

Chapter 18

Configuring Tricolor Marking


Networks police traffic by limiting the input or output transmission rate of a class of traffic on the basis of user-defined criteria. Policing traffic allows you to control the maximum rate of traffic sent or received on an interface and to partition a network into multiple priority levels or classes of service. Tricolor marking (TCM) extends the functionality of class of service (CoS) traffic policing by providing three levels of drop precedence (loss priority or PLP) instead of two. With TCM, you can provision more granular service level agreements (SLAs) across the Differentiated Services (DiffServ) domain.

There are two types of TCM: single-rate and two-rate. The JUNOS software does not support single-rate tricolor marking. For T-series platforms with Enhanced II Flexible PIC Concentrators (FPCs), you can configure two-rate TCM, as defined in RFC 2698, *A Two Rate Three Color Marker*.

Two-rate TCM enables a “color-aware” method of traffic policing. The high, medium, and low loss priorities are mapped to the colors red, yellow, and green. The color of a packet, which is used or set by the TCM policer, corresponds to the packet’s loss priority. Packets with high PLP are marked red, packets with medium PLP are marked yellow, and packets with low PLP are marked green.

Two-rate TCM polices traffic according to the color classification (loss priority) of each packet. Traffic policing is based on two rates: the committed information rate (CIR) and the peak information rate (PIR). The CIR specifies the average rate at which packets are admitted to the network. As each packet enters the network, it is counted. Packets that do not exceed the CIR are marked green, which corresponds to **low** loss priority. Packets that exceed the CIR but are below the PIR are marked yellow, which corresponds to **medium** loss priority. Packets that exceed the PIR are marked red, which corresponds to **high** loss priority. For more information, see “How Two-Rate Tricolor Marking Works” on page 152.

The two-rate TCM policer provides two functions: metering and marking. The policer meters each packet and passes the packet and the metering result to the marker, as shown in Figure 6 on page 150.

Figure 6: Flow of Tricolor Marking Policer Operation

The meter operates in one of two modes. In the color-blind mode, the meter treats the packet stream as uncolored. Any preset loss priorities are ignored. In the color-aware mode, the meter inspects the packet's PLP field, which has been set by an upstream device as green, yellow, or red; the PLP field has already been set by a behavior aggregate (BA) or multifield (MF) classifier to low, medium, or high. The marker changes the color (PLP) of each incoming IP packet according to the results of the meter.

For information about how to use two-rate TCM with BA and MF classifiers, see “Configuring the PLP for a BA Classifier” on page 156 and “Configuring the PLP for a Multifield Classifier” on page 157.

You configure TCM by defining tricolor marking policers, and three levels of packet loss priority (PLP) for classifiers, rewrite rules, random early detection (RED) drop profiles, and firewall filters. To configure TCM, you can include the following statements at the [edit class-of-service] hierarchy level of the configuration:

```

[edit class-of-service]
tri-color;
classifiers {
  (dscp | dscp-ipv6 | exp | ieee-802.1 | inet-precedence) classifier-name {
    import (classifier-name | default);
    forwarding-class class-name {
      loss-priority (low | medium | high) {
        code-points [ aliases ] [ 6-bit-patterns ];
      }
    }
  }
}
rewrite-rules {
  (dscp | dscp-ipv6 | exp | ieee-802.1 | inet-precedence) rewrite-name {
    import (rewrite-name | default);
    forwarding-class class-name {
      loss-priority (low | medium | high) code-point (alias | bits);
    }
  }
}
schedulers {
  scheduler-name {
    drop-profile-map loss-priority (any | low | medium | high) protocol any
    drop-profile profile-name;
  }
}

```

```

[edit firewall]
policer name {
    then loss-priority (low | medium | high);
}
three-color-policer name {
    two-rate {
        (color-aware | color-blind);
        committed-information-rate bps;
        committed-burst-size bytes;
        peak-information-rate bps;
        peak-burst-size bytes;
    }
}
filter filter-name {
    <family family>{
        term rule-name {
            then {
                three-color-policer {
                    two-rate policer-name;
                }
            }
        }
    }
}

```

This chapter discusses the following topics:

- How Two-Rate Tricolor Marking Works on page 152
- Enabling Tricolor Marking on page 152
- Configuring a Tricolor Marking Policer on page 153
- Applying a Tricolor Marking Policer to a Firewall Filter on page 154
- Applying a Tricolor Marking Policer to an Interface on page 155
- Configuring the PLP for a BA Classifier on page 156
- Configuring the PLP for a Multifield Classifier on page 157
- Configuring the PLP for a Drop-Profile Map on page 158
- Configuring the PLP for a Rewrite Rule on page 159
- Verifying Your Configuration on page 159
- Example: Configuring Tricolor Marking on page 160

How Two-Rate Tricolor Marking Works

With two-rate TCM, you can enforce traffic policing according to two separate rates—the committed information rate (CIR) and the peak information rate (PIR). The CIR specifies the average rate at which packets are admitted to the network. The PIR is greater than or equal to the CIR. Packets that exceed the CIR, but are below the PIR are marked yellow. Packets that exceed the PIR are marked red, as shown in Table 22.

