Chapter 20

Configuring CoS on ATM Interfaces

The ATM2 intelligent queuing (IQ) interface allows multiple IP queues into each virtual circuit (VC). On M-series platforms (except the M320 router), a VC tunnel can support four class-of-service (CoS) queues. On the M320 and T-series platforms for all ATM2 IQ PICs except the OC48 PIC, a VC tunnel can support eight CoS queues. Within a VC tunnel, the weighted round-robin (WRR) algorithm schedules the cell transmission of each queue. You can configure the queue admission policies, such as early packet discard (EPD) or weighted random early detection (WRED), to control the queue size during congestion.

For information about CoS components that apply generally to all interfaces, see "CoS Overview" on page 1 and "CoS Configuration" on page 39. For general information about configuring ATM interfaces, see the *JUNOS Network Interfaces Configuration Guide*.

To configure ATM2 IQ VC tunnel CoS components, include the following statements at the [edit interfaces at-*fpc/pic/port*] hierarchy level:

```
[edit chassis fpc slot-number pic pic-number]
max-queues-per-interface number;

[edit interfaces at-fpc/pic/port]
atm-options {
    linear-red-profiles profile-name {
        high-plp-max-threshold percent;
        low-plp-max-threshold percent;
        queue-depth cells high-plp-threshold percent low-plp-threshold percent;
    }
    plp-to-clp;
    scheduler-maps map-name {
        forwarding-class class-name {
            epd-threshold cells plp1 cells;
            linear-red-profile profile-name;
            priority (high | low);
            transmit-weight (cells number | percent number);
        }
        vc-cos-mode (alternate | strict);
    }
}
```

```
unit logical-unit-number {
    atm-scheduler-map (map-name | default);
    family family {
        address address {
            destination address;
        }
    }
    plp-to-clp;
    shaping {
        (cbr rate | rtvbr peak rate sustained rate burst length |
                vbr peak rate sustained rate burst length);
    }
    vci vpi-identifier.vci-identifier;
}
```

This section discusses the following topics:

Configuring Linear RED Profiles

Linear random early detection (RED) profiles define CoS virtual circuit drop profiles. You can configure up to 32 linear RED profiles per port. When a packet arrives, RED checks the queue fill level. If the fill level corresponds to a nonzero drop probability, the RED algorithm determines whether to drop the arriving packet.

To configure linear RED profiles, include the **linear-red-profiles** statement at the [**edit interfaces at-***fpc/pic/port* **atm-options**] hierarchy level:

```
[edit interfaces at-fpc/pic/port atm-options]
linear-red-profiles profile-name {
    high-plp-max-threshold percent;
    low-plp-max-threshold percent;
    queue-depth cells high-plp-threshold percent low-plp-threshold percent;
}
```

The **queue-depth**, **high-plp-threshold**, and **low-plp-threshold** statements are mandatory.

You can define the following options for each RED profile:

- high-plp-max-threshold—Define the drop profile fill-level for the high packet loss priority (PLP) CoS VC. When the fill level exceeds the defined percentage, all packets with high PLP are dropped.
- low-plp-max-threshold—Define the drop profile fill-level for the low PLP CoS VC. When the fill level exceeds the defined percentage, all packets with low PLP are dropped.
- queue-depth—Define maximum queue depth in the CoS VC drop profile. Packets are always dropped beyond the defined maximum. The range you can configure is from 1 through 64,000 cells.
- high-plp-threshold—Define CoS VC drop profile fill-level percentage when linear RED is applied to cells with high PLP. When the fill level exceeds the defined percentage, packets with high PLP are randomly dropped by RED.
- low-plp-threshold—Define CoS VC drop profile fill-level percentage when linear RED is applied to cells with low PLP. When the fill level exceeds the defined percentage, packets with low PLP are randomly dropped by RED.

