How To | Configure QoS To Conform To Standard Marking Schemes

Introduction

This How To Note describes how to deploy a QoS solution across an entire network. It explains how to define per-hop behaviours (PHBs) that each switch in the network should perform, and how to ensure that the switches uniformly interpret the QoS information carried in packets. This means you can achieve a predictable performance for different traffic types across the whole network.

Packets that enter the network's edge may carry no QoS information. If so, the edge switch places such information into the packets before transmitting them to the next node. Thus, QoS information is preserved between nodes within the network and the nodes know what treatment to give each packet.

Contents

Which products and software versions does this How To Note apply to?

This configuration applies to all AlliedWare software versions for the following Allied Telesis switches:

- AT-8948
- AT-9900 series
- AT-9900s series
- x900 series

Related How To Notes

You also may find the following How To Notes useful:

- *Overview of Quality Of Service (QoS) Features On AT-8948, AT-9900, AT-9900s And x900 Series Switches*
- *How To Configure QoS On AT-8948, AT-9900, AT-9900s And x900 Series Switches*
- *How To Configure Filtering Actions On QoS Flow Groups And Traffic Classes*

How To Notes are available from www.alliedtelesis.com/resources/literature/howto.aspx.

Marking QoS information into packets

There are three options for marking QoS information into packets:

- 802.1p Class of Service (CoS) priority field within the VLAN tag of tagged Ethernet frames
- IP Precedence from Type of Service (TOS) field
- Differentiated Services (DiffServ) Code Point (DSCP)

Layer 2 CoS/802.1p

Layer 2 devices can use the 3-bit (binary 000-111; decimal 0-7) CoS field in 802.1q tagged frames to carry priority information. This field is also called the 802.1p field.

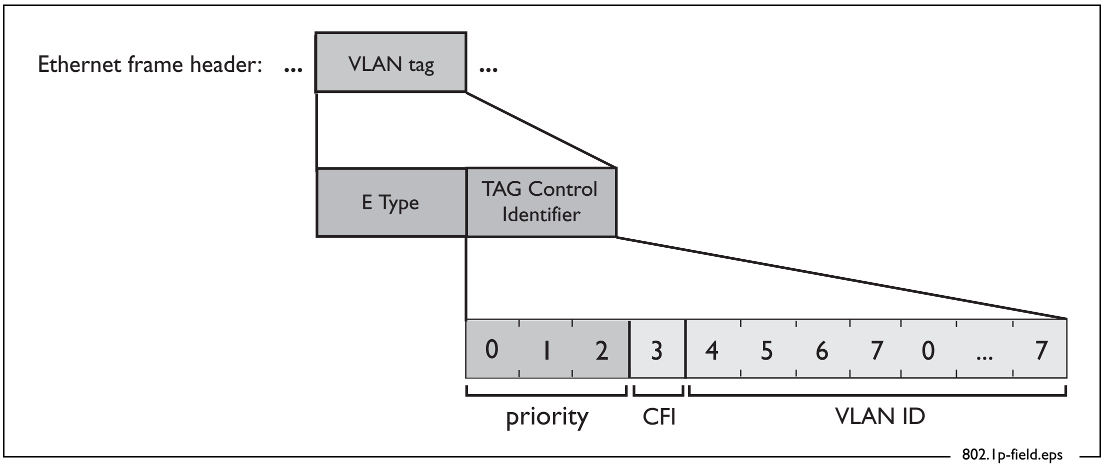

Layer 3 TOS and DSCP

Layer 3 devices can also use the 3-bit (binary: 000-111; decimal: 0-7) IP Precedence field from the TOS byte of the IP header. This is defined in RFC 791 and RFC 1349.

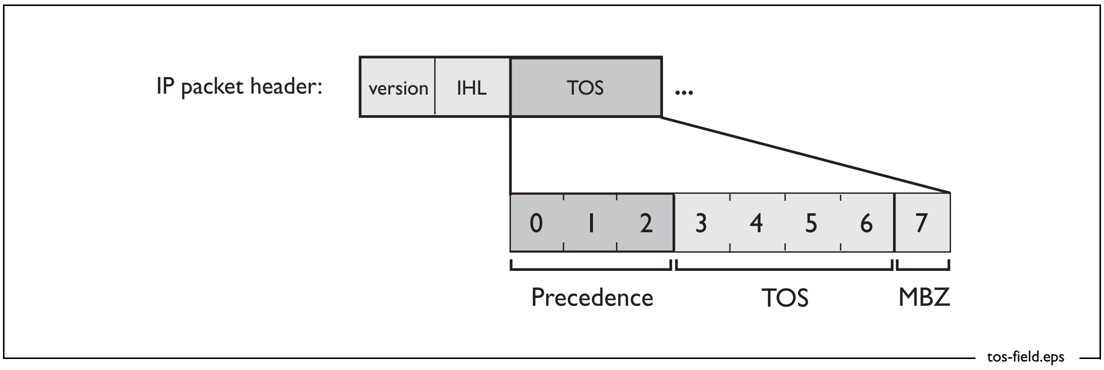

Alternatively, Layer 3 devices can use the Differentiated Services (DiffServ or DS) field from RFC 2474 and RFC 2475. This field is the TOS octet, redefined. The following figure shows the location of the DiffServ field in the IP header.