Table 22: Two-Rate TCM Color-to-PLP Mapping

Color	PLP	Assigned If...
Red	high	Packet exceeds the PIR
Yellow	medium	Packet exceeds the CIR
Green	low	Packet does not exceed the CIR

With colorblind policing, all packets are evaluated by the CIR. If a packet exceeds the CIR, it is evaluated by the PIR.

With color-aware policing, the metering treatment the packet receives depends on its classification. Metering can increase a packet's assigned PLP, but cannot decrease it, as described below:

- Packets belonging to the green class have already been marked by a classifier with low PLP. The marking policer can leave the packet's PLP unchanged or increase the PLP to medium or high. Therefore, these packets are metered against both the CIR and the PIR.
- Packets belonging to the yellow class have already been marked by a classifier with medium PLP. The marking policer can leave the packet's PLP unchanged or increase the PLP to high. Therefore, these packets are metered against the PIR only.
- Packets belonging to the red class have already been marked by a classifier with high PLP. The marking policer can only leave the packet's PLP unchanged. Therefore, these packets are not metered against the CIR or the PIR.

Enabling Tricolor Marking

By default, TCM is not enabled. To enable TCM, include the `tri-color` statement at the `[edit class-of-service]` hierarchy level:

```
[edit class-of-service]
tri-color;
```

If you do not include this statement in the configuration, you cannot configure medium PLP for classifiers, rewrite rules, drop profiles, or firewall filters.

Configuring a Tricolor Marking Policer

A tricolor marking policer polices traffic on the basis of metering rates, including the CIR, the PIR, their associated burst sizes, and any policing actions configured for the traffic. To configure a tricolor marking policer, include the following statements at the [edit firewall] hierarchy level:

```
[edit firewall]
three-color-policer name {
  two-rate {
    (color-aware | color-blind);
    committed-information-rate bps;
    committed-burst-size bytes;
    peak-information-rate bps;
    peak-burst-size bytes;
  }
}
```

You can specify the values for bps and bytes either as complete decimal numbers or as decimal numbers followed by the abbreviation k (1000), m (1,000,000), or g (1,000,000,000).

Table 23 describes the configurable statements.

Table 23: Tricolor Marking Policer Statements

Statement	Meaning	Configurable Values
two-rate	Marking is based on the CIR and the PIR.	N/A
color-aware	Metering depends on the packet's preclassification. Metering can increase a packet's assigned PLP, but cannot decrease it. For more information, see "How Two-Rate Tricolor Marking Works" on page 152.	
color-blind	All packets are evaluated by the CIR. If a packet exceeds the CIR, it is evaluated by the PIR.	
committed-information-rate	Guaranteed bandwidth under normal line conditions and the average rate up to which packets are marked green.	32,000 through 40,000,000,000 bps
committed-burst-size	Maximum number of bytes allowed for incoming packets to burst above the CIR, but still be marked green.	1500 through 100,000,000,000 bytes
peak-information-rate	Maximum achievable rate. Packets that exceed the CIR but are below the PIR are marked yellow. Packets that exceed the PIR are marked red.	32,000 through 40,000,000,000 bps
peak-burst-size	Maximum number of bytes allowed for incoming packets to burst above the PIR, but still be marked yellow.	1500 through 100,000,000,000 bytes

Applying a Tricolor Marking Policer to a Firewall Filter

To rate-limit traffic by attaching a tricolor marking policer to a firewall filter, include the `three-color-policer` statement:

```
three-color-policer two-rate two-rate-policer-name;
```

You can include this statement at the following hierarchy levels:

- [edit firewall family *family* filter *filter-name* term *rule-name* then]
- [edit firewall filter *filter-name* term *rule-name* then]

In the family statement, the protocol family can be any, `ccc`, `inet`, `inet6`, `mpls`, or `vpls`.

