AlliedWare Plus™ OS

How To | Configure QoS on x900-24, x900-12, and SwitchBlade x908 Series Switches

Introduction

This document describes some generic configuration examples for Quality of Service (QoS) on the AlliedWare Plus OS.

What information will you find in this document?

This document provides information in the following sections:

- 1. Setting the egress rate, on page 2
- 2. Setting the priority on a packet, on page 7
- 3. Setting the egress queues, on page 10
- Full QoS scenario - tiered services for a single customer, on page 14

Which product and software version does this information apply to?

This How To Note applies to the following Allied Telesis switches, running the AlliedWare Plus OS software version 5.2.2 or later:

- SwitchBlade x908
- x900 series

1. Setting the egress rate

Example 1-1: Setting the egress limit of a switch port

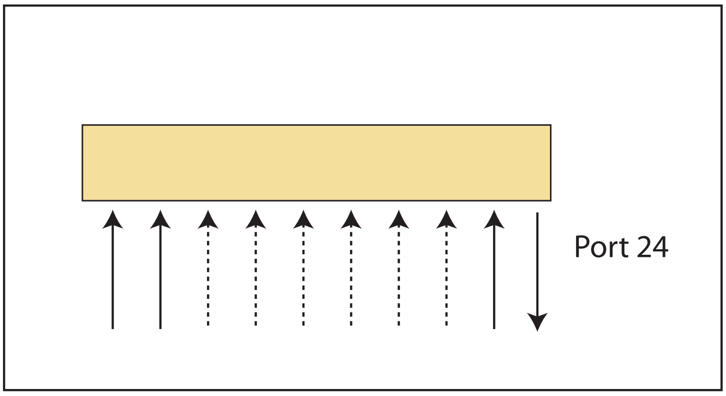

Commands:

```
interface port1.0.24
egress-rate-limit 25m
```

These commands will set the egress limit of port 24 to 25Mbps. The granularity is 651kbps.

Example 1-2: setting the maximum bandwidth limit per ingress port

Ingress rate limiting cannot be configured on the port per se, but is achieved by creating a QoS policy with a bandwidth limited traffic class, and applying that policy to each port.

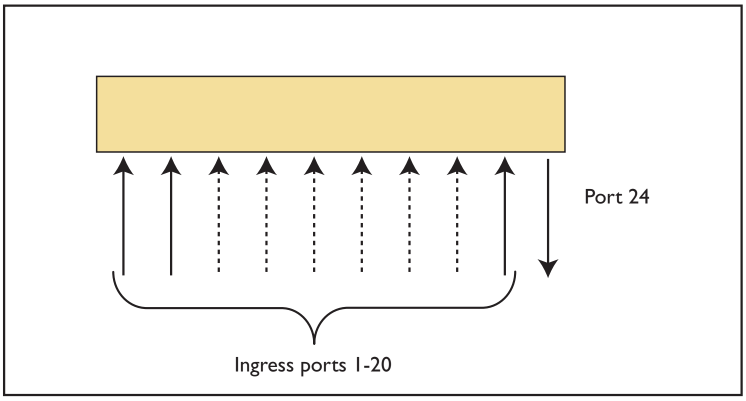

```
mls qos enable

class-map cmap1
policy-map pmap1
 class cmap1
  police single-rate 128 3000 5000 action drop-red
policy-map pmap2
 class cmap1
  police single-rate 256 3000 5000 action drop-red
policy-map pmap3
 class cmap1
  police single-rate 64 3000 5000 action drop-red
.
.
.
policy-map pmap20
 class cmap1
  police single-rate 512 3000 5000 action drop-red

interface port1.0.1
 service-policy input pmap1
interface port1.0.2
 service-policy input pmap2
interface port1.0.3
 service-policy input pmap3
.
.
.
interface port1.0.20
 service-policy input pmap20
```

This set of commands will set the ingress rate of traffic received per port from ports 1 to 20 to various different values. Note that the class-map matches EVERY packet. The granularity of the rate limiting is around 1 kbps.

Example 1-3: Setting the maximum bandwidth limit for each user

In example 1-1, we configured an egress rate on ports. But the granularity of that bandwidth limiting was multiples of 651kbps. It is possible to achieve finer-grained limiting using the devices advanced QoS capability. In this example we assume that there is one device with a known IP address, attached to each port. Policers will be configured to provide bandwidth limiting.

