Chapter 1

CoS Overview

When a network experiences congestion and delay, some packets must be dropped. JUNOS class of service (CoS) allows you to divide traffic into classes and offer various levels of throughput and packet loss when congestion occurs. This allows packet loss to happen according to rules that you configure.

For interfaces that carry IPv4, IPv6, and MPLS traffic, you can configure JUNOS CoS features to provide multiple classes of service for different applications. On the routing platform, you can configure multiple forwarding classes for transmitting packets, define which packets are placed into each output queue, schedule the transmission service level for each queue, and manage congestion using a random early detection (RED) algorithm.

The JUNOS CoS features provide a set of mechanisms that you can use to provide differentiated services when best-effort traffic delivery is insufficient. In designing CoS applications, you must give careful consideration to your service needs, and you must thoroughly plan and design your CoS configuration to ensure consistency across all routing platforms in a CoS domain. You must also consider all the routing platforms and other networking equipment in the CoS domain to ensure interoperability among all equipment.

Because Juniper Networks routing platforms implement CoS in hardware rather than in software, you can experiment with and deploy CoS features without adversely affecting packet forwarding and routing performance.

The standards are defined in the following RFCs:

- RFC 2474, *Definition of the Differentiated Services Field in the IPv4 and IPv6 Headers*
- RFC 2597, *Assured Forwarding PHB Group*
- RFC 2598, *An Expedited Forwarding PHB*
- RFC 2698, *A Two Rate Three Color Marker*

This chapter discusses the following topics:

For hardware capabilities and limitations, see Table 6 on page 30.

Packet Flow Across a Network

CoS works by examining traffic entering at the edge of your network. The edge routers classify traffic into defined service groups, which allow for the special treatment of traffic across the network. For example, voice traffic can be sent across certain links, and data traffic can use other links. In addition, the data traffic streams can be serviced differently along the network path to ensure that higher-paying customers receive better service. As the traffic leaves the network at the far edge, you can reclassify the traffic.

To support CoS, you must configure each router in the network. Generally, each router examines the packets that enter it to determine their CoS settings. These settings then dictate which packets are first transmitted to the next downstream router. In addition, the routers at the edges of the network might be required to alter the CoS settings of the packets that enter the network from the customer or peer networks.

In Figure 1, Router A is receiving traffic from a customer network. As each packet enters, Router A examines the packet's current CoS settings and classifies the traffic into one of the groupings defined by the Internet Service Provider (ISP). This definition allows Router A to prioritize its resources for servicing the traffic streams it is receiving. In addition, Router A might alter the CoS settings (forwarding class and loss priority) of the packets to better match the ISP's traffic groups. When Router B receives the packets, it examines the CoS settings, determines the appropriate traffic group, and processes the packet according to those settings. It then transmits the packets to Router C, which performs the same actions. Router D also examines the packets and determines the appropriate group. Because it sits at the far end of the network, the ISP might decide once again to alter the CoS settings of the packets before Router D transmits them to the neighboring network.

Figure 1: Packet Flow Across the Network

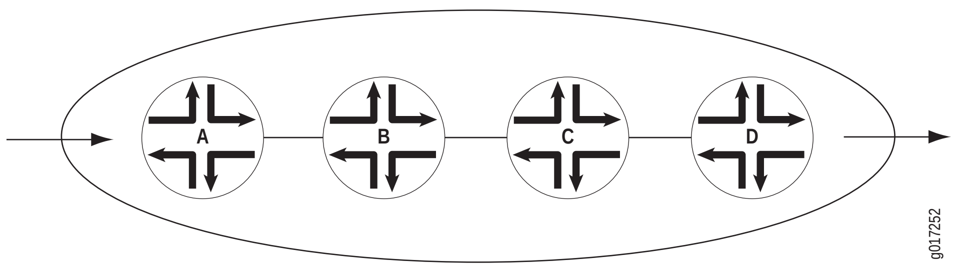

JUNOS CoS Components

The JUNOS CoS components include:

- Code-point aliases—A *code-point alias* assigns a name to a pattern of code-point bits. You can use this name, instead of the bit pattern, when you configure other CoS components, such as classifiers, drop-profile maps, and rewrite rules.

- Classifiers—*Packet classification* refers to the examination of an incoming packet. This function associates the packet with a particular CoS servicing level. In the JUNOS software, classifiers associate incoming packets with a forwarding class and loss priority and, based on the associated forwarding class, assign packets to output queues. Two general types of classifiers are supported:

  - Behavior aggregate or CoS value traffic classifiers—A *behavior aggregate* (BA) is a method of classification that operates on a packet as it enters the router. The CoS value in the packet header is examined, and this single field determines the CoS settings applied to the packet. BA classifiers allow you to set the forwarding class and loss priority of a packet based on the Differentiated Services code point (DSCP) value, DSCP IPv6 value, IP precedence value, MPLS EXP bits, and IEEE 802.1p value. The default classifier is based on the IP precedence value.