Configuring an ATM Scheduler Map

To define a scheduler map, you associate it with a forwarding class. Each class is associated with a specific queue, as follows:

- best-effort—Queue 0
- expedited-forwarding—Queue 1
- assured-forwarding—Queue 2
- network-control—Queue 3

NOTE: For M320 and T-series platforms only, you can configure more than four forwarding classes and queues. For more information, see “Enabling Eight Queues on ATM2 IQ Interfaces” on page 261.

When you configure an ATM scheduler map, the JUNOS software creates these CoS queues for a VC. The JUNOS software prefixes each packet delivered to the VC with the next-hop rewrite data associated with each queue.

To configure an ATM scheduler map, include the scheduler-maps statement at the [edit interfaces at-*fpc/pic/port* atm-options] hierarchy level:

```
[edit interfaces at-fpc/pic/port atm-options]
scheduler-maps map-name {
    forwarding-class class-name {
        epd-threshold cells plp1 cells;
        linear-red-profile profile-name;
        priority (high | low);
        transmit-weight (cells number | percent number);
    }
    vc-cos-mode (alternate | strict);
}
```

You can define the following options for each forwarding class:

- epd-threshold or linear-red-profile—An EPD threshold provides a queue of cells that can be stored with tail drop. When a beginning-of-packet (BOP) cell is received, the VC's queue depth is checked against the EPD threshold. If the VC's queue depth exceeds the EPD threshold, the BOP cell and all subsequent cells in the packet are discarded.

  A linear RED profile defines the number of cells using the queue-depth statement within the RED profile. (You configure the queue-depth statement at the [edit interfaces at-*fpc/pic/port* atm-options linear-red-profiles *profile-name*] hierarchy level.)

  By default, if you include the scheduler-maps statement at the [edit interfaces at-*fpc/pic/port* atm-options] hierarchy level, the interface uses an EPD threshold that is determined by the JUNOS software based on the available bandwidth and other parameters. You can override the default EPD threshold by setting an EPD threshold or a linear RED profile.

  If shaping is enabled, the default EPD threshold is proportional to the shaping rate according to the following formula:

  ```
  default epd-threshold = number of buffers * shaping rate / line rate
  ```

  The minimum value is 48 cells. If the formula results in an EPD threshold less than 48 cells, the result is ignored, and the minimum value of 48 cells is used.

- priority—By default, queue 0 is high priority, and the remaining queues are low priority. You can configure high or low queuing priority for each queue.

- transmit-weight—By default, the transmit weight is 95 percent for queue 0, and 5 percent for queue 3. You can configure the transmission weight in number of cells or percentage. Each CoS queue is serviced in WRR mode. When CoS queues have data to send, they send the number of cells equal to their weight before passing control to the next active CoS queue. This allows proportional bandwidth sharing between multiple CoS queues within a rate-shaped VC tunnel. A CoS queue can send from 1 through 32,000 cells or from 5 through 100 percent of queued traffic before passing control to the next active CoS queue within a VC tunnel.

  The AAL5 protocol prohibits cells from being interleaved on a VC; therefore, a complete packet is always sent. If a CoS queue sends more cells than its assigned weight because of the packet boundary, the deficit is carried over to the next time the queue is scheduled to transmit. If the queue is empty after the cells are sent, the deficit is waived, and the queue's assigned weight is reset.

NOTE: If you include the scheduler-maps statement at the [edit interfaces at-*fpc/pic/port* atm-options] hierarchy level, the epd-threshold statement at the [edit interfaces *interface-name* unit *logical-unit-number*] or [edit interfaces *interface-name* unit *logical-unit-number* address *address* family *family* multipoint-destination *address*] hierarchy level has no effect because either the default EPD threshold, the EPD threshold setting in the forwarding class, or the linear RED profile takes effect instead.

Enabling Eight Queues on ATM2 IQ Interfaces

By default, ATM2 IQ PICs on T-series and M320 platforms are restricted to a maximum of four egress queues per interface. You can enable eight egress queues on ATM2 IQ interfaces by including the max-queues-per-interface statement at the [edit chassis fpc *slot-number* pic *pic-number*] hierarchy level:

```
[edit chassis fpc slot-number pic pic-number]
max-queues-per-interface (4 | 8);
```

The numerical value can be 4 or 8.