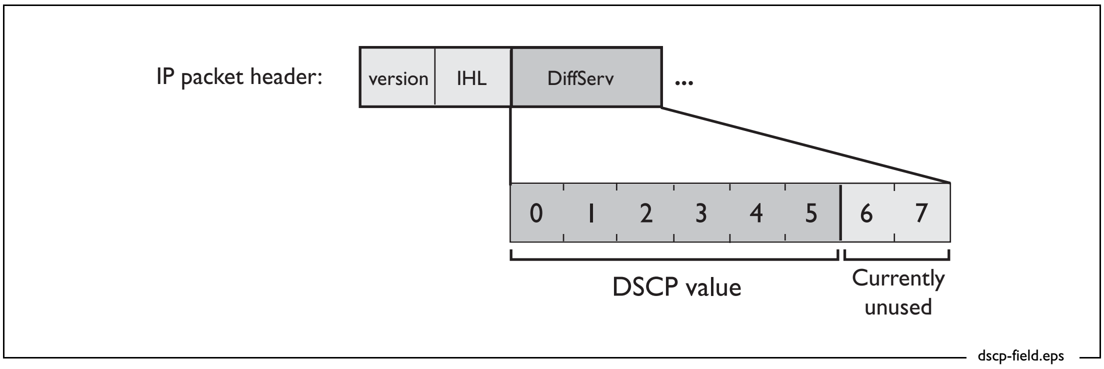

Interpreting the values

802.1p

Interpreting the 802.1p values is quite straightforward: the 8 values define 8 different levels of priority.

Two common mappings of 802.1p value to priority level are:

Low priority			→		high priority		
1	0	2	3	4	5	6	7
0	1	2	3	4	5	6	7

In the first mapping, an 802.1p value of 0 goes into the second-lowest priority queue instead of the lowest priority queue. Unclassified default traffic generally has an 802.1p value of 0, so this mapping reduces the probability of dropping default traffic.

DSCP

Interpreting DSCP values is a little more involved, because there are 64 possible values of DSCP, not just 8.

There are different schemes in common use for interpreting DSCP values:

- "Simple 8-level prioritisation" on page 5
- "More sophisticated per-hop behaviours" on page 6

Simple 8-level prioritisation

In this scheme, the DSCP values are grouped into 8 groups, based upon their first 3 bits. All the values in any given group share the same set of first 3 bits. These first 3 bits, of course, correspond to the IP Precedence field in the older TOS definition of the Layer 3 QoS byte.

These 8 groups of values can be mapped to 8 priority levels (i.e. 8 egress queues) as shown in the table below. Note that the last two values in the table (Internetwork and Network control) should be reserved for "Control Plane" traffic such as routing protocols and network management.

IP Precedence		DSCP range	CoS	Queue
Routine (Default)	000 (0)	000000 (0) - 000111 (7)	0	0
Priority	001 (1)	001000 (8) - 001111 (15)	1	1
Immediate	010 (2)	010000 (16) - 010111 (23)	2	2
Flash	011 (3)	011000 (24) - 011111 (31)	3	3
Flash Override	100 (4)	100000 (32) - 100111 (39)	4	4
Critical	101 (5)	101000 (40) - 101111 (47)	5	5
Internetwork control	110 (6)	110000 (48) - 110111 (55)	6	6
Network control	111 (7)	111000 (56) - 111111 (63)	7	7

More sophisticated per-hop behaviours

Some RFCs have been written that define mappings of DSCP values to quite specific ways in which the network devices should treat packets.

They define three distinct types of treatment (referred to as *per-hop behaviours*, PHBs):

- Expedited Forwarding (EF) PHB
- Assured Forwarding (AF) PHB
- Default PHB

Expedited Forwarding (EF) PHB

This traffic is given highest priority. Packets being given this treatment are prioritized ahead of all other types of data.

This is very suitable for traffic that requires low delay, low loss, and low jitter. Real-time services like voice and video require this kind of treatment.

If too much traffic is put into the top-priority queue, of course, then the packets no longer really get the advantage of being top priority; they end up sitting in the queue, and so suffering from delay and jitter.

Also, if too many packets are being put into the top-priority queue, then the lower-priority queues become starved.

For these reasons, typical networks will limit EF traffic to no more than 30%—and often much less—of the capacity of a link.

Assured Forwarding (AF) PHB

The AF PHB is characterized by a Committed Information Rate (CIR) and Excess Information Rate (EIR).

The network should be able to assure the delivery of all AF traffic up to the CIR.

Once the rate of AF traffic exceeds the CIR, then the network will attempt to deliver the extra traffic up to EIR. But, if the rate of AF traffic exceeds the EIR, then the traffic beyond that rate will almost invariably be dropped.

AF traffic is divided into 4 different classes (effectively 4 different priorities, but more about that below).

Within each class, the traffic is divided into 3 categories:

1. Traffic up to the CIR
2. Traffic that exceeded the CIR, but not the EIR
3. Traffic that exceeded the EIR

These three categories are referred to as *drop precedences*. The category that a packet falls into determines how likely it is to get dropped when congestion occurs.

Four classes, each with 3 drop precedences, defines the following 12 AF values:

	Green (low drop)	Yellow (medium drop)	Red (high drop)
AF Class1	AF11	AF12	AF13
AF Class2	AF21	AF22	AF23
AF Class3	AF31	AF32	AF33
AF Class4	AF41	AF42	AF43

If congestion occurs between classes, the traffic in the higher class is given priority. Rather than using strict priority queueing, more balanced queue servicing algorithms such as fair queueing or weighted fair queuing are likely to be used. This prevents high-priority queues from completely starving lower-priority queues—for example, it can ensure that if video traffic is over-subscribed, database traffic still gets some bandwidth.

