

Chapter 5

Classifying Packets by Behavior Aggregate

The behavior aggregate (BA) classifier maps a class-of-service (CoS) value to a forwarding class and loss priority. The forwarding class determines the output queue. The loss priority is used by schedulers in conjunction with the random early discard (RED) algorithm to control packet discard during periods of congestion.

The types of BA classifiers are based on which part of the incoming packet the classifier examines:

- Differentiated Services code point (DSCP) for IP DiffServ
- DSCP for IPv6 DiffServ
- IP precedence bits
- MPLS EXP bits
- IEEE 802.1p CoS bits

Unlike multifield (MF) classifiers (which are discussed in “Classifying Packets Based on Various Packet Header Fields” on page 61), BA classifiers are based on fixed-length fields, which makes them computationally more efficient than MF classifiers. For this reason, core devices are normally configured to perform BA classification, because of the higher traffic volumes they handle.

In most cases, you need to rewrite a given marker (IP precedence, DSCP, IEEE 802.1P, or MPLS EXP settings) at the ingress node to accommodate BA classification by core and egress devices. For more information about rewrite markers, see “Rewriting Packet Header Information” on page 207.

For M-series routing platforms, four classes can forward traffic independently. For M320 and T-series platforms, eight classes can forward traffic independently. Therefore, you must configure additional classes to be aggregated into one of these classes. You use the BA classifier to configure class aggregation.



NOTE: For a specified interface, you can configure both an multifold (MF) classifier and a BA classifier without conflicts. Because the classifiers are always applied in sequential order, the BA classifier followed by the MF classifier, any BA classification result is overridden by an MF classifier if they conflict. For more information about MF classifiers, see “Classifying Packets Based on Various Packet Header Fields” on page 61.

To configure BA classifiers, you can include the following statements at the [edit class-of-service] hierarchy level of the configuration:

```
class-of-service {
  classifiers {
    (dscp | dscp-ipv6 | exp | ieee-802.1 | inet-precedence) classifier-name {
      import (classifier-name | default);
      forwarding-class class-name {
        loss-priority level {
          code-points [ aliases ] [ 6-bit-patterns ];
        }
      }
    }
  }
  interfaces {
    interface-name {
      unit logical-unit-number {
        classifiers {
          (dscp | dscp-ipv6 | exp | ieee-802.1 | inet-precedence) (classifier-name
            | default);
        }
      }
    }
  }
  routing-instances routing-instance-name {
    classifiers {
      exp (classifier-name | default);
    }
  }
}
```

This chapter discusses the following topics:

- Classifier Types on page 45
- Default Behavior Aggregate Classification on page 45
- Defining Classifiers on page 49
- Applying a Classifier to a Logical Interface on page 50
- Applying DSCP IPv6 Classifiers on page 52

- Applying MPLS EXP Classifiers to Routing Instances on page 52
- Applying MPLS EXP Classifiers for Explicit-Null Labels on page 57
- Setting the PLP on T320 and M320 Platforms on page 57
- Classifying Frame Relay Traffic on page 58

Classifier Types

The simplest way to classify a packet is to use behavior aggregate classification. The DSCP, DSCP IPv6, or IP precedence bits of the IP header convey the behavior aggregate class information. The information might also be found in the MPLS EXP bits or IEEE 802.1p CoS bits.

You can configure the following classifier types:

- DSCP, DSCP IPv6, or IP precedence—IP packet classification (Layer 3 headers)
- MPLS EXP—MPLS packet classification (Layer 2 headers)
- IEEE 802.1p—Packet classification (Layer 2 headers)

If you apply an IEEE 802.1 classifier to a logical interface, this classifier takes precedence and is not compatible with any other classifier type. Classifiers for IP (DSCP or IP precedence) and MPLS (EXP) can coexist on a logical interface if the hardware platform requirements are met. (See Table 14 on page 51.)

Default Behavior Aggregate Classification

The software automatically assigns an implicit default IP precedence classifier to all logical interfaces.