If you include the max-queues-per-interface statement, all ports on the ATM2 IQ PIC use the configured maximum.

When you include the max-queues-per-interface statement and commit the configuration, all physical interfaces on the ATM2 IQ PIC are deleted and re-added. Also, the PIC is taken offline and then brought back online immediately. You do not need to manually take the PIC offline and online. You should change modes between four queues and eight queues only when there is no active traffic going to the ATM2 IQ PIC.

For general information about configuring up to eight forwarding classes and queues on PICs other than ATM2 IQ PICs, see “Configuring Up to Eight Forwarding Classes” on page 98.

NOTE: When you are considering enabling eight queues on an ATM2 IQ interface, you should note the following:

- ATM2 IQ interfaces using Layer 2 circuit trunk transport mode support only four CoS queues.
- ATM2 IQ OC48 interfaces support only four CoS queues.
- ATM2 IQ interfaces with MLPPP encapsulation support only four CoS queues.
- You can configure only four RED profiles for the eight queues. Thus, queue 0 and queue 4 share a single RED profile, as do queue 1 and queue 5, queue 2 and queue 6, and queue 3 and queue 7. There is no restriction on EPD threshold per queue.
- The default chassis scheduler allocates resources for queue 0 through queue 3, with 25 percent of the bandwidth allocated to each queue. When you configure the chassis to use more than four queues, you must configure and apply a custom chassis scheduler to override the default. To apply a custom chassis scheduler, include the scheduler-map-chassis statement at the [edit class-of-service interfaces at-*fpc/pic/**] hierarchy level. For more information about configuring and applying a custom chassis scheduler, see “Associating the Scheduler Map with the Packet Forwarding Component Queues” on page 164.

Example: Enabling Eight Queues on ATM2 IQ Interfaces

In Figure 15, Router A generates IP packets with different IP precedence settings. Router B is an M320 or T-series platform with two ATM2 IQ interfaces. On Router B, interface at-6/1/0 receives traffic from Router A, while interface at-0/1/0 sends traffic to Router C. This example shows the CoS configuration for Router B.