If congestion occurs **within** a class, the packets with the higher drop precedence are discarded first. To prevent issues associated with tail drop, the random early detection (RED) algorithm is usually employed to decide which packets to drop.

The drop precedence is assigned to packets by a metering process in the edge switch where the packets enter the network. Typically, all traffic assigned to a class is initially given a low drop precedence. As the traffic rate exceeds the CIR, and then the EIR, the meter will increase the drop precedence of packets that exceed the threshold.

Default PHB

This is just best-effort forwarding of the lowest-priority traffic. Packets being given this treatment get a very limited amount of the available egress bandwidth in times of congestion.

Priority *versus* drop precedence

It is important to appreciate the difference between priority and drop precedence.

Priorities define which traffic is more important and therefore will go in a higher-priority or higher-weight queue. If a packet has a higher priority, that just means it is handled first, while other traffic waits in a queue.

Drop precedences determine which packets the switch first considers dropping during a period of congestion. A metering and colouring process calculates precedences dynamically as packets pass through the switch.

Traffic with the highest drop precedence may be referred to by any of the following terms: *high drop*, *red*, *exceeded EIR,* or *bandwidth class 3 (bwc 3)*.

Some network administrators decide that for some types of traffic in the network, all traffic in the red zone should be dropped, regardless of congestion levels. For example, this may reflect the level of service a customer has paid for.

PHB mapping table

The DSCP values corresponding to the EF, AF, and Default categories are shown in the table below.

				Queue
Default PHB			00000000	0
AF Class1	AF11 001010 (10)	AF12 001100 (12)	AF13 001110 (14)	1
AF Class2	AF21 010010 (18)	AF22 010100 (20)	AF23 010110 (22)	2
AF Class3	AF31 011010 (26)	AF32 011100 (28)	AF33 011110 (30)	3
AF Class4	AF41 100010 (34)	AF42 100100 (36)	AF43 100110 (38)	4
EF			EF 101110 (46)	5
	Green (low drop)	**Yellow (medium drop)**	**Red (high drop)**	

Implementation example 1: Mapping CoS, TOS, and DSCP to queues

In this example, the switch is configured to map layer 2 CoS (802.1p) values to egress queues, and to map layer 3 DSCP values to queues according to the simple 8-level prioritisation scheme described on page 5.

Configuration options

▶ Map incoming 802.1p priorities to egress queues for incoming **tagged** frames

Use the following command to set the mapping of incoming VLAN Tag User Priorities to the egress queues for incoming packets that include a VLAN tag header.

```
SET QOS PRIO2queuemap=0,1,2,3,4,5,6,7
```

With the above mapping, frames with an 802.1p priority of 0 go into queue 0, frames with an 802.1p priority of 1 go into queue 1, and so on.

▶ Map incoming 802.1p priorities to egress queues for incoming **untagged** frames

Untagged frames get assigned to a queue at ingress, according to the default setting for the port they ingress at. By default, untagged frames go into queue 2. Use the following command to change the default queue:

```
SET QOS POrt={port-list|ALL} [DEFAultqueue=queue-number]
```

Use the following command to set the VLAN Tag User Priority field for such frames.

```
SET QOS DEFAUltpriority=0,1,2,3,4,5,6,7
```

With the above mapping, frames in queue 0 are given an 802.1p priority of 0 at egress, and so on, unless their 802.1p priority value is assigned by another QoS mechanism (such as **dscpmap** or **queue2priomap**).

▶ Map DSCP values to egress queues

The mapping of DSCP values to egress queues is achieved using the Premarking table. For the simple 8-level prioritisation scheme shown in the table on page 5, DSCP values 0-7 map to queue 0; DSCP values 8-15 map to queue 1, on so on.

The following commands do three things:

- map each DSCP range to the correct queue. Note that queue 7 is the highest priority queue
- define a new DSCP value that will be written into the packets (basically, all DSCP values in an 8-value range will be simply replaced by the lowest value of that range)
- define the 802.1p value that will be written into frames if they are sent out tagged at egress.

```
SET QOS DSCPMap=PREmarking DSCP=0-7 NEWDscp=0 NEWQueue=0
   NEWPriority=0

SET QOS DSCPMap=PREmarking DSCP=8-15 NEWDscp=8 NEWQueue=1
   NEWPriority=1

SET QOS DSCPMap=PREmarking DSCP=16-23 NEWDscp=16 NEWQueue=2
   NEWPriority=2

SET QOS DSCPMap=PREmarking DSCP=24-31 NEWDscp=24 NEWQueue=3
   NEWPriority=3

SET QOS DSCPMap=PREmarking DSCP=32-39 NEWDscp=32 NEWQueue=4
   NEWPriority=4

SET QOS DSCPMap=PREmarking DSCP=40-47 NEWDscp=40 NEWQueue=5
   NEWPriority=5

SET QOS DSCPMap=PREmarking DSCP=48-55 NEWDscp=48 NEWQueue=6
   NEWPriority=6

SET QOS DSCPMap=PREmarking DSCP=56-63 NEWDscp=56 NEWQueue=7
   NEWPriority=7
```

Implementation example 2: Implementing per-hop behaviours in a network

In this example, we consider a simple network with a set of access switches, and a couple of core/aggregation switches. On this network, we will implement the default/AF/EF PHB scheme (shown in the table on page 8). The access switches will mark packets and the core/ aggregation switches will apply QoS to packets based on those marked values.