  - Multifield traffic classifiers—A *multifield* (MF) classifier is a second method for classifying traffic flows. Unlike a behavior aggregate, an MF classifier can examine multiple fields in the packet. Examples of some fields that an MF classifier can examine include the source and destination address of the packet as well as the source and destination port numbers of the packet. With MF classifiers, you set the forwarding class and loss priority of a packet based on firewall filter rules.

- Forwarding classes—Affect the forwarding, scheduling, and marking policies applied to packets as they transit a routing platform. The forwarding class plus the loss priority define the per-hop behavior. Four categories of forwarding classes are supported: best effort, assured forwarding, expedited forwarding, and network control. For M-series routing platforms, four forwarding classes are supported; you can configure up to one each of the four types of forwarding classes. For M320 and T-series platforms, 16 forwarding classes are supported, thus allowing you to classify packets more granularly. For example, you can configure multiple classes of EF traffic: EF, EF1, and EF2.

- Loss priorities—Allow you to set the priority of dropping a packet. Loss priority affects the scheduling of a packet without affecting the packet's relative ordering. You can use the packet loss priority (PLP) bit as part of a congestion control strategy. You can use the loss priority setting to identify packets that have experienced congestion. Typically you mark packets exceeding some service level with a high loss priority. You set loss priority by configuring a classifier or a policer. The loss priority is used later in the work flow to select one of the drop profiles used by RED.

- Forwarding policy options—Allow you to associate forwarding classes with next hops. Forwarding policy also allows you to create classification overrides, which assign forwarding classes to sets of prefixes.

- Transmission scheduling and rate control—Provide you with a variety of tools to manage traffic flows:

  - Queuing—After a packet is sent to the outgoing interface on a router, it is queued for transmission on the physical media. The amount of time a packet is queued on the router is determined by the availability of the outgoing physical media as well as the amount of traffic using the interface.

  - Schedulers—An individual router interface has multiple queues assigned to store packets. The router decides which queue to service based on a particular method of scheduling. This process often involves a determination of which type of packet should be transmitted before another. JUNOS schedulers allow you to define the priority, bandwidth, delay buffer size, rate control status, and RED drop profiles to be applied to a particular queue for packet transmission.

  - Fabric schedulers—For M320 and T-series platforms only, fabric schedulers allow you to identify a packet as high or low priority based on its forwarding class, and to associate schedulers with the fabric priorities.
  - Policers for traffic classes—*Policers* allow you to limit traffic of a certain class to a specified bandwidth and burst size. Packets exceeding the policer limits can be discarded, or can be assigned to a different forwarding class, a different loss priority, or both. You define policers with filters that can be associated with input or output interfaces.
- Rewrite rules—A *rewrite rule* sets the appropriate CoS bits in the outgoing packet. This allows the next downstream router to classify the packet into the appropriate service group. Rewriting or *marking* outbound packets is useful when the routing platform is at the border of a network and must alter the CoS values to meet the policies of the targeted peer.

Default CoS

If you do not configure any CoS settings on your routing platform, the software performs some CoS functions to ensure that user traffic and protocol packets are forwarded with minimum delay when the network is experiencing congestion. Some default mappings are automatically applied to each logical interface that you configure. Other default mappings, such as explicit default classifiers and rewrite rules, are in operation only if you explicitly associate them with an interface.

Each CoS component configuration chapter contains a section about default behavior. For more information, see the following sections:

- Default Behavior Aggregate Classification on page 45
- Default Forwarding Classes on page 80
- Default Drop Profile on page 107
- Default Schedulers on page 112
- Default Fabric Priority Queuing on page 162

You can display default CoS settings by issuing the show class-of-service operational mode command. This section includes sample output displaying the default CoS settings. The sample output is truncated for brevity.

show class-of-service

```
user@host> show class-of-service
```

Default Forwarding Classes

```
Forwarding class                    Queue
  best-effort                           0
  expedited-forwarding                  1
  assured-forwarding                    2
  network-control                       3
```

Default Code-Point Aliases

```
Code point type: dscp
  Alias              Bit pattern
  af11               001010
  af12               001100
...
Code point type: dscp-ipv6
...
Code point type: exp
...
Code point type: ieee-802.1
...
Code point type: inet-precedence
...
```

Default Classifiers

```
Classifier: dscp-default, Code point type: dscp, Index: 7
...
Classifier: dscp-ipv6-default, Code point type: dscp-ipv6, Index: 8
...
Classifier: exp-default, Code point type: exp, Index: 9
...

Classifier: ieee8021p-default, Code point type: ieee-802.1, Index: 10
...

Classifier: ipprec-default, Code point type: inet-precedence, Index: 11
...

Classifier: ipprec-compatibility, Code point type: inet-precedence,Index: 12
...
```