Figure 15: Example Topology for Router with Eight Queues

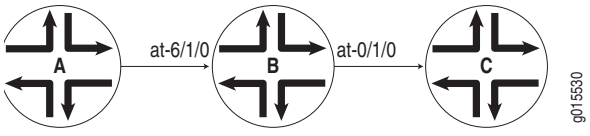

On Router B:

```
[edit chassis]
fpc 0 {
    pic 1 {
        max-queues-per-interface 8;
    }
}
fpc 6 {
    pic 1 {
        max-queues-per-interface 8;
    }
}

[edit interfaces]
at-0/1/0 {
    atm-options {
        linear-red-profiles {
            red_1 queue-depth 1k high-plp-threshold 50 low-plp-threshold 80;
            red_2 queue-depth 2k high-plp-threshold 40 low-plp-threshold 70;
            red_3 queue-depth 3k high-plp-threshold 30 low-plp-threshold 60;
            red_4 queue-depth 4k high-plp-threshold 20 low-plp-threshold 50;
        }
        scheduler-maps {
            sch_red {
                vc-cos-mode strict;
                forwarding-class fc_q0 {
                    priority high;
                    transmit-weight percent 5;
                    linear-red-profile red_1;
                }
                forwarding-class fc_q1 {
                    priority low;
                    transmit-weight percent 10;
                    linear-red-profile red_2;
                }
                forwarding-class fc_q2 {
                    priority low;
                    transmit-weight percent 15;
                    linear-red-profile red_3;
                }
                forwarding-class fc_q3 {
                    priority low;
                    transmit-weight percent 20;
                    linear-red-profile red_4;
                }
                forwarding-class fc_q4 {
                    priority low;
                    transmit-weight percent 5;
                    linear-red-profile red_1;
                }
                forwarding-class fc_q5 {
                    priority low;
                    transmit-weight percent 10;
                    linear-red-profile red_2;
                }
```

```
            forwarding-class fc_q6 {
                priority low;
                transmit-weight percent 15;
                linear-red-profile red_3;
            }
            forwarding-class fc_q7 {
                priority low;
                transmit-weight percent 20;
                linear-red-profile red_4;
            }
        }
        sch_epd {
            vc-cos-mode alternate;
            forwarding-class fc_q0 {
                priority high;
                transmit-weight percent 5;
                epd-threshold 1024;
            }
            forwarding-class fc_q1 {
                priority low;
                transmit-weight percent 10;
                epd-threshold 2048;
            }
            forwarding-class fc_q2 {
                priority low;
                transmit-weight percent 15;
                epd-threshold 3072;
            }
            forwarding-class fc_q3 {
                priority low;
                transmit-weight percent 20;
                epd-threshold 4096;
            }
            forwarding-class fc_q4 {
                priority low;
                transmit-weight percent 5;
                epd-threshold 2048;
            }
            forwarding-class fc_q5 {
                priority low;
                transmit-weight percent 10;
                epd-threshold 3072;
            }
            forwarding-class fc_q6 {
                priority low;
                transmit-weight percent 15;
                epd-threshold 4096;
            }
            forwarding-class fc_q7 {
                priority low;
                transmit-weight percent 20;
                epd-threshold 5120;
            }
        }
    }
}
```

```
        atm-options {
            vpi 0;
        }
        unit 0 {
            vci 0.100;
            shaping {
                cbr 1920000;
            }
            atm-scheduler-map sch_red;
            family inet {
                address 172.16.0.1/24;
            }
        }
        unit 1 {
            vci 0.101;
            shaping {
                vbr peak 1m sustained 384k burst 256;
            }
            atm-scheduler-map sch_epd;
            family inet {
                address 172.16.1.1/24;
            }
        }
    }
    at-6/1/0 {
        atm-options {
            vpi 0;
        }
        unit 0 {
            vci 0.100;
            family inet {
                address 10.10.0.1/24;
            }
        }
        unit 1 {
            vci 0.101;
            family inet {
                address 10.10.1.1/24;
            }
        }
    }
```

```
[edit class-of-service]
classifiers {
    inet-precedence inet_classifier {
        forwarding-class fc_q0 {
            loss-priority low code-points 000;
        }
        forwarding-class fc_q1 {
            loss-priority low code-points 001;
        }
        forwarding-class fc_q2 {
            loss-priority low code-points 010;
        }
        forwarding-class fc_q3 {
            loss-priority low code-points 011;
        }
        forwarding-class fc_q4 {
            loss-priority low code-points 100;
        }
        forwarding-class fc_q5 {
            loss-priority low code-points 101;
        }
        forwarding-class fc_q6 {
            loss-priority low code-points 110;
        }
        forwarding-class fc_q7 {
            loss-priority low code-points 111;
        }
    }
    forwarding-classes {
        queue 0 fc_q0;
        queue 1 fc_q1;
        queue 2 fc_q2;
        queue 3 fc_q3;
        queue 4 fc_q4;
        queue 5 fc_q5;
        queue 6 fc_q6;
        queue 7 fc_q7;
    }
    interfaces {
        at-6/1/0 {
            unit * {
                classifiers {
                    inet-precedence inet_classifier;
                }
            }
        }
    }
}
```

```
[edit routing-options]
static {
    route 10.10.20.2/32 {
        next-hop at-0/1/0.0;
        retain;
        no-readvertise;
    }
    route 10.10.1.2/32 {
        next-hop at-0/1/0.1;
        retain;
        no-readvertise;
    }
}
```

Verifying the Configuration

To see the results of this configuration, you can issue the following operational mode commands:

- show interfaces at-0/1/0 extensive
- show interfaces queue at-0/1/0
- show class-of-service forwarding-class

Configuring VC CoS Mode

VC CoS mode defines the CoS queue scheduling priority. By default, the VC CoS mode is alternate. When it is a queue's turn to transmit, the queue transmits up to its weight in cells as specified by the transmit-weight statement at the [edit interfaces at-*fpc/pic/port* atm-options scheduler-maps *map-name* forwarding-class *class-name*] hierarchy level. The number of cells transmitted can be slightly over the configured or default transmit weight, because the transmission always ends at a packet boundary.