Another term for a network like this is a *Diffserv domain*.

Example network

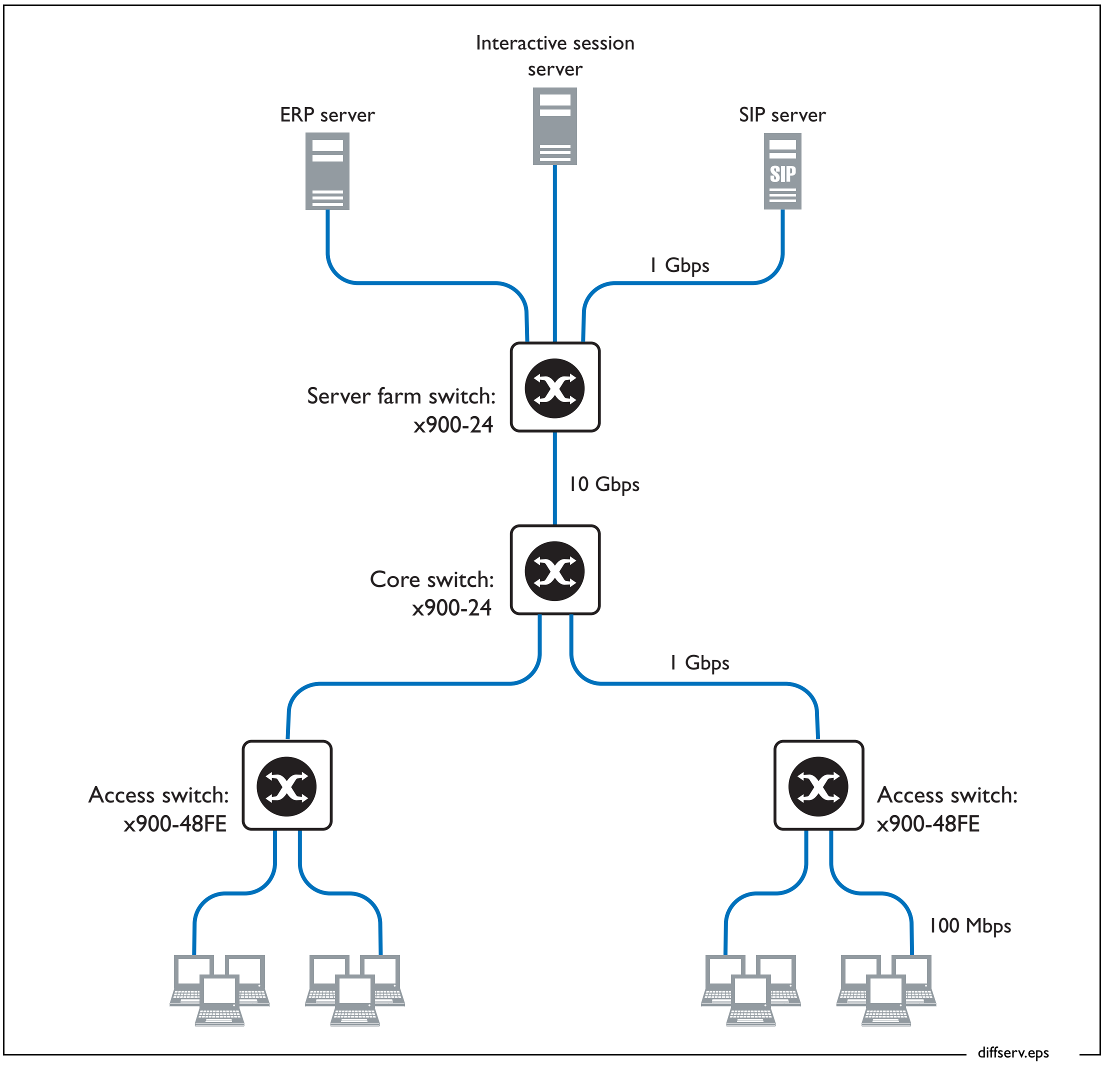

Traffic in the network is classified as shown in the following table. The table shows the DSCP and CoS values for each traffic class, and the egress queue that each type of traffic will be assigned to.

Application	DSCP	CoS	Queue
Management (OSPF, PIM, STP, BGP, SNMP, etc.)	56	7	7
Voice (stream)	EF(46)	5	5
Interactive (remote sessions on server)	AF41(34)	4	4
Voice (call signalling)	AF31(26)	3	3
Critical database (ERP)	AF21(18)	2	2
Default traffic	0	0	0

Note that this network has no traffic assigned to AF11. The command set in the following sections includes the commands for AF11, for demonstration purposes.

Configuration common to all the switches

The core, server farm, and access switches have the same settings for the following:

- "Premark table" on page 13
- "Remarking table" on page 14
- "Egress queue scheduling" on page 16
- "RED curves" on page 17
- "Marking and prioritisation of management traffic" on page 18

As well as these common settings, you need to add switch-specific classifier, flow group, traffic class, and policy settings. These are in:

- "Configuration of access switches" on page 19
- "Configuration of server farm switch" on page 21
- "Configuration of core switch" on page 23

Premark table

This premarking table implements the scheme illustrated in the PHB mapping table on page 8.