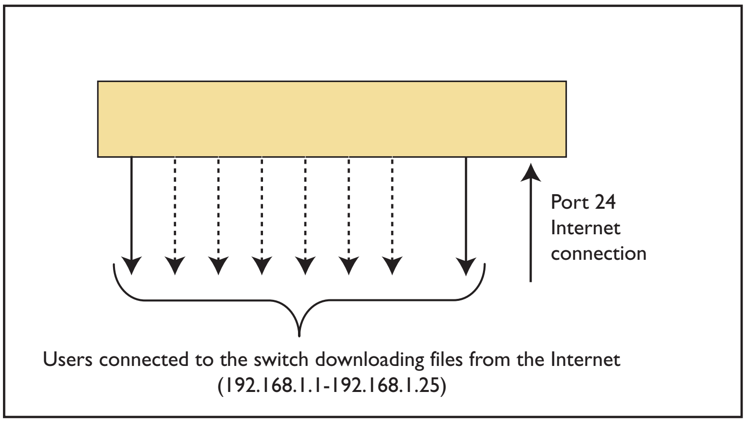
Port 24
Internet
connection
Users connected to the switch downloading files from the Internet
(192.168.1.1-192.168.1.25)

```
mls qos enable

access-list 3011 permit ip any 192.168.1.1/32
access-list 3012 permit ip any 192.168.1.2/32
access-list 3013 permit ip any 192.168.1.3/32
.
.
.

class-map cmap1
 match access-group 3011
class-map cmap2
 match access-group 3012
class-map cmap3
 match access-group 3013
.
.
.

policy-map pmap1
 class cmap1
  police single-rate 256 3000 5000 action drop-red
 class cmap2
  police single-rate 256 3000 5000 action drop-red
 class cmap3
  police single-rate 256 3000 5000 action drop-red
.
.
.

interface port1.0.24
 service-policy input pmap1
```

This set of commands will set the total egress limit (download traffic) for each user to 256kbps. The granularity is around 1kbps.

Example 1-4: Setting the maximum bandwidth limit for each IP subnet

This example is very similar to example 1-3, except that it is assumed that a whole subnet is attached to each port, not just a single device on each port.

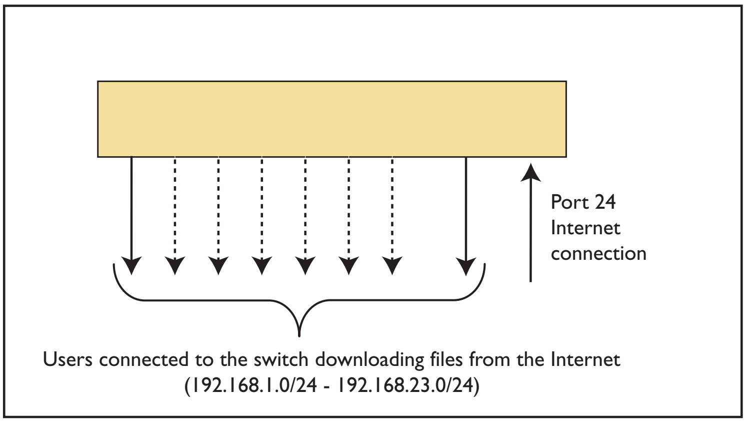
Port 24
Internet
connection
Users connected to the switch downloading files from the Internet
(192.168.1.0/24 - 192.168.23.0/24)

```
mls qos enable

access-list 3011 permit ip any 192.168.1.0/24
access-list 3012 permit ip any 192.168.2.0/24
access-list 3013 permit ip any 192.168.3.0/24
.
.
.

class-map cmap1
 match access-group 3011
class-map cmap2
 match access-group 3012
class-map cmap3
 match access-group 3013
.
.
.

policy-map pmap1
 class cmap1
  police single-rate 256 3000 5000 action drop-red
 class cmap2
  police single-rate 256 3000 5000 action drop-red
 class cmap3
  police single-rate 256 3000 5000 action drop-red
.
.
.

interface port1.0.24
 service-policy input pmap1
```

This set of commands will set the total egress limit (download traffic) for each subnet to 256kbps. The granularity is around 1kbps.