Default Frame Relay Loss Priority Map

```
Loss-priority-map: frame-relay-de-default, Code point type: frame-relay-de,
Index: 13
  Code point         Loss priority
  0                  low
  1                  high
```

Default Rewrite Rules

```
Rewrite rule: dscp-default, Code point type: dscp, Index: 24
  Forwarding class                 Loss priority      Code point
  best-effort                      low                000000
  best-effort                      high               000000
...

Rewrite rule: dscp-ipv6-default, Code point type: dscp-ipv6, Index: 25
...

Rewrite rule: exp-default, Code point type: exp, Index: 26
...

Rewrite rule: ieee8021p-default, Code point type: ieee-802.1, Index: 27
...

Rewrite rule: ipprec-default, Code point type: inet-precedence, Index: 28
...
```

Default Drop Profile

```
Drop profile: <default-drop-profile>, Type: discrete, Index: 1
  Fill level    Drop probability
         100                 100
```

Default Schedulers

```
Scheduler map: <default>, Index: 2

  Scheduler: <default-be>, Forwarding class: best-effort, Index: 17
    Transmit rate: 95 percent, Rate Limit: none, Buffer size: 95 percent,
Priority: low
    Drop profiles:
      Loss priority   Protocol    Index    Name
      Low             non-TCP         1    <default-drop-profile>
      Low             TCP             1    <default-drop-profile>
...
```

CoS Inputs and Outputs

Some CoS components map one set of values to another set of values. Each mapping contains one or more inputs and one or more outputs. When you configure a mapping, you set the outputs for a given set of inputs, as shown in Table 5.

Table 5: CoS Mappings–Inputs and Outputs

CoS Mappings	Inputs	Outputs	Comments
classifiers	code-points	forwarding-class loss-priority	The map sets the forwarding class and PLP for a specific set of code points. See “Classifying Packets by Behavior Aggregate” on page 43.
drop-profile-map	loss-priority protocol	drop-profile	The map sets the drop profile for a specific PLP and protocol type. See “Configuring the Scheduler Drop Profile Map” on page 113.
rewrite-rules	forwarding-class loss-priority	code-points	The map sets the code points for a specific forwarding class and PLP. See “Rewriting Packet Header Information” on page 207.

In the following classifier example, packets with EXP bits 000 are assigned to the data-queue forwarding-class with a low loss priority, and packets with EXP bits 001 are assigned to the data-queue forwarding-class with a high loss priority.

```
class-of-service {
    classifiers {
      exp exp_classifier {
        forwarding-class data-queue {
          loss-priority low code-points 000;
          loss-priority high code-points 001;
        }
      }
    }
}
```

In the following drop-profile map example, the scheduler includes two drop-profile maps, which specify that packets are evaluated by the low-drop drop profile if they have a low loss priority and are from any protocol. Packets are evaluated by the high-drop drop profile if they have a high loss priority and are from any protocol.

```
class-of-service {
    schedulers {
      best-effort {
        drop-profile-map loss-priority low protocol any drop-profile low-drop;
        drop-profile-map loss-priority high protocol any drop-profile high-drop;
      }
    }
}
```

In the following rewrite rule example, packets in the be forwarding class with low loss priority are assigned the EXP bits 000, and packets in the be forwarding class with high loss priority are assigned the EXP bits 001.

```
class-of-service {
    rewrite-rules {
      exp exp-rw {
        forwarding-class be {
          loss-priority low code-point 000;
          loss-priority high code-point 001;
        }
      }
    }
}
```

Packet Flow Within Routing Platforms

When a packet enters an M-series or T-series Juniper Networks router, the Physical Interface Card (PIC) receiving the packet retrieves it from the network and verifies that the link-layer information is valid. The packet is then passed to the Flexible PIC Concentrator (FPC), where the data-link and network layer information is verified. In addition, the FPC is responsible for segmenting the packet into 64-byte J-cells. These cells are then written into packet storage memory while a notification cell is sent to the route lookup engine. The destination address listed in the notification cell is located in the forwarding table, and the next hop of the packet is written into the result cell. This result cell is queued on the appropriate outbound FPC until the outgoing interface is ready to transmit the packet. The FPC then reads the J-cells out of memory, reforms the original packet, and sends the packet to the outgoing PIC, where it is transmitted back into the network.

Packet flow differs by platform type. This section discusses the following topics:

- Packet Flow on J-series Platforms on page 11
- Packet Flow on M-series Platforms on page 11
- Packet Flow on T-series Platforms on page 13.