For example, traffic from AF21 arrives at the switch with a DSCP of 18 and the switch puts it into bandwidth class 1 (green) and queue 2, and puts the value 2 into the 802.1p field of tagged packets.

```
# Default
set qos dscpmap=premark dscp=0 newbwc=3 newq=0 newp=0

# AF11
set qos dscpmap=premark dscp=10 newbwc=1 newq=1 newp=1

# AF12
set qos dscpmap=premark dscp=12 newbwc=2 newq=1 newp=1

# AF13
set qos dscpmap=premark dscp=14 newbwc=3 newq=1 newp=1

# AF21
set qos dscpmap=premark dscp=18 newbwc=1 newq=2 newp=2

# AF22
set qos dscpmap=premark dscp=20 newbwc=2 newq=2 newp=2

# AF23
set qos dscpmap=premark dscp=22 newbwc=3 newq=2 newp=2

# AF31
set qos dscpmap=premark dscp=26 newbwc=1 newq=3 newp=3

# AF32
set qos dscpmap=premark dscp=28 newbwc=2 newq=3 newp=3

# AF33
set qos dscpmap=premark dscp=30 newbwc=3 newq=3 newp=3
```

All switches

```
# AF41
set qos dscpmap=premark dscp=34 newbwc=1 newq=4 newp=4

# AF42
set qos dscpmap=premark dscp=36 newbwc=2 newq=4 newp=4

# AF43
set qos dscpmap=premark dscp=38 newbwc=3 newq=4 newp=4

# EF
set qos dscpmap=premark dscp=46 newbwc=1 newq=5 newp=5

# Network management and control
set qos dscpmap=premark dscp=56 newbwc=1 newq=7 newp=7
```

Remarking table

The remarking table is a little more complex because remarking happens after bandwidth metering. The switch remarks packets' DSCP, queue and priority, on the basis of the DSCP it arrives with and its metered bandwidth class. This implementation is called colour-aware mode.

For example, if traffic from AF21 arrives at the switch with a DSCP of 18 and the metering process finds that it slightly exceeds its bandwidth allocation (bandwidth class 2, yellow), then it gets given a new DSCP of 20 and put into queue 2 with an 802.1p value of 2. This moves it into AF22.

The remark values are based on the incoming DSCP, to stop the switch from over-riding previous switches' marking of packets. If a previous switch found a packet to be non-conformant or semi-conformant, this switch cannot mark that packet as conformant.

Note that there is no need to remark default traffic. In times of congestion, we want the switch to simply drop most default traffic because it has the lowest priority, not to shape it to a CIR.

AF1 remarking

```
# AF11
set qos dscpmap=remark dscp=10 bwc=1 newdscp=10 newbw=1 newq=1 newp=1
set qos dscpmap=remark dscp=10 bwc=2 newdscp=12 newbw=2 newq=1 newp=1
set qos dscpmap=remark dscp=10 bwc=3 newdscp=14 newbw=3 newq=1 newp=1

# AF12
set qos dscpmap=remark dscp=12 bwc=1 newdscp=12 newbw=2 newq=1 newp=1
set qos dscpmap=remark dscp=12 bwc=2 newdscp=12 newbw=2 newq=1 newp=1
set qos dscpmap=remark dscp=12 bwc=3 newdscp=14 newbw=3 newq=1 newp=1

# AF13
set qos dscpmap=remark dscp=14 bwc=1 newdscp=14 newbw=3 newq=1 newp=1
set qos dscpmap=remark dscp=14 bwc=2 newdscp=14 newbw=3 newq=1 newp=1
set qos dscpmap=remark dscp=14 bwc=3 newdscp=14 newbw=3 newq=1 newp=1
```

All switches

AF2 remarking

```
# AF21
set qos dscpmap=remark dscp=18 bwc=1 newdscp=18 newbw=1 newq=2 newp=2
set qos dscpmap=remark dscp=18 bwc=2 newdscp=20 newbw=2 newq=2 newp=2
set qos dscpmap=remark dscp=18 bwc=3 newdscp=22 newbw=3 newq=2 newp=2

# AF22
set qos dscpmap=remark dscp=20 bwc=1 newdscp=20 newbw=2 newq=2 newp=2
set qos dscpmap=remark dscp=20 bwc=2 newdscp=20 newbw=2 newq=2 newp=2
set qos dscpmap=remark dscp=20 bwc=3 newdscp=22 newbw=3 newq=2 newp=2

# AF23
set qos dscpmap=remark dscp=22 bwc=1 newdscp=22 newbw=3 newq=2 newp=2
set qos dscpmap=remark dscp=22 bwc=2 newdscp=22 newbw=3 newq=2 newp=2
set qos dscpmap=remark dscp=22 bwc=3 newdscp=22 newbw=3 newq=2 newp=2
```

AF3 remarking

```
# AF31
set qos dscpmap=remark dscp=26 bwc=1 newdscp=26 newbw=1 newq=3 newp=3
set qos dscpmap=remark dscp=26 bwc=2 newdscp=28 newbw=2 newq=3 newp=3
set qos dscpmap=remark dscp=26 bwc=3 newdscp=30 newbw=3 newq=3 newp=3

# AF32
set qos dscpmap=remark dscp=28 bwc=1 newdscp=28 newbw=2 newq=3 newp=3
set qos dscpmap=remark dscp=28 bwc=2 newdscp=28 newbw=2 newq=3 newp=3
set qos dscpmap=remark dscp=28 bwc=3 newdscp=30 newbw=3 newq=3 newp=3

# AF33
set qos dscpmap=remark dscp=30 bwc=1 newdscp=30 newbw=3 newq=3 newp=3
set qos dscpmap=remark dscp=30 bwc=2 newdscp=30 newbw=3 newq=3 newp=3
set qos dscpmap=remark dscp=30 bwc=3 newdscp=30 newbw=3 newq=3 newp=3
```