Example 1-5: Setting the maximum bandwidth limit of VLANs

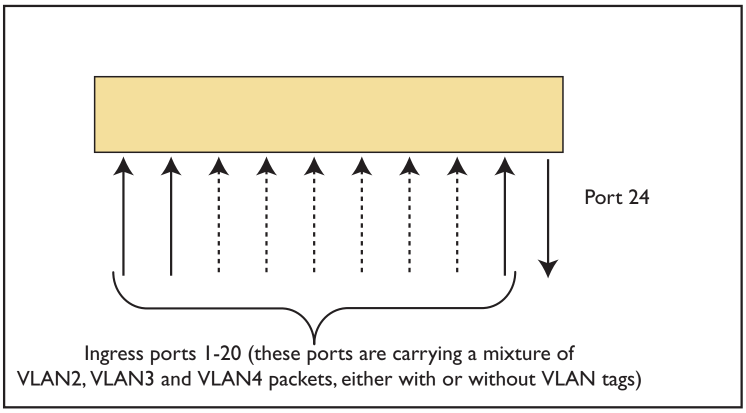
Port 24
Ingress ports 1-20 (these ports are carrying a mixture of VLAN2, VLAN3 and VLAN4 packets, either with or without VLAN tags)

```
mls qos enable

class-map cmap2
 match vlan 2
class-map cmap3
 match vlan 3
class-map cmap4
 match vlan 4

mls qos enable
mls qos aggregate-police agg1 single-rate 256 3000 5000 action drop-red
mls qos aggregate-police agg2 single-rate 256 3000 5000 action drop-red
mls qos aggregate-police agg3 single-rate 256 3000 5000 action drop-red

policy-map pmap1
 class cmap2
  police aggregate agg1
 class cmap3
  police aggregate agg2
 class cmap4
  police aggregate agg3

interface port1.0.1-1.0.20
 service-policy input pmap1
```

2. Setting the priority on a packet

Example 2-1: Setting the Layer 2 (VLAN/802.1p) priority per ingress port

Here we assign different 802.1p values to packets arriving on different ports.

These values are also known as the Layer 2 (L2) or VLAN priority.

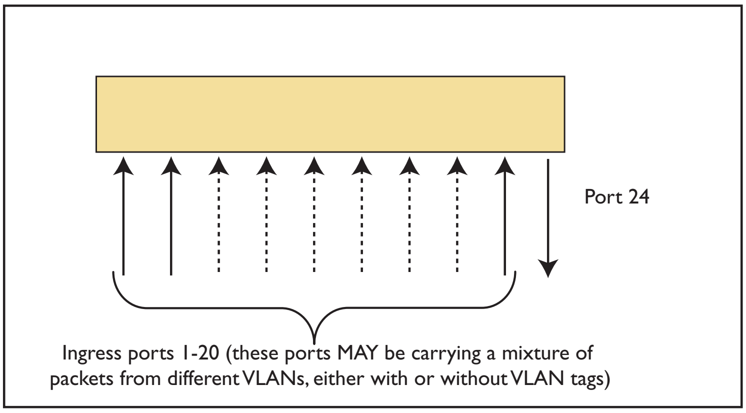
Port 24
Ingress ports 1-20 (these ports MAY be carrying a mixture of packets from different VLANs, either with or without VLAN tags)

```
mls qos enable

class-map cmap1
policy-map pmap1
 class cmap1
  set cos 6
policy-map pmap2
 class cmap1
  set cos 5
policy-map pmap3
 class cmap1
  set cos 4
.
.
.
interface port1.0.1
 service-policy input pmap1
interface port1.0.2
 service-policy input pmap2
interface port1.0.3
 service-policy input pmap3

.
.
.
```

Example 2-2: Setting the Layer 2 (VLAN/802.1p) priority per VLAN

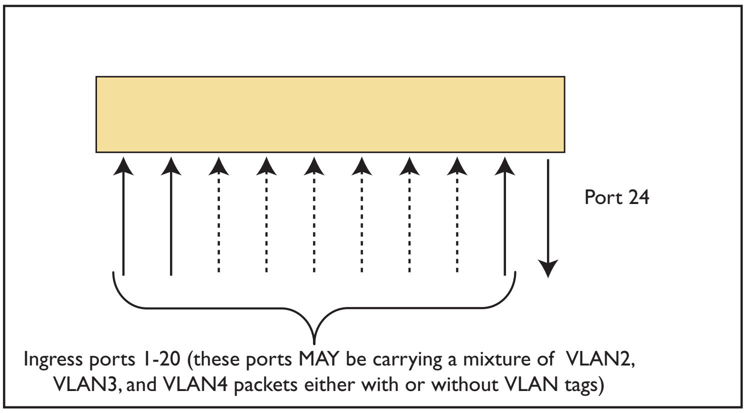
Port 24
Ingress ports 1-20 (these ports MAY be carrying a mixture of VLAN2, VLAN3, and VLAN4 packets either with or without VLAN tags)

```
mls qos enable

class-map cmap2
 match vlan 2
class-map cmap3
 match vlan 3
class-map cmap4
 match vlan 4

policy-map pmap1
 class cmap2
  set cos 6
 class cmap3
  set cos 5
 class cmap4
  set cos 4

interface port1.0.1-1.0.20
 service-policy input pmap1
```