To configure the VC CoS mode, include the vc-cos-mode statement at the [edit interfaces at-*fpc/pic/port* atm-options scheduler-maps] hierarchy level:

```
[edit interfaces at-fpc/pic/port atm-options scheduler-maps]
vc-cos-mode (alternate | strict);
```

Two modes of CoS scheduling priority are supported:

- alternate—Assign high priority to one queue. The scheduling of the queues alternates between the high priority queue and the remaining queues. Every other scheduled packet is from the high priority queue.
- strict—Assign strictly high priority to one queue. A queue with strictly high priority is always scheduled before the remaining queues. The remaining queues are scheduled in round-robin fashion.

Enabling the PLP Setting to Be Copied to the CLP Bit

For a provider-edge (PE) router with customer edge (CE)-facing, egress, ATM2 IQ interfaces configured with standard AAL5 encapsulation, you can enable the PLP setting to be copied into the CLP bit.

NOTE: This configuration setting is not applicable to Layer 2 circuit encapsulations because the control word captures and preserves CLP information. For more information about Layer 2 circuit encapsulations, see the *JUNOS Network Interfaces Configuration Guide*.

By default, at egress ATM2 IQ interfaces configured with standard AAL5 encapsulation, the PLP information is not copied to the CLP bit. This means the PLP information is not carried beyond the egress interface onto the CE router.

You can enable the PLP information to be copied into the CLP bit by including the plp-to-clp statement:

```
plp-to-clp;
```

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* atm-options]
- [edit interfaces *interface-name* unit *logical-unit-number*]
- [edit logical-routers *logical-router-name* interfaces *interface-name* unit *logical-unit-number*]

Configuring ATM CoS on the Logical Interface

To apply the ATM scheduler map to a logical interface, include the atm-scheduler-map statement:

```
atm-scheduler-map (map-name | default);
```

When you add or change a scheduler map, the associated logical interface is taken offline and then brought back online immediately. For ATM CoS to take effect, you must configure the VCI and VPI identifiers and traffic shaping on each VC by including the following statements:

```
vci vpi-identifier.vci-identifier;
shaping {
    (cbr rate | rtvbr peak rate sustained rate burst length |
        vbr peak rate sustained rate burst length);
}
```

You can include these statements at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number*]
- [edit logical-routers *logical-router-name* interfaces *interface-name* unit *logical-unit-number*]

For more information, see the *JUNOS Network Interfaces Configuration Guide*.

You can also apply a scheduler map to the chassis traffic that feeds the ATM interfaces. For more information, see "Associating the Scheduler Map with the Packet Forwarding Component Queues" on page 164.

Example: Configuring ATM2 IQ VC Tunnel CoS Components

Configure ATM2 IQ VC tunnel CoS components:

```
[edit interfaces]
at-1/2/0 {
    atm-options {
        vpi 0;
        linear-red-profiles red-profile-1 {
            queue-depth 35000 high-plp-threshold 75 low-plp-threshold 25;
        }
        scheduler-maps map-1 {
            vc-cos-mode strict;
            forwarding-class best-effort {
                priority low;
                transmit-weight percent 25;
                linear-red-profile red-profile-1;
            }
        }
    }
    unit 0 {
        vci 0.128;
        shaping {
            vbr peak 20m sustained 10m burst 20;
        }
        atm-scheduler-map map-1;
        family inet {
            address 192.168.0.100/32 {
                destination 192.168.0.101;
            }
        }
    }
}
```