AF4 remarking

```
# AF41
set qos dscpmap=remark dscp=34 bwc=1 newdscp=34 newbw=1 newq=4 newp=4
set qos dscpmap=remark dscp=34 bwc=2 newdscp=36 newbw=2 newq=4 newp=4
set qos dscpmap=remark dscp=34 bwc=3 newdscp=38 newbw=3 newq=4 newp=4

# AF42
set qos dscpmap=remark dscp=36 bwc=1 newdscp=36 newbw=2 newq=4 newp=4
set qos dscpmap=remark dscp=36 bwc=2 newdscp=36 newbw=2 newq=4 newp=4
set qos dscpmap=remark dscp=36 bwc=3 newdscp=38 newbw=3 newq=4 newp=4

# AF43
set qos dscpmap=remark dscp=38 bwc=1 newdscp=38 newbw=3 newq=4 newp=4
set qos dscpmap=remark dscp=38 bwc=2 newdscp=38 newbw=3 newq=4 newp=4
set qos dscpmap=remark dscp=38 bwc=3 newdscp=38 newbw=3 newq=4 newp=4
```

EF remarking

```
# EF
set qos dscpmap=remark dscp=46 bwc=1 newdscp=46 newbw=1 newq=5 newp=5
set qos dscpmap=remark dscp=46 bwc=2 newdscp=46 newbw=2 newq=5 newp=5
set qos dscpmap=remark dscp=46 bwc=3 newdscp=46 newbw=3 newq=5 newp=5
```

All switches

Egress queue scheduling

Egress queue scheduling is the stage that most determines the quality of service each traffic type receives.

We use:

- **Strict Priority** scheduling for the high priority queues 5-7. EF traffic is in queue 5 and management traffic is in queue 7.
- **Weighted Round Robin (WRR)** scheduling for queues 0-4, with weights of 5, 15, 20, 25, and 35 respectively. Default traffic is in queue 0 and AF traffic is in queues 1-4.

This queue scheduling means that:

- management traffic has top priority (in queue 7)
- voice traffic has next priority (in queue 5)
- interactive session traffic, voice signalling, database traffic and default traffic share the remaining bandwidth in a weighted round robin fashion (in queues 0-4)

If the standard PHB scheme does not completely suit your network, you can tweak the queue scheduling. For example, the queue scheduling in this section places traffic for the Default PHB in the WRR group instead of leaving it strict priority. This avoids starving the default class. Even when the other low-priority classes are busy, the default class gets to forward some traffic.

For another example, if you had streaming video traffic as well as voice traffic, you could place the video stream into queue 4 and make it a strict priority queue. This would make queue 4 act like it was EF instead of AF.

```
# Configure Weighted Round Robin scheduling on lower-priority queues:

set qos po=all egr=0 sch=wrr1 wrr=5
set qos po=all egr=1 sch=wrr1 wrr=15
set qos po=all egr=2 sch=wrr1 wrr=20
set qos po=all egr=3 sch=wrr1 wrr=25
set qos po=all egr=4 sch=wrr1 wrr=35

# Map all unclassified incoming frames to default queue 0 to prevent
# clients from marking frames:

set qos po=all defaultqueue=0
set qos po=all forcedefaultqueue=yes
```

Note that we do not have to configure Strict Priority queueing for queues 5-7 because Strict Priority is the default setting.

RED curves

We use RED curves for WRR queues 0-4 (the queues used by the default and AF traffic). We need to create a RED curve set for the different queues and bandwidth class (green, yellow, red). Of course, we give the red bandwidth class 3 a higher probability of being dropped.

We do not want to use RED curves on the strict priority queues 5-7, because these queues are used for EF and management traffic. We only ever want to drop red bandwidth class 3 traffic from these queues, not yellow or green traffic.

```
# Create the RED curve set
cre qos red=3

# Queue 0                                                    # drop prob
set qos red=3 que=0 start1=30kB stop1=40kB drop1=3          # 12.5%
set qos red=3 que=0 start2=20kB stop2=30kB drop1=2          # 25%
set qos red=3 que=0 start3=10kB stop3=20kB drop1=1          # 50%

# Queue 1
set qos red=3 que=1 start1=30kB stop1=40kB drop1=3          # 12.5%
set qos red=3 que=1 start2=20kB stop2=30kB drop1=2          # 25%
set qos red=3 que=1 start3=10kB stop3=20kB drop1=1          # 50%

# Queue 2
set qos red=3 que=2 start1=30kB stop1=40kB drop1=3          # 12.5%
set qos red=3 que=2 start2=20kB stop2=30kB drop1=2          # 25%
set qos red=3 que=2 start3=10kB stop3=20kB drop1=1          # 50%

# Queue 3
set qos red=3 que=3 start1=30kB stop1=40kB drop1=3          # 12.5%
set qos red=3 que=3 start2=20kB stop2=30kB drop1=2          # 25%
set qos red=3 que=3 start3=10kB stop3=20kB drop1=1          # 50%

# Queue 4
set qos red=3 que=4 start1=30kB stop1=40kB drop1=3          # 12.5%
set qos red=3 que=4 start2=20kB stop2=30kB drop1=2          # 25%
set qos red=3 que=4 start3=10kB stop3=20kB drop1=1          # 50%

# Disable RED curves for Strict Priority Scheduling queues by setting
# the start value very high (higher than the queue is ever likely to
# fill to).
# Queue 5
set qos red=3 que=5 start1=100kB stop1=200kB drop1=15       # ~0%
set qos red=3 que=5 start2=100kB stop2=200kB drop1=15       # ~0%
set qos red=3 que=5 start3=100kB stop3=200kB drop1=15       # ~0%

# Queue 6
set qos red=3 que=6 start1=100kB stop1=200kB drop1=15       # ~0%
set qos red=3 que=6 start2=100kB stop2=200kB drop1=15       # ~0%
set qos red=3 que=6 start3=100kB stop3=200kB drop1=15       # ~0%

# Queue 7
set qos red=3 que=7 start1=100kB stop1=200kB drop1=15       # ~0%
set qos red=3 que=7 start2=100kB stop2=200kB drop1=15       # ~0%
set qos red=3 que=7 start3=100kB stop3=200kB drop1=15       # ~0%

# Apply the RED curve set to all ports
set qos po=all red=3
```