This set of commands will set the Layer 2 priority of VLAN2, VLAN3 and VLAN4 traffics packets, received on ports 1 to 20, to 6, 5 and 4 respectively.

Example 2-3: Setting the Layer 3 (TOS/DCSP) priority per ingress port

Use the same configuration provided in Example 2-1 on page 7, but change the following line:

```
set cos <cos-value>
```

to

```
set dscp <dscp-value>
```

Example 2-4: Setting the Layer 3 (TOS/DSCP) priority per VLAN

Use the same configuration provided in Example 2-2 on page 8, but change the following line:

```
set cos <cos-value>
```

to

```
set dscp <dscp-value>
```

3. Setting the egress queues

In this section we look at methods for directing certain packets into certain queues on the egress port.

Example 3-1: Setting the egress queue according to the L2 priority of the incoming packet

The priority-to-queue map is a straightforward method for assigning packets to egress queues on the basis of the packets' 802.1p values.

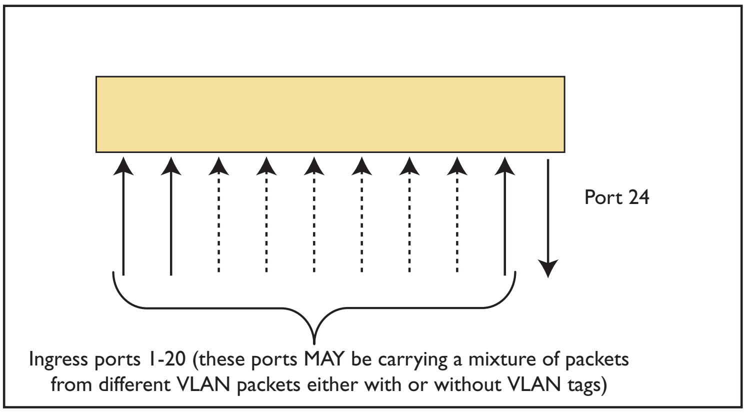
Port 24
Ingress ports 1-20 (these ports MAY be carrying a mixture of packets from different VLAN packets either with or without VLAN tags)

```
mls qos enable

mls qos map cos-queue 0 to 2
mls qos map cos-queue 1 to 3
mls qos map cos-queue 2 to 1
mls qos map cos-queue 3 to 0
mls qos map cos-queue 4 to 4
mls qos map cos-queue 5 to 5
mls qos map cos-queue 6 to 6
mls qos map cos-queue 7 to 7

interface port1.0.1-port1.0.20
 mls qos queue 2
```

Command settings

The **map cos-queue** commands in the above example set the mapping between the VLAN Tag User Priorities of the packets and the egress queues. Priorities 0 to 7 are mapped to queues 2, 3, 1, 0, 4, 5, 6 and 7 respectively.

The Interface mode command sets the incoming ports to send untagged packets to queue number 2, which means the untagged packets will use the same queue as the tagged packets with a User Priority of 1.

Notes: *The **cos-queue map** does not operate until you have entered the **mls qos enable** command–this means that packets will not be mapped to queues until this occurs.*
*It is also possible for packets to be "priority tagged". This means that the packets contain a VLAN tag with VID=0, but with a priority value set in the 802.1p field of the tag. The **cos-queue map** treats these packets like any other tagged packets, and will map them to the queue that corresponds to their 802.1p value.*

Example 3-2: Setting the egress queue according to the DSCP value of the incoming packet

The premark-dscp map allows you to specify a new queue for traffic based on the DSCP value of the incoming packet. The traffic you want to prioritize this way must be processed by a class that has **trust dscp** applied to it.