If you enable the MPLS protocol family on a logical interface, a default MPLS EXP classifier is automatically applied to that logical interface.

Other default classifiers (such as those for IEEE 802.1p bits and DSCP) require you to explicitly associate a default classification table with a logical interface. When you explicitly associate a default classifier with a logical interface, you are in effect overriding the implicit default classifier with an explicit default classifier.

The following sections describe the implicit and explicit default BA classifiers:

- Default IP Precedence Classifier (ipprec-compatibility) on page 46
- Default MPLS EXP Classifier on page 47
- Default DSCP and DSCP IPv6 Classifier on page 47

- Default IEEE 802.1p Classifier on page 48
- Default IP Precedence Classifier (ipprec-default) on page 48



NOTE: Although several code points map to the expedited-forwarding (ef) and assured-forwarding (af) classes, by default no resources are assigned to these forwarding classes. All af classes other than af1x are mapped to best-effort, because RFC 2597, *Assured Forwarding PHB Group*, prohibits a node from aggregating classes.

Default IP Precedence Classifier (ipprec-compatibility)

By default, all logical interfaces are automatically assigned an implicit IP precedence classifier called `ipprec-compatibility`. The `ipprec-compatibility` IP precedence classifier maps IP precedence bits to forwarding classes and loss priorities, as shown in Table 9.

Table 9: Default IP Precedence Classifier

IP Precedence CoS Values	Forwarding Class	Loss Priority
000	best-effort	low
001	best-effort	high
010	best-effort	low
011	best-effort	high
100	best-effort	low
101	best-effort	high
110	network-control	low
111	network-control	high

Default MPLS EXP Classifier

For all PICs except PICs mounted on M-series standard (non-enhanced) FPCs, if you enable the MPLS protocol family on a logical interface, the default MPLS EXP classifier is automatically applied to that logical interface. The default MPLS classifier maps EXP bits to forwarding classes and loss priorities, as shown in Table 10.

Table 10: Default MPLS Classifier

Code Point	Forwarding Class	Loss Priority
000	best-effort	low
001	best-effort	high
010	expedited-forwarding	low
011	expedited-forwarding	high
100	assured-forwarding	low
101	assured-forwarding	high
110	network-control	low
111	network-control	high

Default DSCP and DSCP IPv6 Classifier

Table 11 shows the forwarding class and packet loss priority (PLP) that are assigned to each well-known DSCP when you apply the explicit default DSCP or DSCP IPv6 classifier. To do this, include the default statement at the [edit class-of-service interfaces *interface-name* unit *logical-unit-number* classifiers (dscp | dscp-ipv6)] hierarchy level:

```
[edit class-of-service interfaces interface-name unit logical-unit-number classifiers
(dscp | dscp-ipv6)]
default;
```

Table 11: Default DSCP Classifier (1 of 2)

DSCP and DSCP IPv6	Forwarding Class	PLP
ef	expedited-forwarding	low
af11	assured-forwarding	low
af12	assured forwarding	high
af13	assured forwarding	high
af21	best-effort	low
af22	best-effort	low
af23	best-effort	low
af31	best-effort	low
af32	best-effort	low
af33	best-effort	low
af41	best-effort	low
af42	best-effort	low
af43	best-effort	low
be	best-effort	low

Table 11: Default DSCP Classifier (2 of 2)

DSCP and DSCP IPv6	Forwarding Class	PLP
cs1	best-effort	low
cs2	best-effort	low
cs3	best-effort	low
cs4	best-effort	low
cs5	best-effort	low
nc1/cs6	network-control	low
nc2/cs7	network control	low
other	best-effort	low

Default IEEE 802.1p Classifier

Table 12 shows the forwarding class and PLP that are assigned to the IEEE 802.1p CoS bits when you apply the explicit default IEEE 802.1p classifier. To do this, include the `default` statement at the [edit class-of-service interfaces *interface-name* unit *logical-unit-number* classifiers *ieee-802.1*] hierarchy level:

```
[edit class-of-service interfaces interface-name unit logical-unit-number classifiers
  ieee-802.1]
  default;
```