Packet Flow on J-series Platforms

On J-series Services Routers, some of the hardware components associated with larger platforms are virtualized. These virtualized components include Packet Forwarding Engines, Routing Engines, and their associated ASICs. For this reason, packet flow on J-series platforms cannot be described in terms of discrete hardware components.

Packet Flow on M-series Platforms

On M-series platforms, CoS actions are performed in several locations in a Juniper Networks router: the incoming I/O Manager ASIC, the Internet Processor II ASIC, and the outgoing I/O Manager ASIC. These locations are shown in Figure 2.

Figure 2: M-series Packet Forwarding Engine Components and Data Flow

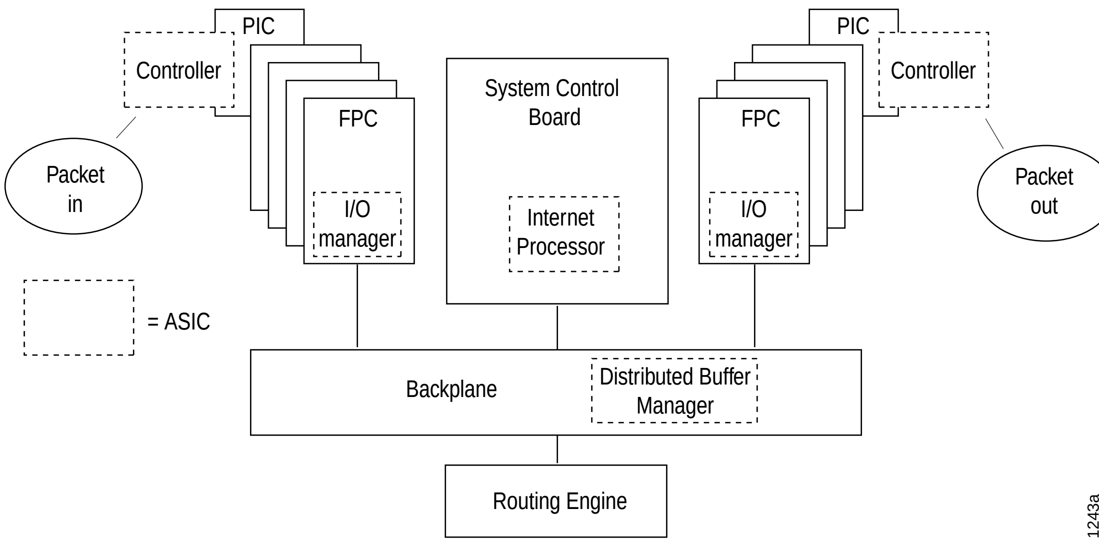

The following sections describe the packet flow in more detail:

- Incoming I/O Manager ASIC on page 11
- Internet Processor ASIC on page 12
- Outgoing I/O Manager ASIC on page 12

Incoming I/O Manager ASIC

When a data packet is passed from the receiving interface to its connected FPC, it is received by the I/O Manager ASIC on that specific FPC. During the processing of the packet by this ASIC, the information in the packet's header is examined by a behavior aggregate (BA) classifier. This classification action associates the packet with a particular forwarding class. In addition, the value of the packet's loss priority bit is set by this classifier. Both the forwarding class and loss priority information are placed into the notification cell, which is then transmitted to the Internet Processor II ASIC.

Internet Processor ASIC

The Internet Processor II ASIC receives notification cells representing inbound data packets and performs route lookups in the forwarding table. This lookup determines the outgoing interface on the router and the next-hop IP address for the data packet. While the packet is being processed by the Internet Processor II ASIC, it might also be evaluated by a firewall filter, which is configured on either the incoming or outgoing interface. This filter can perform the functions of an multifield (MF) classifier by matching on multiple elements within the packet and overwriting the forwarding class and/or loss priority settings within the notification cell. Once the route lookup and filter evaluations are complete, the notification cell, now called the result cell, is passed to the I/O manager ASIC on the FPC associated with the outgoing interface.

Outgoing I/O Manager ASIC

When the result cell is received by the I/O Manager ASIC, it is placed into a queue to await transmission on the physical media. The specific queue used by the ASIC is determined by the forwarding class associated with the data packet. The configuration of the queue itself helps determine the service the packet receives while in this queued state. This functionality guarantees that certain packets be serviced and transmitted before other packets. In addition, the queue settings and the packet’s loss priority setting determine which packets might be dropped from the network during periods of congestion.

In addition to queuing the packet, the outgoing I/O manager ASIC is responsible for ensuring that CoS bits in the packet’s header are correctly set before it is transmitted. This rewrite function helps the next downstream router perform its CoS function in the network.

Packet Flow on T-series Platforms

On T-series platforms, CoS actions are performed in several locations in a Juniper Networks router: the incoming and outgoing Switch Interface ASICs, the T-series Internet Processor, and the Queuing and Memory Interface ASICs. These locations are shown in Figure 3.