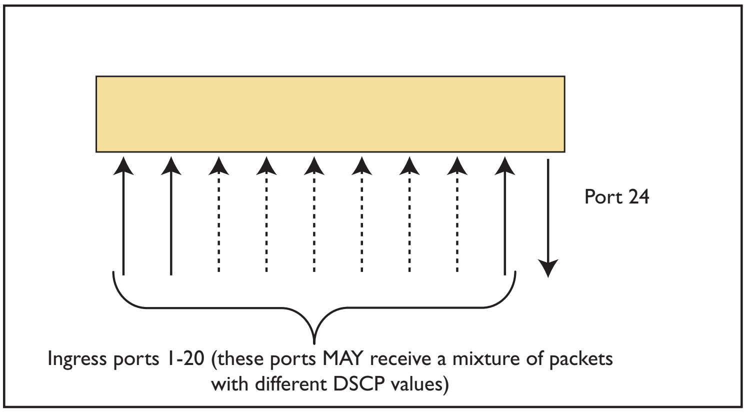
Port 24
Ingress ports 1-20 (these ports MAY receive a mixture of packets with different DSCP values)

```
mls qos map premark-dscp 12 to new-queue 1
mls qos map premark-dscp 24 to new-queue 2
mls qos map premark-dscp 27 to new-queue 3
mls qos map premark-dscp 35 to new-queue 4
mls qos map premark-dscp 40 to new-queue 5
mls qos map premark-dscp 49 to new-queue 6

mls qos enable
class-map cmap1
policy-map pmap1
 class cmap1
   trust dscp
interface port1.0.1-1.0.20
service-policy input pmap1
```

This will result in the queue mappings in the following table.

DSCP value	Queue
12	1
24	2
27	3
35	4
40	5
49	6
All others	Default value as set in the cos-queue map.

Example 3-3: Setting the egress queue according to the ingress port

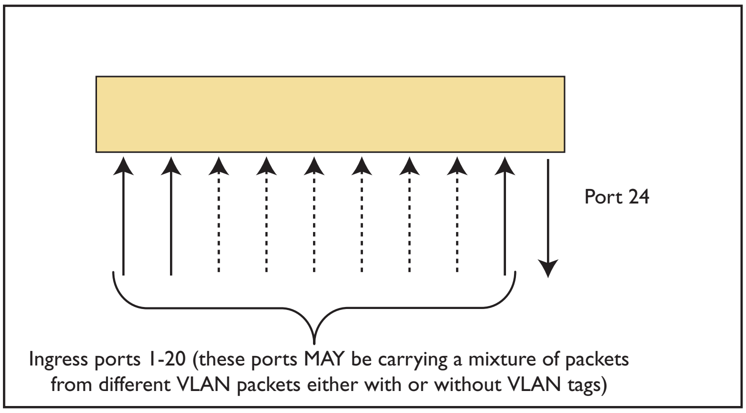
Port 24
Ingress ports 1-20 (these ports MAY be carrying a mixture of packets
from different VLAN packets either with or without VLAN tags)

```
mls qos enable

class-map cmap1
policy-map pmap1
 class cmap1
  set queue 6
policy-map pmap2
 class cmap1
  set queue 5
policy-map pmap3
 class cmap1
  set queue 4
interface port1.0.1
 service-policy input pmap1
interface port1.0.2
 service-policy input pmap2
interface port1.0.3
 service-policy input pmap3
.
.
.
```

This set of commands will set the egress queue of any traffic received from port 1 to egress queue 6, from port 2 to egress queue 5 … etc.

Example 3-4: Configuring WRR for egress queues

This example is used for configuring WRR on egress queues according to the ingress port of the traffic.