Table 12: Default IEEE 802.1p Classifier

Code Point	Forwarding Class	PLP
000	best-effort	low
001	best-effort	high
010	expedited-forwarding	low
011	expedited-forwarding	high
100	assured-forwarding	low
101	assured-forwarding	high
110	network-control	low
111	network-control	high

Default IP Precedence Classifier (*ipprec-default*)

There are two separate tables for default IP precedence classification. All logical interfaces are implicitly assigned the `ipprec-compatibility` classifier by default, as shown in Table 9 on page 46.

The other default IP precedence classifier (called `ipprec-default`) overrides the `ipprec-compatibility` classifier when you explicitly associate it with a logical interface. To do this, include the `default` statement at the [edit class-of-service interfaces *interface-name* unit *logical-unit-number* classifiers *inet-precedence*] hierarchy level:

```
[edit class-of-service interfaces interface-name unit logical-unit-number classifiers
  inet-precedence]
  default;
```

Table 13 shows the forwarding class and PLP that are assigned to the IP precedence CoS bits when you apply the default the IP precedence classifier.

Table 13: Default IP Precedence (ipprec-default) Classifier

Code Point	Forwarding Class	PLP
000	best-effort	low
001	assured-forwarding	low
010	best-effort	low
011	best-effort	low
100	best-effort	low
101	expedited-forwarding	low
110	network-control	low
111	network-control	high

Defining Classifiers

You can override the default IP precedence classifier by defining a classifier and applying it to a logical interface. To define new classifiers for all code-point types, include the `classifiers` statement at the `[edit class-of-service]` hierarchy level:

```
[edit class-of-service]
classifiers {
  (dscp | dscp-ipv6 | exp | ieee-802.1 | inet-precedence) classifier-name {
    import [classifier-name | default];
    forwarding-class class-name {
      loss-priority level {
        code-points [ aliases ] [ 6-bit-patterns ];
      }
    }
  }
}
```

The map sets the forwarding class and PLP for a specific set of code-point aliases and bit patterns. The inputs of the map are code-point aliases and bit patterns. The outputs of the map are the forwarding class and the PLP. For more information about how CoS maps work, see Table 5 on page 9.

The classifiers work as follows:

- `dscp`—Handles incoming IPv4 packets.
- `dscp-ipv6`—Handles incoming IPv6 packets. For more information, see “Applying DSCP IPv6 Classifiers” on page 52.
- `exp`—Handles MPLS packets using Layer 2 headers.
- `ieee-802.1`—Handles Layer 2 CoS.
- `inet-precedence`—Handles incoming IPv4 packets. IP precedence mapping requires only the upper three bits of the DSCP field.

A classifier takes a specified bit pattern as either the literal pattern or as a defined alias and attempts to match it to the type of packet arriving on the interface. If the information in the packet's header matches the specified pattern, the packet is sent to the appropriate queue, defined by the forwarding class associated with the classifier.

The code-point aliases and bit patterns are the input for the map. The loss priority and forwarding class are outputs of the map. In other words, the map sets the PLP and forwarding class for a given set of code-point aliases and bit patterns.



NOTE: On T-series and M320 platforms that do not have tricolor marking enabled, the loss priority can be configured only by setting the PLP within an MF classifier. This setting can then be used by the appropriate drop profile map and rewrite rule. For more information, see “Setting the PLP on T320 and M320 Platforms” on page 57.

Importing a Classifier

You can use any table, including the default, in the definition of a new classifier by including the `import` statement. The imported classifier is used as a template and is not modified. Whenever you commit a configuration that assigns a new `class-name` and `loss-priority` value to a code-point alias or set of bits, it replaces that entry in the imported classifier template. As a result, you must explicitly specify every CoS value in every designation that requires modification.