Figure 3: T-series Packet Forwarding Engine Components and Data Flow

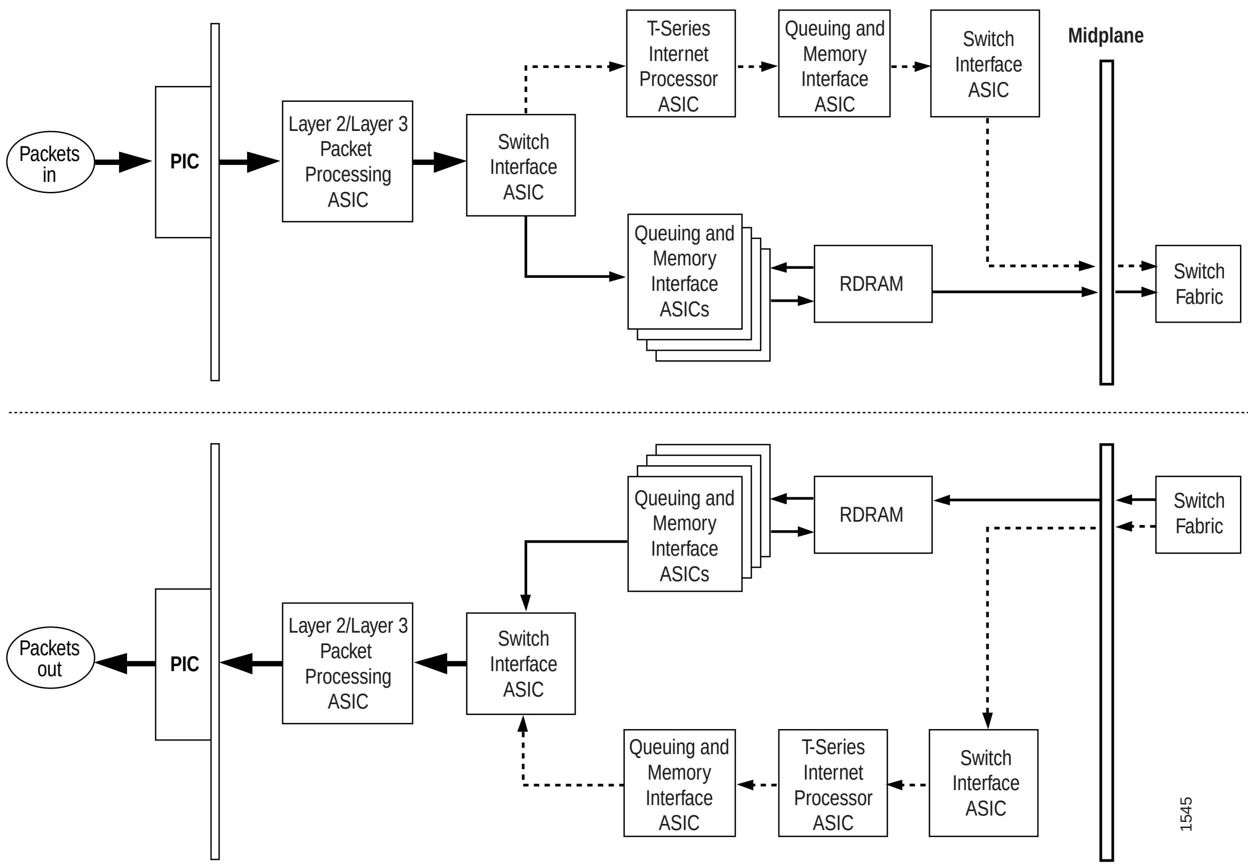

The following sections describe the packet flow in more detail:

- Incoming Switch Interface ASICs on page 14
- T-series Internet Processor ASIC on page 14
- Queuing and Memory Interface ASICs on page 14
- Outgoing Switch Interface ASICs on page 15

Incoming Switch Interface ASICs

When a data packet is passed from the receiving interface to its connected FPC, it is received by the incoming Switch Interface ASIC on that specific FPC. During the processing of the packet by this ASIC, the information in the packet's header is examined by a BA classifier. This classification action associates the packet with a particular forwarding class. In addition, the value of the packet's loss priority bit is set by this classifier. Both the forwarding class and loss priority information are placed into the notification cell, which is then transmitted to the T-series Internet Processor ASIC.

T-series Internet Processor ASIC

The T-series Internet Processor ASIC receives notification cells representing inbound data packets and performs route lookups in the forwarding table. This lookup determines the outgoing interface on the router and the next-hop IP address for the data packet. While the packet is being processed by the T-series Internet Processor ASIC, it might also be evaluated by a firewall filter, which is configured on either the incoming or outgoing interface. This filter can perform the functions of an MF classifier by matching on multiple elements within the packet and overwriting the forwarding class and/or loss priority settings within the notification cell. Once the route lookup and filter evaluations are complete, the notification cell, now called the result cell, is passed to the Queuing and Memory Interface ASICs.