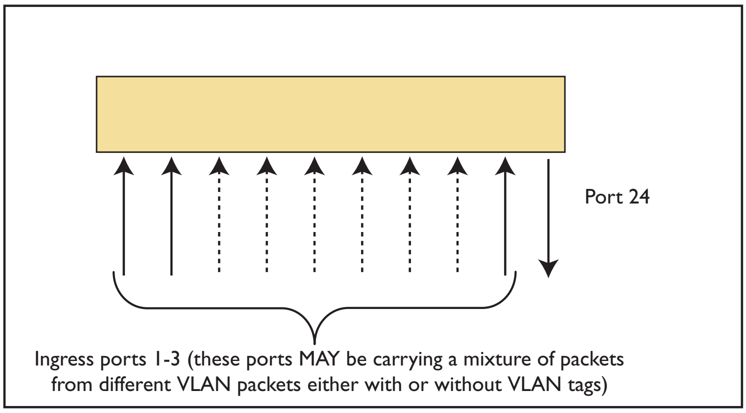
Port 24
Ingress ports 1-3 (these ports MAY be carrying a mixture of packets from different VLAN packets either with or without VLAN tags)

```
class-map cmap1
policy-map pmap1
 class cmap1
  set queue 6
policy-map pmap2
 class cmap1
  set queue 5
policy-map pmap3
 class cmap1
  set queue 4
interface port1.0.1
 service-policy input pmap1
interface port1.0.2
 service-policy input pmap2
interface port1.0.3
 service-policy input pmap3

interface port1.0.24
 wrr-queue group 1 weight 6 queues 4
 wrr-queue group 1 weight 12 queues 5
 wrr-queue group 1 weight 24 queues 6
```

This set of commands will set the egress queue of any traffic received from port 1 to egress queue 6, from port 2 to egress queue 5 and from port 3 to egress queue 4. And on egress, the WRR algorithm will do the following: for every 4 packets from queue 6, 2 packets will leave queue 5 and 1 packet will leave queue 4.

Full QoS scenario - tiered services for a single customer

In this section, we will build up a relatively complex QoS configuration to support a scenario requiring quite precise control over the traffic passing through the switch.

The scenario is an ISP providing connectivity for a customer, and offering different levels of service for different types of traffic.

The customer is connected to port 1 of the switch, and the uplink to the ISP is on port 24 of the switch.

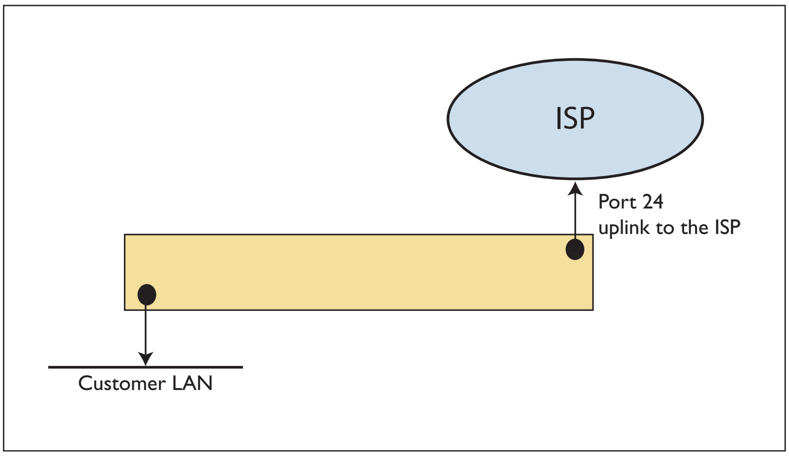

▶ **Step 1** - Set the egress bandwidth limiting

The service offered by the ISP puts a limit on the total bandwidth of traffic that the customer can send to the ISP.

This is achieved by setting a maximum bandwidth on the uplink port:

```
interface port1.0.24
 egress-rate-limit 10m
```

▶ **Step 2** - Give better service to some types of traffic than to others

The deal offered to the customer is that their traffic will be treated as belonging to three categories, and each category of traffic will be given a different level of service:

Gold traffic will be limited to, say, 2Mbps, but the ISP will guarantee delivery of the traffic across their network, with a low latency.

Silver traffic will have a much higher limit, but when congestion occurs, it will be throttled back in favour of Gold traffic, if necessary. The ISP will guarantee to deliver up to, say 5Mbps of Silver traffic across its network 90% of the time, but will give no guarantees about latency.

Bronze traffic will also have a high burst limit, but when congestion occurs, will be throttled back in favour of Gold traffic, and will share the remaining bandwidth with Silver traffic in a Weighted Round Robin fashion. The ISP makes no guarantees at all with regard to delivery of Bronze traffic across their network; it will be delivered on a best-effort basis.

Identify the types of traffic

The different categories of traffic will be identified by the DSCP values in the packets' headers. It is up to the customer to mark the packets with the appropriate DSCP values.