Marking and prioritisation of management traffic

We want to set the egress queue for the packets generated by the switch's CPU (such as STP, BGP, OSPF, SNMP, and so on) to queue 7, and make this the highest priority queue of all, to prevent any loss of these packets.

Also, we want to mark IP packets among this traffic with DSCP=56, so that they are given EF treatment when being forwarded by other switches.

```
set switch cputxque=7

set ip dscpoverride=56
```

Configuration of access switches

As well as the "Configuration common to all the switches" on page 13, add this section's configuration to each access switch.

The access switches use classifiers matching on various values to assign traffic to different traffic classes. Then they mark the packets, and forward them based on the traffic class they belong to.

```
# Create classifiers to classify traffic by application.
# We assume that all the VoIP equipment (handsets and SIP server) has
# been configured to mark SIP signaling packets with DSCP=26 and VoIP
# RTP packets with DSCP=46.

# VoIP RTP traffic
cre class=11 ipdscp=46

# SIP signaling traffic
cre class=12 ipdscp=26,28,30

# Traffic to interactive server
cre class=21 ipda=<your-value> tcpdp=<your-value>

# Traffic from interactive server (this traffic is already
# DSCP-marked by the server farm switch)
cre class=22 ipdscp=34,36,38

# Traffic to database server
cre class=31 ipda=<your-value> tcpdp=<your-value>

# Traffic from database server (this traffic is already
# DSCP-marked by the server farm switch)
cre class=32 ipdscp=18,20,22

# Create QoS flow groups and add classifiers to them

cre qos flow=11 descr="RTP to core" premark=usedscp
cre qos flow=12 descr="signaling to core" premark=usedscp
add qos flow=11 class=11
add qos flow=12 class=12

cre qos flow=13 descr="RTP from core" premark=usedscp
cre qos flow=14 descr="signaling from core" premark=usedscp
add qos flow=13 class=11
add qos flow=14 class=12

# The access switch identifies traffic to each server by IP
# settings (see classifiers) and marks it with the appropriate DSCP.

cre qos flow=21 descr="to interactive server" premark=usemark mark=34
cre qos flow=22 descr="from interactive server" premark=usedscp
add qos flow=21 class=21
add qos flow=22 class=22

cre qos flow=31 descr="to database server" premark=usemark mark=18
cre qos flow=32 descr="from database server" premark=usedscp
add qos flow=31 class=31
add qos flow=32 class=32
```

```
# Create QoS traffic classes, and add flow groups to them

cre qos traf=11 descr="voip RTP and signaling to core"
   dropbwclass3=yes
# Drop bandwidth class 3 to avoid high flows of this traffic from
# starving lower-priority queues:
set qos traf=11 max=100Mbps maxbu=1000K min=80Mbps minbu=800K
   remark=usedscpm
add qos traf=11 flow=11,12

cre qos traf=13 descr="voip RTP and signaling from core"
   dropbwclass3=yes
set qos traf=13 max=100Mbps maxbu=1000K min=80Mbps minbu=800K
   remark=usedscpm
add qos traf=13 flow=13,14

cre qos traf=21 descr="to interactive server"
set qos traf=21 max=10Mbps maxbu=100K min=8Mbps minbu=80K
   remark=usedscpm
add qos traf=21 flow=21

cre qos traf=22 descr="from interactive server"
set qos traf=22 max=10Mbps maxbu=100K min=8Mbps minbu=80K
   remark=usedscpm
add qos traf=22 flow=22

cre qos traf=31 descr="to database server"
set qos traf=31 max=300Mbps maxbu=3000K min=240Mbps minbu=2400K
   remark=usedscpm
add qos traf=31 flow=31

cre qos traf=32 descr="from database server"
set qos traf=32 max=300Mbps maxbu=3000K min=240Mbps minbu=2400K
   remark=usedscpm
add qos traf=32 flow=32


# Create QoS policies, add traffic classes to them, and apply them to
# ports

cre qos poli=1 descr="to core" dtcpremark=usemarkvalue mark=0
add qos poli=1 traff=11,21,31
cre qos portgroup=1 port=1-48
set qos portgroup=1-48 poli=1

cre qos poli=2 descr="from core" dtcpremark=usemarkvalue mark=0
add qos poli=2 traff=13,22,32
set qos port=49-52 poli=2
```