Example: Applying a Tricolor Marking Policer to a Firewall Filter

Apply the `trtc1` policer to a firewall filter:

```
firewall {
  three-color-policer trtc1 { # Configure the trtc1 policer.
    two-rate {
      color-blind;
      committed-information-rate 1048576;
      committed-burst-size 65536;
      peak-information-rate 10485760;
      peak-burst-size 131072;
    }
  }
  filter fil { # Configure the fil firewall filter, which attaches the trtc1 policer.
    term default {
      then {
        three-color-policer {
          two-rate trtc1;
        }
      }
    }
  }
}
```

For more information about applying policers to firewall filters, see the *JUNOS Policy Framework Configuration Guide*.

Applying a Tricolor Marking Policer to an Interface

To apply a tricolor marking policer to an interface, you must reference the filter name in the interface configuration. To do this, include the `filter` statement:

```
filter {
  input filter-name;
  output filter-name;
}
```

You can include these statements at the following hierarchy levels:

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

The filter name that you reference should have an attached tricolor marking policer, as shown in “Applying a Tricolor Marking Policer to a Firewall Filter” on page 154.

Example: Applying a Tricolor Marking Policer to an Interface

Apply the `trtcm1` policer to an interface:

```
firewall {
  three-color-policer trtcm1 { # Configure the trtcm1 policer.
    two-rate {
      color-blind;
      committed-information-rate 1048576;
      committed-burst-size 65536;
      peak-information-rate 10485760;
      peak-burst-size 131072;
    }
  }
  filter fil { # Configure the fil firewall filter, which attaches the trtcm1 policer.
    term default {
      then {
        three-color-policer {
          two-rate trtcm1;
        }
      }
    }
  }
}
interfaces { # Configure the interface, which attaches the fil firewall filter.
  so-1/0/0 {
    unit 0 {
      family inet {
        filter {
          input fil;
        }
      }
    }
  }
}
```

Configuring the PLP for a BA Classifier

Behavior aggregate (BA) classifiers take action on incoming packets. When TCM is enabled, T-series platforms support three classifier PLP designations: low, medium, and high.

To configure the PLP for a classifier, include the following statements 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 (low | medium | high) {
        code-points [ aliases ] [ 6-bit-patterns ];
      }
    }
  }
}
```

The inputs for a classifier are the code points. The outputs for a classifier are the forwarding class and the loss priority (PLP). A classifier sets the forwarding class and the PLP for each packet entering the interface with a specific set of code points.

For example, in the following configuration, the **assured-forwarding** forwarding class and **medium** PLP are assigned to all packets entering the interface with the **101110** code points:

```
class-of-service {
  classifiers {
    dscp dscp-cl {
      forwarding-class assured-forwarding{
        loss-priority medium {
          code-points 101110;
        }
      }
    }
  }
}
```

To use this classifier, you must configure the settings for the **assured-forwarding** forwarding class at the [edit class-of-service forwarding-classes queue queue-number assured-forwarding] hierarchy level. For more information, see “Configuring Forwarding Classes” on page 33.

Configuring the PLP for a Multifield Classifier

Multifield classifiers take action on incoming or outgoing packets, depending whether the firewall rule is applied as an input filter or an output filter. When TCM is enabled, T-series platforms support three multifield classifier PLP designations: low, medium, and high.

To configure the PLP for a multifield classifier, include the `loss-priority` statement in a policer or firewall filter that you configure at the at the `[edit firewall]` hierarchy level:

```
[edit firewall]
family family-name {
  filter filter-name {
    term term-name {
      from {
        match-conditions;
      }
      then {
        loss-priority (low | medium | high);
        forwarding-class class-name;
      }
    }
  }
}
```

The inputs (match conditions) for a multifield classifier are one or more of the six packet header fields: destination address, source address, IP protocol, source port, destination port, and DSCP. The outputs for a multifield classifier are the forwarding class and the loss priority (PLP). In other words, a multifield classifier sets the forwarding class and the PLP for each packet entering or exiting the interface with a specific destination address, source address, IP protocol, source port, destination port, or DSCP.

For example, in the following configuration, the forwarding class `expedited-forwarding` and PLP `medium` are assigned to all IPv4 packets with the `10.1.1.0/24` or `10.1.2.0/24` source address:

```
firewall {
  family inet {
    filter classify-customers {
      term isp1-customers {
        from {
          source-address 10.1.1.0/24;
          source-address 10.1.2.0/24;
        }
        then {
          loss-priority low;
          forwarding-class expedited-forwarding;
        }
      }
    }
  }
}
```

To use this classifier, you must configure the settings for the `expedited-forwarding` forwarding class at the `[edit class-of-service forwarding-classes queue queue-number expedited-forwarding]` hierarchy level. For more information, see “Configuring Forwarding Classes” on page 33.

Configuring the PLP for a Drop-Profile Map

RED drop profiles take action on outgoing packets. When TCM is enabled, T-series platforms support three drop-profile map PLP designations: low, medium, and high.

To configure the PLP for the drop-profile map, include the following statements at the `[edit class-of-service]` hierarchy level:

```
[edit class-of-service]
schedulers {
  scheduler-name {
    drop-profile-map loss-priority (any | low | medium | high) protocol any
    drop-profile profile-name;
  }
}
```

When you configure TCM, the drop-profile map’s protocol type must be `any`.

The inputs for a drop-profile map are the loss priority and the protocol type. The output for a drop-profile map is the drop profile name. In other words, the map sets the drop profile for each packet with a specific PLP and protocol type exiting the interface.

For example, in the following configuration, the `dp` drop profile is assigned to all packets exiting the interface with a medium PLP and belonging to any protocol:

```
class-of-service {
  schedulers {
    af {
      drop-profile-map loss-priority medium protocol any drop-profile dp;
    }
  }
}
```

To use this drop-profile map, you must configure the settings for the `dp` drop profile at the `[edit class-of-service drop-profiles dp]` hierarchy level. For more information, see “Configuring RED Drop Profiles” on page 85.

Configuring the PLP for a Rewrite Rule

Rewrite rules take action on outgoing packets. When TCM is enabled, T-series platforms support three rewrite PLP designations: low, medium, and high.

To configure the PLP for a rewrite rule, include the following statements at the [edit class-of-service] hierarchy level:

```
[edit class-of-service]
rewrite-rules {
  (dscp | dscp-ipv6 | exp | ieee-802.1 | inet-precedence) rewrite-name {
    import (rewrite-name | default);
    forwarding-class class-name {
      loss-priority (low | medium | high) code-point (alias | bits);
    }
  }
}
```

The inputs for a rewrite rule are the forwarding class and the loss priority (PLP). The output for a rewrite rule are the code points. In other words, a rewrite rule sets the code points for each packet exiting the interface with a specific forwarding class and PLP.

For example, if you configure the following, the 000000 code points are assigned to all packets exiting the interface with the `assured-forwarding` forwarding class and medium PLP:

```
class-of-service {
  rewrite-rules {
    dscp dscp-rw {
      forwarding-class assured-forwarding {
        loss-priority medium code-point 000000;
      }
    }
  }
}
```

To use this classifier, you must configure the settings for the `assured-forwarding` forwarding class at the [edit class-of-service forwarding-classes queue *queue-number* assured-forwarding] hierarchy level. For more information, see “Configuring Forwarding Classes” on page 33.

Verifying Your Configuration

The following operational mode commands are useful for checking the results of your configuration:

- `show class-of-service`
- `show interfaces interface-name extensive`

For information about these commands, see the *JUNOS Interfaces Command Reference* and *JUNOS System Basics and Services Command Reference*.

Example: Configuring Tricolor Marking

Two-Rate Tricolor Policer's Effect on Traffic

Configure two SONET/SDH interfaces and apply a two-rate tricolor policer. The objective is to have packets with three different code points corresponding to the green, yellow, and red markings (low, medium, and high loss priorities) of the policer.

```

class-of-service {
  tri-color;
  drop-profiles {
    dp-low {
      interpolate {
        fill-level 95;
        drop-probability 100;
      }
    }
    dp-high {
      interpolate {
        fill-level 20;
        drop-probability 100;
      }
    }
    dp-medium {
      interpolate {
        fill-level 50;
        drop-probability 100;
      }
    }
  }
}
scheduler-maps {
  cp-plp {
    forwarding-class assured-forwarding scheduler af;
    forwarding-class expedited-forwarding scheduler ef;
    forwarding-class best-effort scheduler be;
  }
}
schedulers {
  af {
    transmit-rate percent 10;
    buffer-size percent 10;
    drop-profile-map loss-priority high protocol any drop-profile dp-high;
    drop-profile-map loss-priority low protocol any drop-profile dp-low;
    drop-profile-map loss-priority medium protocol any drop-profile dp-medium;
  }
  ef {
    transmit-rate percent 30;
    buffer-size percent 30;
    drop-profile-map loss-priority low protocol any drop-profile dp-low;
    drop-profile-map loss-priority high protocol any drop-profile dp-high;
  }
  be {
    transmit-rate remainder;
    buffer-size remainder;
  }
}

```

