### Step 3: Enter interface configuration mode

Enter interface configuration mode for port1.0.1. Use the command:

```
awplus(config)#interface port1.0.1
```

### Step 4: Enable DCBX for ETS

Enable DCBX for ETS on the interface. Use the command:

```
awplus(config-if)#dcbx ets
```

**Step 5: Return to global configuration mode**

Return to global configuration mode for the device. Use the command:

```
awplus(config-if)#exit
```

**Step 6: Enable LLDP**

Enable LLDP for the interface. Use the command:

```
awplus(config)#lldp run
```

So now, in our example, port1.0.1 is using DCBX to advertise its ETS settings.

**Step 7: Check LLDP settings**

You can view the currently advertised ETS parameters with the **show lldp local-info** command:

```
awplus#show lldp local-info dot1 interface port1.0.1

LLDP Local Information:

Local port1.0.1:
  Chassis ID Type ..... MAC address
  Chassis ID ..... 889d.98df.23b2
  Port ID Type ..... Interface name
  Port ID ..... port1.0.1
  TTL ..... 120
  Port VLAN ID (PVID) ..... 1
  Port & Protocol VLAN - Supported . Yes
                        - Enabled ... No
                        - VIDs ..... 0
  VLAN Names ..... default
  Protocol IDs ..... 9000, 0026424203000000, 0027424203000002,
                    0069424203000003, 888e01,
                    aaaa0300e02b00bb, 88090101, 00540000e302,
                    000a424203000101, 0800, 0806, 86dd,
                    89020001, 89020128

ETS Configuration:
  Willing ..... 0
  CBS ..... 0
  Max TCs ..... 0
  Priority Assignment Table ... [0, 1, 2, 3, 4, 5, 6, 7]
  TC Bandwidth Table ..... [0, 0, 10, 10, 0, 0, 0, 0]
  TSA Assignment Table ..... [0, 2, 0, 2, 0, 2, 0, 2]

ETS Recommendation:
  Priority Assignment Table ... [0, 0, 0, 0, 0, 0, 0, 0]
  TC Bandwidth Table ..... [0, 0, 0, 0, 0, 0, 0, 0]
  TSA Assignment Table ..... [0, 0, 0, 0, 0, 0, 0, 0]

PFC:
  Willing ..... 0
  MBC ..... 0
  PFC Capability..... 8
  PFC Enable..... 4
```

This example is using the **dot1** parameter to show 802.1 TLVs, and the **interface** parameter to show the details for port1.0.1.

You can view the ETS parameters received from neighbors with the **show lldp neighbors detail** command:

```
awplus#show lldp neighbors detail dot1 interface port1.0.1

LLDP Detailed Neighbor Information:

Local port1.0.1:
  Neighbors table last updated 0 hrs 0 mins 2 secs ago

Chassis ID Type ..... MAC address
Chassis ID ..... 0000.cd37.0005
Port ID Type ..... Interface name
Port ID ..... port1.0.1
TTL ..... 120 (secs)
Port VLAN ID (PVID) ..... [not advertised]
Port & Protocol VLAN ..... [not advertised]
VLAN Names ..... [not advertised]
Protocol IDs ..... [not advertised]
ETS Configuration:
  Willing ..... 0
  CBS ..... 0
  Max TCs ..... 8
  Priority Assignment Table ... [0, 1, 2, 3, 4, 5, 6, 7]
  TC Bandwidth Table ..... [0, 0, 10, 10, 0, 0, 0, 0]
  TSA Assignment Table ..... [0, 0, 0, 2, 0, 2, 0, 0]
ETS Recommendation ..... [not advertised]
PFC:
  Willing ..... 0
  MBC ..... 0
  PFC Capability..... 8
  PFC Enable..... 3, 5
Application Priority ..... TCP/SCTP port 3260
  Priority ..... 5
Application VLAN ..... [not advertised]
```

This example is using the **dot1** parameter to show 802.1 TLVs, and the **interface** parameter to show the details for port1.0.1.

## Removing a scheduler-set

To remove a scheduler-set, use the following command:

```
no mls qos scheduler-set
```