To do this, include the `import default` statement at the `[edit class-of-service classifiers type classifier-name]` hierarchy level:

```
[edit class-of-service classifiers type classifier-name]
import default;
```

For instance, to import the default DSCP classifier, include the `dscp default` statement at the `[edit class-of-service classifiers dscp classifier-name]` hierarchy level:

```
[edit class-of-service classifiers dscp classifier-name]
import default;
```

Applying a Classifier to a Logical Interface

You can apply the classification map to a logical interface by including the `classifiers` statement at the `[edit class-of-service interfaces interface-name unit logical-unit-number]` hierarchy level:

```
[edit class-of-service interfaces interface-name unit logical-unit-number]
classifiers (dscp | dscp-ipv6 | exp | ieee-802.1 | inet-precedence) (classifier-name
| default);
```

You can use interface wildcards for `interface-name` and `logical-unit-number`.

For most PICs, if you apply an IEEE 802.1p classifier to a logical interface, you cannot apply non-IEEE classifiers to other logical interfaces on the same physical interface. This restriction does not apply to Gigabit Ethernet IQ2 PICs.

There are some restrictions on applying multiple BA classifiers to a single logical interface. Table 14 shows the supported combinations, by platform.

Table 14: Logical Interface Classifier Combinations by Platform

Classifier Combinations	Gigabit Ethernet IQ2 PICs	Other PICs on T-series and M320	Other M-series with Regular FPCs	Other M-series with Enhanced FPCs
dscp and inet-precedence	No	No	No	No
dscp-ipv6 and (dscp inet-precedence)	Yes	Yes	No	No
exp and ieee 802.1	Yes	No	No	No
ieee 802.1 and (dscp dscp-ipv6 exp inet-precedence)	Yes	No	No	Yes
exp and (dscp dscp-ipv6 inet-precedence)	Yes	Yes	No	Yes

For Gigabit Ethernet IQ2 interfaces, IEEE 802.1p classifiers are evaluated after other BA classifiers. For example, if you configure a logical interface to use both an MPLS EXP and an IEEE 802.1p classifier, the EXP classifier takes precedence. MPLS-labeled packets are evaluated by the EXP classifier, and all other packets are evaluated by the IEEE 802.1p classifier. The same is true about other classifiers when combined with IEEE 802.1p classifiers on the same logical interface.



NOTE: If an interface is mounted on an M-series FPC, you can apply to the interface the default `exp` classifier only. If an interface is mounted on an enhanced FPC, you can create a new `exp` classifier and apply it to an interface.

Applying DSCP IPv6 Classifiers

For M320 and T-series platforms, you can apply separate classifiers for IPv4 and IPv6 packets per logical interface by including the `classifiers` statement at the `[edit class-of-service interfaces interface-name unit logical-unit-number]` hierarchy level and specifying the `dscp` and `dscp-ipv6` classifier types:

```
[edit class-of-service interfaces interface-name unit logical-unit-number]
classifiers dscp (classifier-name | default);
classifiers dscp-ipv6 (classifier-name | default);
```

For M-series enhanced FPCs, you cannot apply separate classifiers for IPv4 and IPv6 packets on a single logical interface. Instead, classifier assignment works as follows:

- If you assign a DSCP classifier only, IPv4 and IPv6 packets are classified using the DSCP classifier.
- If you assign an IP precedence classifier only, IPv4 and IPv6 packets are classified using the IP precedence classifier. In this case, the lower three bits of the DSCP field are ignored because IP precedence mapping requires the upper three bits only.
- If you assign either the DSCP or the IP precedence classifier in conjunction with the DSCP IPv6 classifier, the commit fails.
- If you assign a DSCP IPv6 classifier only, IPv4 and IPv6 packets are classified using the DSCP IPv6 classifier, but the commit displays a warning message.

For more information, see Table 14 on page 51. For a complex configuration example, see the *JUNOS Feature Guide*.