The DSCP values belonging to the different traffic categories are:

Gold	40
Silver	30
Bronze	0

The class maps to match these DSCP values are:

```
class-map cmap1
 match dscp 0
class-map cmap2
 match dscp 30
class-map cmap3
 match dscp 40
```

▶ Put the different categories of trafffic into different queues

This involves creating a policy-map, assigning the 3 class maps, and assigning the policy-map to a port.

```
policy-map pmap1
 class cmap1
  set queue 1
 class cmap2
  set queue 2
 class cmap3
  set queue 6

interface port1.0.1
 service-policy input pmap1
```

▶ Set the required properties on the egress queues

What is required is that:

- Gold traffic ALWAYS has precedence over Silver or Bronze traffic. So, when a Gold packet arrives at the egress port, it is transmitted immediately, irrespective of how many Silver or Bronze packets might be queued up.
- When there are Silver and Bronze packets queued up, they are transmitted according to a Weighted Round Robin (WRR) scheme.

This is achieved by ensuring that the egress queue to which Gold traffic is directed to is a priority queue, and the egress queues to which the silver and bronze traffic are directed are WRR queues.

So, we need to specify the queue types of queues 1, 2, and 6 on port 24. The relative weights to give to Silver and Bronze traffic are set by specifying the WRR weight for their queues. For example, to give a 4:1 ratio of Silver to Bronze traffic:

```
interface port1.0.24
 priority-queue 6
 wrr-queue group 1 weight 6 queues 1
 wrr-queue group 1 weight 24 queues 2
```

▶ Step 3 - set the bandwidth limits

The Gold traffic must be strictly limited to 2Mbps.

This is achieved by configuring a maxbandwidth on that traffic class, and dropping bandwidthclass 3 traffic:

```
policy-map pmap1
 class cmap1
  police single-rate 2000 15000 20000 action drop-red
```

For silver traffic, there is preferential treatment for the first 5Mbps of traffic. So, when there is congestion, you want to still be getting 5Mbps of Silver traffic through, if possible.

This is achieved by putting bandwidth limits on both Silver and Bronze traffic, and using RED curves to shape the throughput back to those limits when congestion occurs.

```
 class cmap2
  police single-rate 5000 25000 30000 action policed-dscp-transmit
 class cmap3
  police single-rate 10000 100000 125000 action policed-dscp-transmit

mls qos queue-set 1 queues 1 threshold 20000 50000 5000 10000 2000 6000
mls qos queue-set 1 queues 2 threshold 40000 60000 7000 12000 5000 10000
interface port1.0.24
 mls qos queue-set 1 random-detect
```

The **action** configured on the policers for each of the classes cmap2 and cmap3 is **policed-dscp-transmit**. That means that QoS passes the packets to the policed-dscp map for processing. If there are no entries in the policed-dscp map configured to alter the queue, DSCP, or priority of the packets, then all packets (green, yellow, and red) are passed straight through to be candidates for admission into the egress queues. In uncongested conditions, all packets are admitted into the egress queues and then transmitted. In congested conditions, as the lengths of the queues build up, the RED curves will admit green packets into the egress queues, in preference to yellow and red packets.

▶ Full configuration script

```
mls qos enable
mls qos queue-set 1 queues 1 threshold 20000 50000 5000 10000 2000 6000
mls qos queue-set 1 queues 2 threshold 40000 60000 70000 120000 5000 10000
class-map cmap1
 match dscp 0
!
class-map cmap2
 match dscp 30
!
class-map cmap3
 match dscp 40
!
policy-map pmap1
 class default
 class cmap1
  set queue 1
  police single-rate 2000 15000 20000 action drop-red
 class cmap2
  set queue 2
  police single-rate 5000 25000 30000 policed-dscp-transmit
 class cmap3
  set queue 3
  police single-rate 10000 100000 125000 policed-dscp-transmit
interface port1.0.1
 service-policy input pmap1
interface port1.0.24
 egress-rate-limit 10416
 mls qos queue-set 1 random-detect
 wrr-queue group 1 weight 6 queues 1
 wrr-queue group 1 weight 24 queues 2
```

USA Headquarters | 19800 North Creek Parkway | Suite 100 | Bothell | WA 98011 | USA | T: +1 800 424 4284 | F: +1 425 481 3895
European Headquarters | Via Motta 24 | 6830 Chiasso | Switzerland | T: +41 91 69769.00 | F: +41 91 69769.11
Asia-Pacific Headquarters | 11 Tai Seng Link | Singapore | 534182 | T: +65 6383 3832 | F: +65 6383 3830
www.alliedtelesis.com

Connecting The IP World

Allied Telesis™