Queuing and Memory Interface ASICs

The Queuing and Memory Interface ASICs pass the data cells to memory for buffering. The data cells are placed into a queue to await transmission on the physical media. The specific queue used by the ASICs is determined by the forwarding class associated with the data packet. The configuration of the queue itself helps determine the service the packet receives while in this queued state. This functionality guarantees that certain packets be serviced and transmitted before other packets. In addition, the queue settings and the packet's loss priority setting determine which packets might be dropped from the network during periods of congestion.

In addition to queuing the packet, the outgoing I/O manager ASIC is responsible for ensuring that CoS bits in the packet's header are correctly set before it is transmitted. This rewrite function helps the next downstream router perform its CoS function in the network.

The Queuing and Memory Interface ASIC sends the notification to the Switch Interface ASIC facing the switch fabric, unless the destination is on the same Packet Forwarding Engine. In this case, the notification is sent back to the Switch Interface ASIC facing the outgoing ports, and the packets are sent to the outgoing port without passing through the switch fabric. The default behavior is for fabric priority queuing on egress interfaces to match the scheduling priority you assign. High-priority egress traffic is automatically assigned to high-priority fabric queues.

The Queuing and Memory Interface ASIC forwards the notification, including next-hop information, to the outgoing Switch Interface ASIC.

Outgoing Switch Interface ASICs

The destination Switch Interface ASIC sends bandwidth grants through the switch fabric to the originating Switch Interface ASIC. The Queuing and Memory Interface ASIC forwards the notification, including next-hop information, to the Switch Interface ASIC. The Switch Interface ASIC sends read requests to the Queuing and Memory Interface ASIC to read the data cells out of memory, and passes the cells to the Layer 2/Layer 3 Packet Processing ASIC. The Layer 2/Layer 3 Packet Processing ASIC reassembles the data cells into packets, adds Layer 2 encapsulation, and sends the packets to the outgoing PIC interface. The outgoing PIC sends the packets out into the network.

Packet Flow Through the CoS Process

Perhaps the best way to understand JUNOS CoS is to examine how a packet is treated on its way through the CoS process. This section includes a description of each step, some figures illustrating the process, and a configuration example.

The following steps describe the CoS process:

1. A logical interface has one or more classifiers of different types applied to it (at the [edit class-of-service interfaces] hierarchy level). The types of classifiers are based on which part of the incoming packet the classifier examines (for example, EXP bits, IEEE 802.1p bits, or DSCP bits).

2. The classifier assigns the packet to a forwarding class and a loss-priority (at the [edit class-of-service classifiers] hierarchy level).

3. Each forwarding class is assigned to a queue (at the [edit class-of-service forwarding-classes] hierarchy level).

4. Input (and output) policers meter traffic and might change the forwarding class and loss priority if a traffic flow exceeds its service level.

5. The physical or logical interface has a scheduler map applied to it (at the [edit class-of-service interfaces] hierarchy level).

   At the [edit class-of-service interfaces] hierarchy level, the scheduler-map and rewrite-rules statements affect the outgoing packets, and the classifiers statement affects the incoming packets.

6. The scheduler defines how traffic is treated in the output queue—for example, the transmit rate, buffer size, priority, and drop-profile (at the [edit class-of-service schedulers] hierarchy level).

7. The scheduler map assigns a scheduler to each forwarding class (at the [edit class-of-service scheduler-maps] hierarchy level).

8. The drop-profile defines how aggressively to drop packets that are using a particular scheduler (at the [edit class-of-service drop-profiles] hierarchy level).

9. The rewrite rule takes effect as the packet leaves a logical interface that has a rewrite rule configured (at the [edit class-of-service rewrite-rules] hierarchy level). The rewrite rule writes information to the packet (for example, EXP or DSCP bits) according to the forwarding class and loss priority of the packet.

Figure 4 and Figure 5 show the components of the JUNOS CoS features, illustrating the sequence in which they interact.