Applying MPLS EXP Classifiers to Routing Instances

When you enable VRF table labels and you do not explicitly apply a classifier configuration to the routing instance, the default MPLS EXP classifier is applied to the routing instance. For detailed information about VRF table labels, see the *JUNOS VPNs Configuration Guide*.

The default MPLS EXP classification table contents are shown in Table 15.

Table 15: Default MPLS EXP Classification Table

Forwarding Class	Loss Priority	CoS Value
best-effort	low	000
best-effort	high	001
expedited-forwarding	low	010
expedited-forwarding	high	011
assured-forwarding	low	100
assured-forwarding	high	101
network-control	low	110
network-control	high	111

For PICs that are installed on enhanced FPCs, you can override the default MPLS EXP classifier and apply a custom classifier to the routing instance. To do this, perform the following configuration tasks:

1. Filter traffic based on the IP header by including the `vrf-table-label` statement at the `[edit routing-instances routing-instance-name]` hierarchy level:

```
[edit routing-instances routing-instance-name]  
vrf-table-label;
```

2. Configure a custom MPLS EXP classifier by including the following statements at the `[edit class-of-service]` hierarchy level:

```
[edit class-of-service]  
classifiers {  
  exp classifier-name {  
    import (classifier-name | default);  
    forwarding-class class-name {  
      loss-priority level {  
        code-points [ aliases ] [ 6-bit-patterns ];  
      }  
    }  
  }  
}  
forwarding-classes {  
  queue queue-number class-name priority (high | low);  
}
```

3. Configure the routing instance to use the custom MPLS EXP classifier by including the `exp` statement at the `[edit class-of-service routing-instances routing-instance-name classifiers]` hierarchy level:

```
[edit class-of-service routing-instances routing-instance-name classifiers]  
exp classifier-name;
```

To display the MPLS EXP classifiers associated with all routing instances, issue the `show class-of-service routing-instances` command.



NOTE: The following caveats apply to custom MPLS EXP classifiers for routing instances:

- An enhanced FPC is required.
 - Logical routers are not supported.
-

For more details, see the following sections:

- [Configuring Global Classifiers and Wildcard Routing Instances on page 54](#)
- [Examples: Applying MPLS EXP Classifiers to Routing Instances on page 55](#)

Configuring Global Classifiers and Wildcard Routing Instances

To configure a global routing instance classifier, include the `all` statement at the `[edit class-of-service routing-instances]` hierarchy level:

```
[edit class-of-service routing-instances]
all {
  classifiers {
    exp classifier-name;
  }
}
```

For routing instances associated with specific classifiers, the global configuration is ignored.

To use a wildcard in the routing instance classifier configuration, include an asterisk (*) in the name of the routing instance:

```
[edit class-of-service routing-instances]
instance-name* {
  classifiers {
    exp classifier-name;
  }
}
```

The wildcard configuration follows the longest match. If there is a specific configuration, it is given precedence over the wildcard configuration.



NOTE: Wildcards and the `all` keyword are supported at the `[edit class-of-service routing-instances]` hierarchy level but not at the `[edit routing-instances]` hierarchy level.

If you configure a routing instance at the `[edit routing-instances]` hierarchy level with, for example, the name `vpn*`, the JUNOS software treats `vpn*` as a valid and distinct routing instance name. If you then try to apply a classifier to the `vpn*` routing instance at the `[edit class-of-service routing-instances]` hierarchy level, the JUNOS software treats the `vpn*` routing instance name as a wildcard, and all the routing instances that start with `vpn` and do not have a specific classifier applied receive the classifier associated with `vpn*`. This same behavior applies with the `all` keyword.

