Configuring Gigabit Ethernet Policers

On Gigabit Ethernet IQ and Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i platform), you can define rate limits for premium and aggregate traffic received on the interface. These policers allow you to perform simple traffic policing without configuring a firewall filter. First you configure the Ethernet policer profile, next you classify ingress and egress traffic, then you can apply the policer to a logical interface.

For Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i platform), the policer rates you configure can be different than the rates on the Packet Forward Engine. The difference results from Layer 2 overhead. The PIC accounts for this difference.

This section contains the following topics:

Configuring a Policer

To configure an Ethernet policer profile, include the ethernet-policer-profile statement at the [edit interfaces *interface-name* gigether-options ethernet-switch-profile] hierarchy level:

```
[edit interfaces interface-name gigether-options ethernet-switch-profile]
ethernet-policer-profile {
    [Unresolved xref] cos-policer-name {
        [Unresolved xref] {
            bandwidth-limit (ethernet) bps;
            burst-size-limit (ethernet) bytes;
        }
        [Unresolved xref] {
            bandwidth-limit (ethernet) bps;
            burst-size-limit (ethernet) bytes;
        }
    }
}
```

In the Ethernet policer profile, the aggregate-priority policer is mandatory; the premium-priority policer is optional.

For aggregate and premium policers, you specify the bandwidth limit in bits per second. You can specify the value as a complete decimal number or as a decimal number followed by the abbreviation k (1000), m (1,000,000), or g (1,000,000,000).

There is no absolute minimum value for bandwidth limit, but any value below 61,040 bps will result in an effective rate of 30,520 bps. The maximum bandwidth limit is 4.29 Gbps.

The maximum burst size controls the amount of traffic bursting allowed. To determine the burst-size limit, you can multiply the bandwidth of the interface on which you are applying the filter by the amount of time you allow a burst of traffic at that bandwidth to occur:

burst size = bandwidth x *allowable time for burst traffic*

If you do not know the interface bandwidth, you can multiply the maximum MTU of the traffic on the interface by 10 to obtain a value. For example, the burst size for an MTU of 4700 would be 47,000 bytes. The burst size should be at least 10 interface MTUs. The maximum value for the burst-size limit is 100 MB.

Specifying an Input Priority Map

An input priority map identifies ingress traffic with specified IEEE 802.1p priority values, and classifies that traffic as premium.

If you include a premium-priority policer, you can specify an input priority map by including the **ieee802.1 premium** statement at the **[edit interfaces** *interface-name* **gigether-options ethernet-policer-profile input-priority-map]** hierarchy level:

```
[edit interfaces interface-name gigether-options ethernet-policer-profile input-priority-map]
ieee802.1p premium [ values ];
```

The priority values can be from 0 through 7. The remaining traffic is classified as nonpremium (or aggregate). For a configuration example, see “Example: Configuring Gigabit Ethernet Policers” on page 5.

NOTE: On IQ2 and IQ2-E interfaces and MX-series interfaces, when a VLAN tag is pushed, the inner VLAN IEEE 802.1p bits are copied to the IEEE bits of the VLAN or VLANs being pushed. If the original packet is untagged, the IEEE bits of the VLAN or VLANs being pushed are set to 0.

Specifying an Output Priority Map

An output priority map identifies egress traffic with specified queue classification and packet loss priority (PLP), and classifies that traffic as premium.

If you include a premium-priority policer, you can specify an output priority map by including the **classifier** statement at the **[edit interfaces** *interface-name* **gigether-options ethernet-policer-profile output-priority-map]** hierarchy level:

```
[edit interfaces interface-name gigether-options ethernet-policer-profile
   output-priority-map]
classifier {
   [Unresolved xref] {
      [Unresolved xref] class-name {
```

```
            loss-priority (high | low);
        }
    }
}
```

You can define a forwarding class, or you can use a predefined forwarding class. Table 1 shows the predefined forwarding classes and their associated queue assignments.

Table 1: Default Forwarding Classes

Forwarding Class Name	Queue
best-effort	Queue 0
expedited-forwarding	Queue 1
assured-forwarding	Queue 2
network-control	Queue 3

For more information about CoS forwarding classes, see the *JUNOS Class of Service Configuration Guide*. For a configuration example, see "Example: Configuring Gigabit Ethernet Policers" on page 5.

Applying a Policer

On Gigabit Ethernet IQ and Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i platform), you can apply input and output policers that define rate limits for premium and aggregate traffic received on the logical interface. Aggregate policers are supported on Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i platform).

These policers allow you to perform simple traffic policing without configuring a firewall filter. For information about defining these policers, see "Configuring Gigabit Ethernet Policers" on page 1.