```

interfaces {
  so-3/2/0 {
    unit 0 {
      rewrite-rules {
        exp default;
        inet-precedence 8q;
      }
    }
  }
  so-3/0/0 {
    unit 0 {
      classifiers {
        exp default;
        inet-precedence 8q;
      }
    }
  }
  so-4/0/0 {
    scheduler-map cp-plp;
    unit 0 {
      rewrite-rules {
        exp default;
        inet-precedence 8q;
      }
    }
  }
  so-4/2/0 {
    scheduler-map cp-plp;
    unit 0 {
      rewrite-rules {
        exp default;
        inet-precedence 8q;
      }
    }
  }
  so-1/* {
    scheduler-map sched;
    unit 0 {
      classifiers {
        inet-precedence color-cl;
      }
      rewrite-rules {
        inet-precedence color-rw;
      }
    }
  }
}

```

```

interfaces {
  so-1/0/0 {
    no-keepalives;
    mtu 4474;
    sonet-options {
      fcs 32;
    }
    unit 0 {
      family inet {
        filter {
          input fil;
        }
        address 192.7.1.1/32 {
          destination 192.7.1.2;
        }
      }
    }
  }
  so-1/1/0 {
    no-keepalives;
    sonet-options {
      fcs 32;
    }
    unit 0 {
      family inet {
        filter {
          input fil;
        }
        address 192.2.1.1/32 {
          destination 192.2.1.2;
        }
      }
    }
  }
  so-3/0/0 {
    no-keepalives;
    mtu 4474;
    encapsulation ppp;
    sonet-options {
      fcs 32;
    }
    unit 0 {
      family inet {
        filter {
          input out-plp;
        }
      }
    }
  }
}

```

```

firewall {
  three-color-policer trtcm1 {
    two-rate {
      color-blind;
      committed-information-rate 1048576;
      committed-burst-size 65536;
      peak-information-rate 10485760;
      peak-burst-size 131072;
    }
  }
  filter fil {
    term default {
      then {
        three-color-policer {
          two-rate trtcm1;
        }
      }
    }
  }
}

class-of-service {
  rewrite-rules {
    dscp plp {
      forwarding-class assured-forwarding {
        loss-priority low code-point 010110;
        loss-priority high code-point 010101;
      }
      forwarding-class expedited-forwarding {
        loss-priority high code-point 110101;
        loss-priority low code-point 110100;
      }
    }
  }
  inet-precedence 8q {
    forwarding-class best-effort {
      loss-priority low code-point 010;
      loss-priority high code-point 001;
    }
    forwarding-class expedited-forwarding {
      loss-priority low code-point 100;
      loss-priority high code-point 110;
    }
    forwarding-class assured-forwarding {
      loss-priority low code-point 010;
      loss-priority high code-point 110;
    }
    forwarding-class network-control {
      loss-priority low code-point 110;
      loss-priority high code-point 111;
    }
  }
}

```

```

inet-precedence color-rw {
  forwarding-class best-effort {
    loss-priority low code-point 010;
    loss-priority high code-point 001;
  }
  forwarding-class expedited-forwarding {
    loss-priority low code-point 100;
    loss-priority high code-point 110;
    loss-priority medium code-point 101;
  }
}
}
classifiers {
  dscp plp {
    forwarding-class expedited-forwarding {
      loss-priority low code-points 110010;
      loss-priority high code-points 110011;
    }
    forwarding-class assured-forwarding {
      loss-priority low code-points 101000;
      loss-priority high code-points 101001;
    }
  }
  dscp-ipv6 plp {
    forwarding-class expedited-forwarding {
      loss-priority low code-points 110010;
      loss-priority high code-points 100110;
    }
  }
  inet-precedence 8q {
    forwarding-class best-effort {
      loss-priority low code-points 001;
    }
    forwarding-class expedited-forwarding {
      loss-priority high code-points 100;
      loss-priority low code-points 110;
    }
    forwarding-class assured-forwarding {
      loss-priority high code-points 011;
      loss-priority low code-points 010;
    }
  }
  inet-precedence color-cl {
    forwarding-class expedited-forwarding {
      loss-priority high code-points 110;
      loss-priority low code-points 100;
    }
    forwarding-class assured-forwarding {
      loss-priority medium code-points 101;
    }
  }
}
}
}

```