Figure 4: CoS Classifier, Queues, and Scheduler

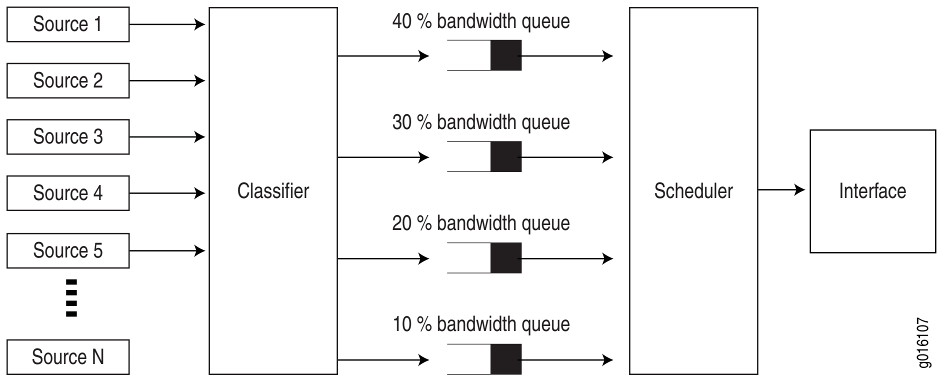

Figure 5: Packet Flow Through CoS Configurable Components

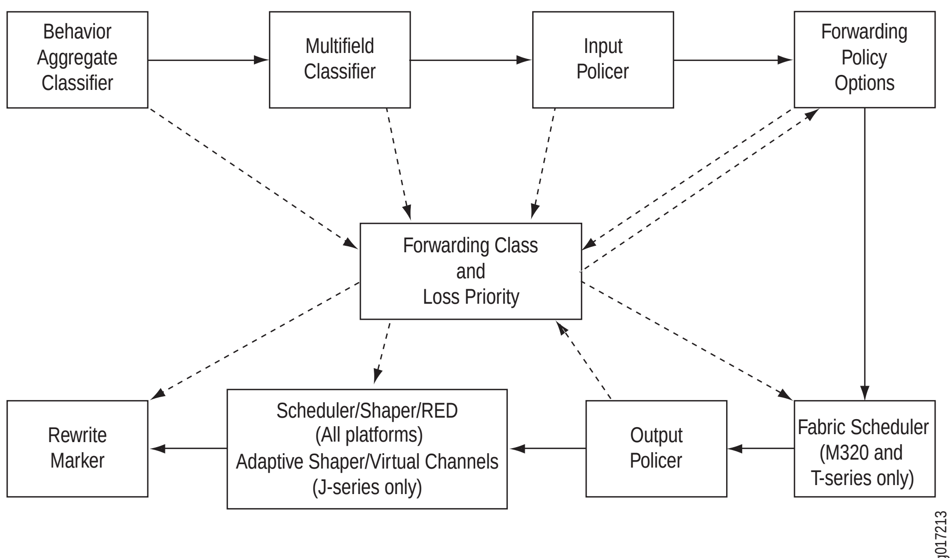

Each outer box in Figure 5 represents a process component. The components in the upper row apply to inbound packets, and the components in the lower row apply to outbound packets. The arrows with the solid lines point in the direction of packet flow.

The middle box (forwarding class and loss priority) represents two data values that can either be inputs to or outputs of the process components. The arrows with the dotted lines indicate inputs and outputs (or settings and actions based on settings). For example, the multifield classifier sets the forwarding class and loss priority of incoming packets. This means that the forwarding class and loss priory are outputs of the classifier; thus, the arrow points away from the classifier. The scheduler receives the forwarding class and loss priority settings, and queues the outgoing packet based on those settings. This means that the forwarding class and loss priority are inputs to the scheduler; thus, the arrow points to the scheduler. For more information, see "CoS Inputs and Outputs" on page 9.

Typically, only a combination of some components (not all) is used to define a CoS service offering.

The following configuration demonstrates the concepts described in this section:

```
class-of-service {
    interfaces {
        so-* {
            scheduler-map sched1;
            unit 0 {
                classifiers {
                    exp exp_classifier;
                }
            }
        }
        t3-* {
            scheduler-map sched1;
            unit 0 {
                classifiers {
                    exp exp_classifier;
                }
            }
        }
    }
    classifiers { # Step 2: Define classifiers.
        exp exp_classifier {
            forwarding-class data-queue {
                loss-priority low code-points 000;
                loss-priority high code-points 001;
            }
            forwarding-class video-queue {
                loss-priority low code-points 010;
                loss-priority high code-points 011;
            }
            forwarding-class voice-queue {
                loss-priority low code-points 100;
                loss-priority high code-points 101;
            }
            forwarding-class nc-queue {
                loss-priority high code-points 111;
                loss-priority low code-points 110;
            }
        }
    }
```

```
        drop-profiles {
            be-red {
                fill-level 50 drop-probability 100;
            }
        }
        forwarding-classes {
            queue 0 data-queue;
            queue 1 video-queue;
            queue 2 voice-queue;
            queue 3 nc-queue;
        }
        schedulers {
            data-scheduler {
                transmit-rate percent 50;
                buffer-size percent 50;
                priority low;
                drop-profile-map loss-priority high protocol any drop-profile be-red;
            }
            video-scheduler {
                transmit-rate percent 25;
                buffer-size percent 25;
                priority strict-high;
            }
            voice-scheduler {
                transmit-rate percent 20;
                buffer-size percent 20;
                priority high;
            }
            nc-scheduler {
                transmit-rate percent 5;
                buffer-size percent 5;
                priority high;
            }
        }
        scheduler-maps {
            sched1 {
                forwarding-class data-queue scheduler data-scheduler;
                forwarding-class video-queue scheduler video-scheduler;
                forwarding-class voice-queue scheduler voice-scheduler;
                forwarding-class nc-queue scheduler nc-scheduler;
            }
        }
    }
```