Configuration of server farm switch

As well as the "Configuration common to all the switches" on page 13, add this section's configuration to the server farm switch.

```
# Create classifiers to classify traffic by application.
# We assume that SIP server have been configured to mark SIP
# signaling packets with DSCP=26.

# SIP signaling.
# Signaling to and from the SIP server uses the same classifier. This
# is okay because it goes into two flow groups that end up in
# different policies on different ports.
cre class=11 ipdscp=26,28,30

# Traffic to interactive server (this traffic is already DSCP-marked
# by the access switch)
cre class=21 ipdscp=34,36,38

# Traffic from interactive server
cre class=22 ipda=<ipadd> tcpdp=<ipadd>

# Traffic to database server (this traffic is already DSCP-marked by
# the access switch)
cre class=31 ipdscp=18,20,22

# Traffic from the database server
cre class=32 ipda=<ipadd> tcpdp=<ipadd>

# Create QoS flow groups, and add classifiers to them

cre qos flow=11 descr="signaling to SIP server" premark=usedscp
cre qos flow=12 descr="signaling from SIP server" premark=usedscp
add qos flow=11 class=11
add qos flow=12 class=11

# The server farm switch identifies traffic from each server by IP
# settings (see classifiers) and marks it with the appropriate DSCP.

cre qos flow=21 descr="to interactive server" premark=usedscp
cre qos flow=22 descr="from interactive server" premark=usemark
   mark=34
add qos flow=21 class=21
add qos flow=22 class=22

cre qos flow=31 descr="to database server" premark=usedscp
cre qos flow=32 descr="from database server" premark=usemark mark=18
add qos flow=31 class=31
add qos flow=32 class=32
```

```
# Create QoS traffic classes, and add flow groups to them

cre qos traf=11 descr="signaling to SIP server" dropbwclass3=yes
set qos traf=11 max=100Mbps maxbu=1000K min=80Mbps minbu=800K
   remark=usedscpm
add qos traf=11 flow=11

cre qos traf=12 descr="signaling from SIP server" dropbwclass3=yes
set qos traf=12 max=100Mbps maxbu=1000K min=80Mbps minbu=800K
   remark=usedscpm
add qos traf=12 flow=12

cre qos traf=21 descr="to interactive server"
set qos traf=21 max=10Mbps maxbu=100K min=8Mbps minbu=80K
   remark=usedscpm
add qos traf=21 flow=21

cre qos traf=22 descr="from interactive server"
set qos traf=22 max=10Mbps maxbu=100K min=8Mbps minbu=80K
   remark=usedscpm
add qos traf=22 flow=22

cre qos traf=31 descr="to database server"
set qos traf=31 max=300Mbps maxbu=3000K min=240Mbps minbu=2400K
   remark=usedscpm
add qos traf=31 flow=31

cre qos traf=32 descr="from database server"
set qos traf=32 max=300Mbps maxbu=3000K min=240Mbps minbu=2400K
   remark=usedscpm
add qos traf=32 flow=32


# Create QoS policies, add traffic classes to them, and apply them to
# ports

cre qos poli=1 descr="to core" dtcpremark=usemarkvalue mark=0
add qos poli=1 traff=12,22,32
set qos port=1-24 poli=1

cre qos poli=2 descr="from core" dtcpremark=usemarkvalue mark=0
add qos poli=2 traff=11,21,31
set qos port=101 poli=2
```

Core switch

Configuration of core switch

As well as the "Configuration common to all the switches" on page 13, add this section's configuration to the core switch.

```
# Create classifiers to classify traffic by application

# VoIP RTP stream
cre class=11 ipdscp=46
# VoIP signaling
cre class=12 ipdscp=26,28,30
# Interactive server traffic
cre class=21 ipdscp=34,36,38
# Database server traffic
cre class=31 ipdscp=18,20,22

# Create QoS flow groups and add classifiers to them

cre qos flow=11 descr="RTP stream" premark=usedscp
cre qos flow=12 descr="SIP signaling" premark=usedscp
add qos flow=11 class=11
add qos flow=12 class=12

cre qos flow=21 descr="interactive server traffic" premark=usedscp
add qos flow=21 class=21

cre qos flow=31 descr="database server traffic" premark=usedscp
add qos flow=31 class=31

# Create QoS traffic classes and add flow groups to them

cre qos traf=11 descr="RTP stream" dropbwclass3=yes
set qos traf=11 max=100Mbps maxbu=1000K min=80Mbps minbu=800K
   remark=usedscpm
add qos traf=11 flow=11

cre qos traf=12 descr="SIP signaling" dropbwclass3=yes
set qos traf=12 max=100Mbps maxbu=1000K min=80Mbps minbu=800K
   remark=usedscpm
add qos traf=12 flow=12

cre qos traf=21 descr="interactive server traffic"
set qos traf=21 max=10Mbps maxbu=100K min=8Mbps minbu=80K
   remark=usedscpm
add qos traf=21 flow=21

cre qos traf=31 descr="database server traffic"
set qos traf=31 max=10Mbps maxbu=100K min=8Mbps minbu=80K
   remark=usedscpm
add qos traf=31 flow=31

# Create a QoS policy, add traffic classes, and apply it to all ports

cre qos poli=1 descr="bi-directional" dtcpremark=usemarkvalue mark=0
add qos poli=1 traff=11,12,21,31
set qos port=all poli=1
```

USA Headquarters | 19800 North Creek Parkway | Suite 200 | 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

C613-16104-00 REV A