Examples: Applying MPLS EXP Classifiers to Routing Instances

Configuring a Global Classifier

Configure a global classifier for all routing instances and override the global classifier for a specific routing instance. In this example, there are three routing instances: `vpn1`, `vpn2`, and `vpn3`, each with VRF table label enabled. The classifier `exp-classifier-global` is applied to `vpn1` and `vpn2`. The classifier `exp-classifier-3` is applied to `vpn3`.

```
[edit routing-instances]
vpn1 {
  vrf-table-label;
}
vpn2 {
  vrf-table-label;
}
vpn3 {
  vrf-table-label;
}

[edit class-of-service routing-instances]
all {
  classifiers {
    exp exp-classifier-global;
  }
}
vpn3 {
  classifiers {
    exp exp-classifier-3;
  }
}
```

Configuring a Wildcard Routing Instance

Configure a wildcard routing instance and override the wildcard with a specific routing instance. In this example, there are three routing instances: `vpn-red`, `vpn-yellow`, and `vpn-green`, each with VRF table label enabled. The classifier `exp-class-wildcard` is applied to `vpn-yellow` and `vpn-green`. The classifier `exp-class-red` is applied to `vpn-red`.

```
[edit routing-instances]
vpn-red {
  vrf-table-label;
}
vpn-yellow {
  vrf-table-label;
}
vpn-green {
  vrf-table-label;
}

[edit class-of-service routing-instances]
vpn* {
  classifiers {
    exp exp-class-wildcard;
  }
}
vpn-red {
  classifiers {
    exp exp-class-red;
  }
}
```

Monitoring a Configuration

Display the MPLS EXP classifiers associated with two routing instances:

```
[edit class-of-service routing-instances]
vpn1 {
  classifiers {
    exp default;
  }
}
vpn2 {
  classifiers {
    exp class2;
  }
}
```

`user@router>show class-of-service routing-instances`

```
Routing Instance : vpn1
  Object      Name           Type      Index
  Classifier   exp-default    exp       8

Routing Instance : vpn2
  Object      Name           Type      Index
  Classifier   class2         exp       57507
```

Applying MPLS EXP Classifiers for Explicit-Null Labels

When you configure MPLS explicit-null labels, label 0 is advertised to the egress router of an LSP. When label 0 is advertised, the egress router (instead of the penultimate router) removes the label. Ultimate-hop popping ensures that any packets traversing an MPLS network include a label. For more information about explicit-null labels and ultimate-hop popping, see the *JUNOS MPLS Applications Configuration Guide*.

On M320 and T-series platforms, when you configure MPLS explicit-null labels with an MPLS EXP classifier, the MPLS EXP classifier can be different from an IPv4 or IPv6 classifier configured on the same logical interface. In other words, you can apply separate classifiers for MPLS EXP, IPv4, and IPv6 packets per logical interface. To combine an EXP classifier with a distinct IPv6 classifier, the PIC must be mounted on an Enhanced FPC.



NOTE: For J-series and other M-series platforms, MPLS explicit-null labels with MPLS EXP classification are supported if you set the same classifier for EXP and IPv4 traffic, or EXP and IPv6 traffic.

For more information about how IPv4 and IPv6 packet classification is handled, see “Applying DSCP IPv6 Classifiers” on page 52.

To configure an MPLS EXP classifiers for explicit-null labels, include the `exp` statement at the [edit class-of-service classifiers] and [edit class-of-service interfaces *interface-name* unit *logical-unit-number* classifiers] hierarchy levels:

```
[edit class-of-service classifiers]
exp classifier-name {
  import (classifier-name | default);
  forwarding-class class-name {
    loss-priority level {
      code-points [ aliases ] [ 6-bit-patterns ];
    }
  }
}
[edit class-of-service interfaces interface-name unit logical-unit-number classifiers]
exp (classifier-name | default);
```

Setting the PLP on T320 and M320 Platforms

By default, the least significant bit of the CoS value sets the packet loss priority (PLP) value. For example, CoS value 000 is associated with PLP low, and CoS value 001 is associated with PLP high. In general, you can change the PLP by configuring a BA or MF classifier, as discussed in “Classifier Types” on page 45.