CoS Applications

You can configure CoS features to meet your application needs. Because the components are generic, you can use a single CoS configuration syntax across multiple platforms. CoS mechanisms are useful for two broad classes of applications. These applications can be referred to as *in the box* and *across the network*.

In-the-box applications use CoS mechanisms to provide special treatment for packets passing through a single node on the network. You can monitor the incoming traffic on each interface, using CoS to provide preferred service to some interfaces (that is, to some customers) while limiting the service provided to other interfaces. You can also filter outgoing traffic by the packet's destination, thus providing preferred service to some destinations.

Across-the-network applications use CoS mechanisms to provide differentiated treatment to different classes of packets across a set of nodes in a network. In these types of applications, you typically control the ingress and egress routing platforms to a routing domain and all the routing platforms within the domain. You can use JUNOS CoS features to modify packets traveling through the domain to indicate the packet's priority across the domain.

Specifically, you modify the CoS code points in packet headers, remapping these bits to values that correspond to levels of service. When all routing platforms in the domain are configured to associate the precedence bits with specific service levels, packets traveling across the domain receive the same level of service from the ingress point to the egress point. For CoS to work in this case, the mapping between the precedence bits and service levels must be identical across all routing platforms in the domain.

JUNOS CoS applications support the following range of mechanisms:

- Differentiated Services (DiffServ)—The CoS application supports DiffServ, which uses six-bit IPv4 and IPv6 header type-of-service (ToS) byte settings. The configuration uses CoS values in the IP and IPv6 ToS fields to determine the forwarding class associated with each packet.

- Layer 2 to Layer 3 CoS mapping—The CoS application supports mapping of Layer 2 (IEEE 802.1p) packet headers to routing platform forwarding class and loss-priority values.

  Layer 2 to Layer 3 CoS mapping involves setting the forwarding class and loss priority based on information in the Layer 2 header. Output involves mapping the forwarding class and loss priority to a Layer 2-specific marking. You can mark the Layer 2 and Layer 3 headers simultaneously.

- MPLS EXP—Supports configuration of mapping of MPLS experimental (EXP) bit settings to routing platform forwarding classes and vice versa.

- VPN outer-label marking—Supports setting of outer-label EXP bits, also known as CoS bits, based on MPLS EXP mapping.

Interface Types That Do Not Support CoS

You can configure CoS on all interfaces, except the following:

- cau4—Channelized STM1 IQ interface (configured on the Channelized STM1 IQ PIC).
- coc1—Channelized OC1 IQ interface (configured on the Channelized OC12 IQ PIC).
- coc12—Channelized OC12 IQ interface (configured on the Channelized OC12 IQ PIC).
- cstm-1—Channelized STM1 IQ interface (configured on the Channelized STM1 IQ PIC).
- ct1—Channelized T1 IQ interface (configured on the Channelized DS3 IQ PIC or Channelized OC12 IQ PIC).
- ct3—Channelized T3 IQ interface (configured on the Channelized DS3 IQ PIC or Channelized OC12 IQ PIC).
- ce1—Channelized E1 IQ interface (configured on the Channelized E1 IQ PIC or Channelized STM1 IQ PIC).
- dsc—Discard interface.
- fxp—Management and internal Ethernet interfaces.
- lo—Loopback interface. This interface is internally generated.
- pe—Encapsulates packets destined for the rendezvous point routing platform. This interface is present on the first-hop routing platform.
- pd—De-encapsulates packets at the rendezvous point. This interface is present on the rendezvous point.
- vt—Virtual loopback tunnel interface.

NOTE: For channelized interfaces, you can configure CoS on channels, but not at the controller level. For a complex configuration example, see the *JUNOS Feature Guide*.

For original Channelized OC12 PICs, limited CoS functionality is supported. For more information, contact Juniper Networks customer support.

The standard JUNOS CoS hierarchy is not supported on ATM interfaces. ATM has traffic-shaping capabilities that would override CoS, because ATM traffic shaping is performed at the ATM layer and CoS is performed at the IP layer. For more information about ATM traffic shaping and ATM CoS components, see "Configuring CoS on ATM Interfaces" on page 239 and the *JUNOS Network Interfaces Configuration Guide*.