To apply policers to specific source MAC addresses, include the accept-source-mac statement:

```
accept-source-mac {
    mac-address mac-address {
        [Unresolved xref] {
            input cos-policer-name;
            output cos-policer-name;
        }
    }
}
```

You can include these statements at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

You can specify the MAC address as *nn*:*nn*:*nn*:*nn*:*nn*:*nn* or *nnnn*.*nnnn*.*nnnn*, where *n* is a hexadecimal number. You can configure up to 64 source addresses. To specify more than one address, include multiple mac-address statements in the logical interface configuration.

NOTE: On untagged Gigabit Ethernet interfaces you should not configure the source-address-filter statement at the [edit interfaces *ge-fpc/pic/port* gigether-options] hierarchy level and the accept-source-mac statement at the [edit interfaces *ge-fpc/pic/port* gigether-options unit *logical-unit-number*] hierarchy level simultaneously. If these statements are configured for the same interfaces at the same time, an error message is displayed.

On tagged Gigabit Ethernet interfaces you should not configure the source-address-filter statement at the [edit interfaces *ge-fpc/pic/port* gigether-options] hierarchy level and the accept-source-mac statement at the [edit interfaces *ge-fpc/pic/port* gigether-options unit *logical-unit-number*] hierarchy level with an identical MAC address specified in both filters. If these statements are configured for the same interfaces with an identical MAC address specified, an error message is displayed.

NOTE: If the remote Ethernet card is changed, the interface does not accept traffic from the new card because the new card has a different MAC address.

The MAC addresses you include in the configuration are entered into the routing platform's MAC database. To view the routing platform's MAC database, enter the show interfaces mac-database *interface-name* command:

```
user@host> show interfaces mac-database interface-name
```

In the input statement, list the name of one policer template to be evaluated when packets are received on the interface.

In the output statement, list the name of one policer template to be evaluated when packets are transmitted on the interface.

NOTE: On IQ2 and IQ2-E PIC interfaces, the default value for maximum retention of entries in the MAC address table has changed, for cases in which the table is not full. The new holding time is 12 hours. The previous retention time of 3 minutes is still in effect when the table is full.

You can use the same policer one or more times.

If you apply both policers and firewall filters to an interface, input policers are evaluated before input firewall filters, and output policers are evaluated after output firewall filters.

Configuring MAC Address Filtering

You cannot explicitly define traffic with specific source MAC addresses to be rejected; however, for Gigabit Ethernet IQ and Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i platform), you can block all incoming packets that do not have a source address specified in the accept-source-mac statement. For more information about the accept-source-mac statement, see "Applying a Policer" on page 3.

To enable this blocking, include the source-filtering statement at the [edit interfaces *interface-name* gigether-options] hierarchy level:

```
[edit interfaces interface-name gigether-options]
source-filtering;
```

For more information about the source-filtering statement, see Enabling Ethernet MAC Address Filtering.

To accept traffic even though it does not have a source address specified in the accept-source-mac statement, include the no-source-filtering statement at the [edit interfaces *interface-name* gigether-options] hierarchy level:

```
[edit interfaces interface-name gigether-options]
no-source-filtering;
```

For more information about the accept-source-mac statement, see "Applying a Policer" on page 3.

Example: Configuring Gigabit Ethernet Policers

Configure interface ge-6/0/0 to treat priority values 2 and 3 as premium. On ingress, this means that IEEE 802.1p priority values 2 and 3 are treated as premium. On egress, it means traffic that is classified into queue 0 or 1 with PLP of low and queue 2 or 3 with PLP of high, is treated as premium.

Define a policer that limits the premium bandwidth to 100 Mbps and burst size to 3 k, and the aggregate bandwidth to 200 Mbps and burst size to 3 k.

Specify that frames received from the MAC address 00:01:02:03:04:05 and the VLAN ID 600 are subject to the policer on input and output. On input, this means frames received with the source MAC address 00:01:02:03:04:05 and the VLAN ID 600 are subject to the policer. On output, this means frames transmitted from the routing platform with the destination MAC address 00:01:02:03:04:05 and the VLAN ID 600 are subject to the policer.

```
[edit interfaces]
ge-6/0/0 {
   gigether-options {
      ether-switch-profile {
```

```
        ether-policer-profile {
            input-priority-map {
                ieee-802.1p {
                    premium [ 2 3 ];
                }
            }
            output-priority-map {
                classifier {
                    premium {
                        forwarding-class best-effort {
                            loss-priority low;
                        }
                        forwarding-class expedited-forwarding {
                            loss-priority low;
                        }
                        forwarding-class assured-forwarding {
                            loss-priority high;
                        }
                        forwarding-class network-control {
                            loss-priority high;
                        }
                    }
                }
            }
            policer policer-1 {
                premium {
                    bandwidth-limit 100m;
                    burst-size-limit 3k;
                }
                aggregate {
                    bandwidth-limit 200m;
                    burst-size-limit 3k;
                }
            }
        }
    }
    unit 0 {
        accept-source-mac {
            mac-address 00:01:02:03:04:05 {
                policer {
                    input policer-1;
                    output policer-1;
                }
            }
        }
    }
}
```