However, on T-series and M320 platforms that do not have tricolor marking enabled, the loss priority can be configured only by setting the PLP within an MF classifier. This setting can then be used by the appropriate drop profile map and rewrite rule.

For T-series and M320 platforms with Enhanced II Flexible PIC Concentrators (FPCs) and tricolor marking enabled, you can set the PLP with a BA or MF classifier, as described in “Setting the PLP with a BA Classifier” on page 176 and “Setting the PLP with a Multifield Classifier” on page 177.

Example: Overriding the Default PLP on M320 Platforms

Override the default PLP.

1. The following example specifies that while the DSCP code points are 110, the loss priority is set to high; however, on M320 platforms, overriding the default PLP this way has no effect.

```
class-of-service {
  classifiers {
    dscp ba-classifier {
      forwarding-class expedited-forwarding {
        loss-priority high code-points 110;
      }
    }
  }
}
```

2. For M320 platforms, this MF classifier sets the PLP.

```
firewall {
  filter ef-filter {
    term ef-multifield {
      from {
        precedence 6;
      }
      then {
        loss-priority high;
        forwarding-class expedited-forwarding;
      }
    }
  }
}
```

Classifying Frame Relay Traffic

For J-series Services Router interfaces with Frame Relay encapsulation, you can set the loss priority of Frame Relay traffic, based on the discard eligibility (DE) bit. For each incoming frame with the DE bit containing the CoS value 0 or 1, you can configure a Frame Relay loss priority value of low, medium-low, medium-high, or high.

You can apply a classifier to the same interface on which you configure a Frame Relay loss priority value. The Frame Relay loss priority map is applied first, followed by the classifier. The classifier can change the loss priority to a higher value only (for example, from low to high). If the classifier specifies a loss priority with a lower value than the current loss priority of a particular packet, the classifier does not change the loss priority of that packet.

This section is organized as follows:

- Assigning the Default Frame Relay Loss Priority Map to an Interface on page 59
- Defining a Custom Frame Relay Loss Priority Map on page 59
- Verifying Your Configuration on page 60

Assigning the Default Frame Relay Loss Priority Map to an Interface

The default Frame Relay loss priority map contains the following settings:

```
loss-priority low code-point 0;
loss-priority high code-point 1;
```

This default map sets the loss priority to **low** for each incoming frame with the DE bit containing the 0 CoS value. The map sets the loss priority to **high** for each incoming frame with the DE bit containing the 1 CoS value.

To assign the default map to an interface, include the `frame-relay-de default` statement at the `[edit class-of-service interfaces interface-name unit logical-unit-number loss-priority-maps]` hierarchy level:

```
[edit class-of-service interfaces interface-name unit logical-unit-number
loss-priority-maps]
frame-relay-de default;
```

Defining a Custom Frame Relay Loss Priority Map

To define a custom Frame Relay loss priority map, include the following statements at the `[edit class-of-service]` hierarchy level:

```
[edit class-of-service]
loss-priority-maps {
  frame-relay-de map-name {
    loss-priority (low | medium-low | medium-high | high) code-point (0 | 1);
  }
}
```

A custom loss priority map sets the loss priority to **low**, **medium-low**, **medium-high**, or **high** for each incoming frame with the DE bit containing the specified 0 or 1 CoS value.

Applying the Map to a Logical Interface

The map does not take effect until you apply it to a logical interface. To apply a map to a logical interface, include the `frame-relay-de map-name` statement at the `[edit class-of-service interfaces interface-name unit logical-unit-number loss-priority-maps]` hierarchy level:

```
[edit class-of-service interfaces interface-name unit logical-unit-number
loss-priority-maps]
frame-relay-de map-name;
```

Verifying Your Configuration

To verify your configuration, you can issue the following operational mode commands:

- `show class-of-service forwarding-table loss-priority-map`
- `show class-of-service forwarding-table loss-priority-map mapping`
- `show chassis forwarding`
- `show pfe fwdd`